



OPERATING SUPPLY CHAINS IN AN ERA OF UNCERTAINTY

ABOUT THE REPORT	2
EXECUTIVE SUMMARY	3
SCENE SETTING	4
PRINCIPLES FOR SUPPLY CHAINS OPERATING IN UNCERTAIN TIMES	9
1. Demand-Driven, Agile & Responsive Operations	11
2. Digital Infrastructure, Intelligence & Governance	16
3. Transparency, Traceability & Compliance	19
4. Circularity & Resource Stewardship	22
5. Fair Value Distribution	25
6. Climate Adaptation & Mitigation	28
7. Ecosystem Building	31
The Principles as a System	34
INTERDEPENDENCIES BETWEEN THE PRINCIPLES	35
THE OPERATING DECISIONS LEADERS MUST OPTIMISE	37
CONCRETE APPLICATION OF THE PRINCIPLES	44
APPENDIX	58

ABOUT THE REPORT

Authors and project team

This report was led by Hyran, with analytical support from Swanstant and with support from SHEIN. The core research team comprised Ahmed Zaidi, François Souchet and Ankit Tejpal, with the lead authors being Ahmed Zaidi and François Souchet.

Disclosure and independence

Editorial control rested with Hyran and Swanstant throughout. Hyran designed and managed the research, and the framing, analysis, conclusions and recommendations—together with the trade-offs, limitations and critical observations the report sets out—are its own. Alongside material made available by SHEIN, Hyran drew on its own expertise and prior work, independent data analysis, and a programme of interviews with external experts engaged solely by Hyran. Where the evidence raised tensions or left questions open, the report says so.

SHEIN contributed funding to the research and provided the Hyran team access to its internal supply chain and sustainability functions as well as its external manufacturing partners whose operations, data and first-hand accounts form part of the evidence base. That access allowed the analysis to draw on primary material, including site visits and interviews. For transparency, no company financial data were furnished for inclusion in the study.

Acknowledgements

We are grateful to the external experts who gave their time and insight, including:

- Anant Ahuja - Director of Environment, Social and Governance (ESG) and Sustainability, Shahi Export
- Azfar Hasan - CEO of Matrix Sourcing and IAF Business Innovation Committee
- Brad Ballentine - CEO, MAS ACME USA

- Dr. Christoph Wolff - CEO of Smart Freight Centre and Honorary Professor Economics & Social Sciences, University of Cologne
- Lutz Water - Secretary General, Textile ETP (European Technology Platform for the Future of Textiles and Clothing)
- Roland Van Os - Former Chief Commercial Officer (CCO), Otto Group

Their perspectives helped shape and stress-test the analysis; their participation does not imply endorsement of its conclusions. Responsibility for the final text rests with Hyran alone. We also thank the manufacturing and supplier partners who welcomed the research team for site visits and spoke with us candidly, and the teams at SHEIN who facilitated access and answered our questions throughout.

About the organisations



Hyran is the developer of a fashion technology platform designed to optimise textile and brand management activities. The company's AI-based platform optimises product development and manufacturing while reducing excess production, textile waste, and inventory planning, bringing collaboration between factories, suppliers, and brands with value chain alignment and shorter design-to-shelf cycles, enabling brands and suppliers to respond to consumer demand, increasing profitability, and reducing waste.



Swanstant is an economic forecasting firm that prices the climate signal. It builds interpretable models, combining econometrics and machine learning, to estimate how climate affects economic outcomes over a one- to six-month horizon. Its current work spans agricultural yields and industrial activity, tracing the channels through which those impacts reach costs, margins, and earnings.

EXECUTIVE SUMMARY

The problem: growth without profit.

Across the past decade, most segments of the fashion industry have grown revenue, while profits did not follow. Our benchmark of 47 listed companies over that period shows that:

1. Premium market's operating margins halved even as revenue rose by nearly a third.
2. Sportswear margins fell despite a doubling of revenue.
3. Luxury margins slipped from a 27% peak in 2022 to 24% in 2025.

An analysis of 36,036 investor questions across 1,000 earnings calls locates the cause between inventory overhang, gross-margin compression, and markdown dependency, which together account for over 40% of all problem-related questions put to management every year since 2018. These are signals that a forecast-driven, long-lead-time, large-batch operating model is increasingly eroding value.

Why now: four forces compound.

Four forces are converging to upend the current forecast-driven model:

1. Demand volatility has compressed trend cycles from years to weeks.
2. Trade disruption has made landed cost a moving target, with China-specific apparel surcharges peaking at 145% and investor concern about geographic concentration rising to roughly seven times its prior-decade average in 2025.
3. Regulation is tightening: the EU's ban on destroying unsold apparel took effect in July 2026, the Digital Product Passport lands from 2028 to 2029, and California's textile Extended Producer Responsibility (EPR) scheme is phasing in.
4. Climate disruption is affecting raw-material prices and manufacturing hubs.

Agility is often mistaken for producing faster. We make the case that it is better understood as a different operating infrastructure and a different way of managing relationships and decisions that go alongside it.

This paper's core argument: agility is an infrastructure and a mindset, not a speed setting.

We map three changes that sit at the centre, on which companies can start tomorrow:

1. **Live and two-way data flows sit at the centre.** Demand signals, sell-through velocity and reorder probability need to flow down to the mill and the factory. Capacity and material positions need to flow up to the brand. Making this work requires standardisation: common data definitions, standardised order fields, real-time visibility and the governance orchestrating who sees what and who decides what. Most brands do not start from nothing; the work already done to achieve traceability requirements provides foundations to build on.
2. **Suppliers are treated as trusted partners and real decisions are delegated to them.** To make agility work, decisions need to be delegated to suppliers: pre-positioning fabric at the right stage of completion, adjusting capacity, triggering replenishment against demand signals. Making this work requires three conditions: payment terms that align with fast inventory turnaround, the capability to convert a demand signal into action, and technology built as a functional asset for the supplier. Most brands already hold direct supplier relationships; the change is to evolve those by delegating more than volumes.
3. **Businesses optimise for total system margins rather than unit cost.** An initial run of 80–100 units can look irrational through a unit-cost lens; however, once markdowns, unsold stock, and disruption risk are counted, it can look very rational. Making this work requires redefining what buyers are measured on: full-price sell-through and inventory days of supply need to be integrated alongside cost per unit.

The framework: Seven interdependent principles.

This paper sets out seven interdependent principles for operating under uncertainty: demand-driven and agile operations; digital infrastructure and governance; transparency, traceability and compliance; circularity and resource stewardship; fair value distribution; climate adaptation and mitigation; and ecosystem building. These function as a system: digital infrastructure enables agility; transparency makes climate risk visible and circularity possible; fair value distribution is the financial condition under which suppliers can build the capabilities the others depend on.

While implementing those solutions, leaders will need to navigate key decisions.

Agility is not a blanket switch, and the report is explicit about where it helps and where it costs.

- > **Not all uncertainty is equal.** The demand-driven model delivers most for high-uncertainty SKUs, while predictable basics can remain efficient on a traditional forecast-driven model. The mature version is not a fixed classification but a discipline: defer each commitment to the latest responsible moment and size it to the strength of the signal.
- > **Transport emissions need to be optimised.** Agility can result in higher transport emissions, as it can increase reliance on air freight. A balanced model is broadly carbon-comparable to a planned one, with its advantages in margin, working capital and waste rather than emissions. A clear carbon benefit is available, but it relies on agility being delivered through multi-modal transport systems minimising air distances, production located closer to retail markets, volume consolidation and route optimisation, and sustainable aviation fuels.
- > **Supplier capital is a real constraint.** Agility can shift working-capital burden onto Tier 2 by requiring more fabric to be pre-positioned. This burden scales with fabric value, potentially creating barriers to premium and luxury. Several approaches lower the barrier: greige fabric postponement, fabric portfolio rationalisation, forward-looking demand data sharing with mills, Tier 2 financing mechanisms, and risk-sharing through rolling purchase orders and shared markdown liability.

The evidence: Three boardroom questions.

The case is organised around three starting points, corresponding to three questions encountered in investor or regulatory conversations, each illustrated by a practical case study:

- > **Inventory and working capital.** A demand-driven architecture with initial orders around 100 units and reorders triggered by confirmed demand can produce roughly ten times less unsold inventory than the reported industry average (around 2%, against an industry norm of 25–40% more produced than sold), generating associated benefits in immobilised capital, margin erosion and inventory costs.
- > **Margin and markdown.** Markdown dependency can erode 10–15% of potential revenue. Agility can significantly cut markdown or clearance reliance. A few companies are already illustrating this with direct cash-flow benefits.
- > **Circularity.** In the EU, the Ecodesign for Sustainable Products Regulation (ESPR) bans the destruction of unsold apparel from July 2026, the Digital Product Passport is coming, and California's textile EPR scheme starts mid-2026. Traditionally, those regulations represent a financial and compliance burden for companies. The capabilities agility requires (data integration, multi-tier traceability, and a tight link between design and manufacturing) are the same ones these regulations demand, turning a compliance mandate into a performance opportunity.

What the report offers.

Grounded in primary interviews with experts, operators, and suppliers running these models at scale, combined with a decade of investor-question data and a financial benchmark of listed companies, the report aims to be a working guide for senior leaders. It does not argue that agility is a universal answer, but that the infrastructure and management changes it depends on, like live two-way data, deep supplier partnership, and metrics that reward system margin, are available to start building now.



SCENE SETTING

Over the past decade, most segments of the fashion industry have grown their top line while failing to convert that growth into proportional profit.

Our analysis of financial performance across 47 listed fashion and apparel companies reveals a persistent and widening gap between revenue growth and operating profitability.

The pattern is stark. The premium segment has seen its operating margin halve despite growing revenue by nearly a third. Sportswear doubled its revenue but saw margins decline. Only luxury has managed to grow operating profit faster than revenue, but even there, our analysis shows that margins have contracted from their 2022 peak of 27% to 24% in 2025¹ with further pressure visible in the most recent results.

Financial markets have identified probable causes. An analysis of 36,036 questions asked by investors during 1,000 fashion and luxury earnings calls over the past decade reveals that three direct symptoms of this operating model failure—inventory overhang, gross margin compression, and markdown dependency—together account for over 40% of all problem-related questions put to management teams, consistently, every year since 2018². Inventory overhang alone roughly doubled to 192 mentions per 1,000 as an investor concern between 2014 and its post-pandemic peak in 2022–23, and remains the single most-discussed operational problem in the dataset. These are structural signals that the prevailing forecast-driven, long-lead-time, large-batch production operating model is systematically destroying a growing share of the value that revenue growth creates.

Segment	Revenue growth (2014–2025)	Operating profit growth	Op. margin 2014	Op. margin 2025	Margin change
Premium	↗ 30%	↘ 41%	13.2%	6.1%	↘ 7.2 pp
Sportswear	↗ 103%	↗ 81%	11.1%	9.9%	↘ 1.2 pp
Mass	↗ 71%	↗ 55%	13.1%	11.9%	↘ 1.2 pp
Luxury	↗ 140%	↗ 194%	19.6%	24.0%	↗ 4.4 pp

Until the late 1700s, fashion was made to measure³. Customers walked into a store, got measured, and a one-of-one garment was made just for them. By the early 1800s, the emergence of mass production in fashion was enabled not primarily by mechanisation alone, but by size standardisation, first developed for military uniforms during the Napoleonic Wars, which allowed garments to be produced in batches and stocked in standard sizes⁴.

However, even in the 1800s, mass production posed challenges that remain familiar today. Producing garments before knowing exact demand meant that retailers carried inventory risk.

If products did not sell, the margins on the garments were not sufficient to allow for deep discounting in order to clear stock. This resulted in pressure from retailers to reduce the cost of production. In the East End of London, garment workers organised strikes, most notably the 1889 East London Tailors' Strike involving 10,000 workers, as factory owners facing retail price pressure sought to reduce costs through the 'sweating system' of long hours and low piecework wages⁵. These cost challenges and structural inventory overhang didn't go away; they were simply exported overseas after World War II, when the move towards globalisation and free market systems accelerated⁶.

¹ Analysis based on public financial disclosure

² Authors' proprietary analysis of 1,000 fashion earnings calls, 2014–2025, details in Appendix A

³ Miles Lambert, "The Work of the Tailor in Eighteenth-Century Britain," *Costume* 44, no. 1 (2010), <https://doi.org/10.1179/174963010x12662396505761>.

⁴ Fit Analytics, "Vanity Sizing and Size Charts: A Brief History," 2024.

⁵ Oxford University Press, "Investigating Work in late Nineteenth-century London," *History Workshop Journal*, 2011 (on the House of Lords inquiry into the Sweating System).

⁶ Abernathy, F.H. et al., *A Stitch in Time: Lean Retailing and the Transformation of Manufacturing*, Harvard Business Press, 1999

The root cause is a model that is no longer fit to address the current challenges facing brands. For decades, fashion has operated on forecast-based production: brands designed seasons in advance, manufacturers produced in large batches to achieve unit economies, and the system absorbed its errors through markdowns and destruction. That model's core assumption that demand is predictable enough to forecast months in advance has been undermined by the emergence of forces that are reshaping the global operating landscape simultaneously.

➤ **Accelerating demand volatility:**

Investor questions about demand unpredictability and forecast failure have grown from near-zero in 2014 to a sustained 20-30 mentions per 1,000 analyst questions per year since 2022⁷. Consumer trend cycles have compressed from years to weeks, driven by an increasing number of forces competing for the finite resource of human attention. This is largely driven by social media and commerce. Long-range forecasting is increasingly unreliable, placing pressure on brands to respond to demand as it emerges rather than predicting it months in advance.

➤ **Geopolitical and trade disruption:**

Investor concern about geographic concentration and geopolitical overexposure surged to roughly seven times its prior-decade average in 2025, the sharpest single-year jump of any signal in the dataset. The fashion industry has absorbed five major supply chain shocks⁸ in as many years. The effective tariff burden on Chinese-origin apparel imports rose sharply through 2025, with China-specific surcharges reaching 145% on top of base rates at their peak⁹. The landed cost of goods has become a moving target. Supply chain disruption and lead-time volatility, barely mentioned before 2017, have become a persistent concern since 2019.

➤ **Mounting regulatory pressure:**

Regulatory pressure is intensifying, driven by the EU's ban on the destruction of unsold apparel that took effect in July 2026 and the

Digital Product Passport (DPP), which will mandate product-level transparency from raw material origin to end-of-life recyclability for textiles from 2028 to 2029¹⁰. In the US, California's SB 707, signed into law in September 2024, establishes the nation's first textile Extended Producer Responsibility (EPR) scheme. Producers were required to join an approved Producer Responsibility Organisation by 1 July 2026, with full programme compliance obligations phased in through 2030¹¹. Investor questioning on compliance remains modest today, but these incoming regulations are poised to move it from a background consideration to a boardroom priority.

➤ **Intensifying climate disruption:**

Extreme weather is disrupting the availability and quality of raw materials like cotton, while rising temperatures and water scarcity pose physical risks to manufacturing hubs. Investor questioning about climate risk to operations, while still nascent in absolute terms, has grown from near-zero before 2016 to a small but recurring presence, registering fewer than five mentions per 1,000 questions in any year.

Critically, these forces do not operate independently; they compound. Tariff volatility makes geographic diversification of sourcing urgent, but climate risk constrains where companies can credibly diversify. Regulatory pressure demands traceability across whichever new geographies are chosen, at precisely the moment when trade disruption is reshuffling the supplier base. And demand volatility shortens the window in which all these decisions must be made. This compounding dynamic, rather than any single force in isolation, makes incremental adjustments to the legacy model insufficient.

The strategic imperative is clear: shift from a supply chain that pushes forecasted products into the market to one that is agile enough to pull production in response to real-time demand signals, while building the transparency, resilience, and supplier partnerships needed to operate under sustained uncertainty.

⁷ Authors' proprietary analysis of 1,000 fashion earnings calls, 2014–2025, details Appendix A

⁸ The five shocks, in sequence: the COVID-19 pandemic (2020), which closed factories and ports and triggered demand whiplash; the six-day blockage of the Suez Canal by the Ever Given (March 2021), which delayed some 400 ships at an estimated \$9.6bn of trade per day (Lloyd's List); the container and ocean-freight crisis of 2021–22, when freight rates rose roughly fivefold (Drewry World Container Index); Russia's invasion of Ukraine (February 2022) and the energy and input-cost inflation that followed; and the Red Sea shipping crisis from late 2023, which forced carriers to reroute around the Cape of Good Hope and cut Red Sea container volumes by around 65% (Reuters). Sources: Business of Fashion; Lloyd's List; Drewry; Reuters.

⁹ Yale Budget Lab, "State of U.S. Tariffs: July 28, 2025," July 2025.

¹⁰ EU DPP Timeline, edigitalproductpassport.org; Seismic Change, "Digital Product Passports in Fashion," March 2026; digiproductpass.com practical guide, November 2025.

¹¹ California Responsible Textile Recovery Act (SB 707), signed 22 September 2024. CalRecycle: <https://calrecycle.ca.gov/epr/textiles/>

WHAT THIS REPORT OFFERS

This report is built on a premise: that the gap between the industry's operating model and its operating environment is now wide enough to demand a fundamentally different approach, and that the building blocks for that approach already exist and are being deployed by early movers with measurable success.

The analysis is grounded in three sources. First, primary research, including in-depth expert interviews with operators and suppliers running agile supply chain models at scale, providing operational detail that is rarely visible in public reporting.

Second, a set of seven interdependent principles for designing supply chains built for uncertainty, anchored in published literature and previous research on the topic. These span demand-driven operations, digital infrastructure, transparency, circularity, fair value distribution, climate adaptation, and building ecosystems.

Third, a quantitative diagnosis of what financial markets care about, drawn from an analysis of a decade of investor questions, and a financial benchmark tracking the operating performance of listed fashion companies.

For each, we examine how the principles translate into concrete operational decisions, illustrate them with case studies drawn from our research, and provide a practical starting point: where to begin implementation and what to monitor as they scale. The aim is to offer senior leaders a working guide grounded in evidence and operational reality for building supply chains that turn uncertainty from a source of value destruction into a source of competitive advantage.

The report is structured around the three questions that dominate boardroom and investor conversations:

1

How can we reduce inventory risk and the associated working capital trap?

2

How can we protect gross margins from the pressures of markdown dependency?

3

How can we build a business model that is compliant with—and benefits from—the transition to a circular economy?

A photograph of a port at dusk. Large gantry cranes are visible, with their lights glowing. Shipping containers are stacked on the left and right. The sky is a deep blue. The overall scene is industrial and captures the end of the day.

PRINCIPLES FOR SUPPLY CHAINS OPERATING IN UNCERTAIN TIMES

The bullwhip effect was first named at Procter & Gamble in the early 1990s. Managers noticed that orders for Pampers diapers swung wildly up and down, even though babies use diapers at a steady rate. Tracing it backward, they found that small changes in consumer demand became amplified into larger and larger swings as orders moved upstream toward distributors, the manufacturer, and suppliers, growing at each step like the crack of a whip.

These inflated signals were costly: fooled into thinking demand was surging, manufacturers often ramped up production and expanded capacity, only to find those investments were unnecessary once real demand settled. The concept was popularised by Hau Lee and colleagues in a landmark 1997 paper¹² though Jay Forrester had described the same dynamic back in 1961¹³.

In the first quarter of 2023, yarn exports across seven of the world's largest textile-producing countries fell 40 percent compared to the year before. This did not happen because consumers stopped buying clothes, but because the forecast-driven model that connects consumer demand to raw material production had over-amplified a market signal.

Brands had over-ordered in 2021 to compensate for pandemic disruptions. When demand normalised, they cancelled. Factories that had been running at full capacity were operating at 30 to 40 percent below it within eighteen months.¹⁴

The cost was borne across the supply chain: brands carried excess finished goods inventory, manufacturers absorbed cancelled orders and carried dead stock in the form of raw materials, and raw material producers faced a demand collapse they had no role in creating. This is the outcome of a system in which information is concentrated at the point of sale, while risk is borne across the chain.

The seven principles that follow describe the operational foundations of supply chains built for different conditions. They are not a checklist, and they are not sequential. They are a system: each principle amplifies or constrains the others. Transparency enables risk management. Fair value distribution enables agility. Digital infrastructure enables the coordination layer on which the system depends.

Taken together, they aim to distribute risk more openly across the supply chain, replacing a model in which forecast error cascades downstream with one in which information flows and the ability to act on it are shared across production tiers.

¹² Lee, H. L., Padmanabhan, V., & Whang, S. (1997). The Bullwhip Effect in Supply Chains. *Sloan Management Review*, 38(3), 93–102.

¹³ Forrester, J. W. (1961). *Industrial Dynamics*. Cambridge, MA: MIT Press.

¹⁴ McKinsey & Company / Business of Fashion, *The State of Fashion 2024*, December 2023.

1. DEMAND-DRIVEN, AGILE & RESPONSIVE OPERATIONS

1.1 From Forecast to Signal

A demand-driven operating model moves the primary production decision from assumption to evidence, but it does not necessarily apply uniformly across a product range. For SKUs with genuinely predictable demand like core basics, replenishment lines, or proven carryovers, the traditional forecast-commit model remains efficient. The demand-driven approach delivers its greatest value for SKUs where demand uncertainty is high: new styles, trend-driven items, seasonal introductions, and untested categories. This approach may also extend to core basics or replenishment lines where there is measurable size-related risk (e.g. XXL or XXS).

For these, instead of committing to volume upfront based on a forecast, brands commit to a small initial quantity across a wider range of options, observe what actually sells, and scale production behind real demand signals. The forecast does not disappear; it shrinks in scope and consequence. Its job is no longer to determine total volume, but rather to determine where to start.

1.2 Why It Matters Now

The Brand Reality

Industry data shows that between 25 and 40 percent of fashion inventory¹⁵ is never sold at full price, leading to margin erosion through markdowns, warehousing costs and immobilised capital.

Demand-driven models have the potential to reduce this erosion. By limiting the scale of initial production, the financial exposure of any single bet is capped. Successful products are replenished swiftly based on actual sales velocity. Unsuccessful designs are discontinued with minimal residual stock. The result is a fundamentally different risk profile. The demand-driven model examined in our research reports an unsold inventory rate of approximately 2 percent¹⁶.

The relationship between inventory efficiency and profitability is not theoretical. An analysis of Days Inventory Outstanding (DIO) against operating margin across ten listed brands over 2013-2025 in two key segments—sportswear and mass-market apparel—reveals a consistent inverse relationship¹⁷. In sportswear, companies with lower DIO (faster inventory turns) operate at materially higher margins, with a cross-sectional correlation of $r = -0.60$. Deckers, with approximately 85 days of inventory, operates at a 24% margin; Adidas and Puma, carrying 155–160 days, sit at 5–7%.


¹⁵ McKinsey & Company / Business of Fashion, "Tackling Fashion's Excess Inventory Problem," State of Fashion 2025, 2025.

¹⁶ Boston Consulting Group, "Agility Is Fashion's New Source of Competitive Advantage," 2023, <https://www.bcg.com/publications/2023/creating-agile-supply-chains-in-the-fashion-industry>.

¹⁷ Source: Authors' analysis of public company financial filings, 2013–2025.

► The inverted relationship between inventory and margins

Sportswear | DIO vs operating margin | most recent reported fiscal year



$r = -0.60$

Stable 2013—2024
pooled $r = -0.54$


Source: Swanstant PCI database | actuals | periods ending Nov 2024—Aug 2025 | 5 brands per tier

The pattern holds in mass-market apparel ($r = -0.37$), where Inditex—the segment's most agile operator—carries the lowest inventory (approximately 70 days) and the highest margin (20%).

Correlations do not prove causation; companies with stronger brand relevance may simultaneously command higher margins and manage inventory more tightly. But the consistency of the pattern across segments and over a decade suggests that inventory discipline is, at minimum, a necessary condition for margin leadership.

► The inverted relationship between inventory and margins

Mass-market apparel | DIO vs operating margin | most recent reported fiscal year



$r = -0.37$

Stable 2013—2025
pooled $r = -0.69$

Source: Swanstant PCI database | actuals | periods ending Nov 2024—Aug 2025 | 5 brands per tier

The Manufacturer Reality

When a forecast leaves a demand-supply gap, the cost impact goes beyond the brand's balance sheet. Manufacturers often absorb the impact of cancelled orders on fabric that has already been cut, with buyers refusing to pay raw material and production costs already incurred. During periods of demand shock, this extends to payment terms: industry surveys have documented buyers stretching settlement from a pre-disruption average of around 43 days to 60, 90 or even 120 days on work already completed; terms that, before the pandemic, were considered highly exceptional¹⁸. Margins are compressed on the next order to compensate for losses on the last.

A demand-driven model shifts this dynamic for both sides. Factories produce smaller batches linked to actual demand rather than speculative forecasts. While individual orders are smaller, more frequent shipments mean manufacturers experience healthier cash flow cycles. This applies not only to finished goods orders being triggered by a purchase order but also to raw material flows of yarn and fabric. To unlock the demand-driven model, mechanisms such as raw material pre-positioning and capacity blocking are necessary. This creates a more stable and predictable flow

of production, replacing the financial and volume uncertainty of seasonal bulk ordering with something closer to a working relationship built on shared visibility, shared risk, and shared value.


Demand-driven supply chains in fashion require a decoupling of raw-material and finished goods supply-chain decisions. Brands work closely with Tier 2 (raw material suppliers) to intelligently pre-position raw materials in different stages of completion. In operations research, this approach is known as postponement. For example, for a given high-risk SKU, material may be kept in the yarn or grey fabric state (greige). As confidence in the demand signal increases, additional production actions are taken, e.g. dye, cut and sew. The deadline for these steps is driven by production lead times and expected ex-factory delivery dates.

For lower-risk SKUs, brands may pre-position at more advanced stages of production, e.g. dyed fabric. This enables shorter lead times from purchase order issuance to delivery.

The objective of postponement in a demand-driven supply chain is to delay the decision of converting materials into finished goods to the last responsible moment. The last responsible moment is determined by the relative risk of each SKU.

Two thousand years ago, the Pyrrhonist philosopher Aenesidemus coined the term epoché. The principle is simple: when evidence is insufficient, defer the irreversible decision until the signal is clearer.

The Risk Escalation Pyramid



Early-stage commitments offer speed to market but lock in costs and reduce flexibility. Brands constantly navigate this tension between agility and financial liability, particularly as market volatility increases.

¹⁸ Center for Global Workers' Rights (CGWR), Penn State University, Abandoned? The Impact of COVID-19 on Workers and Businesses at the Bottom of Global Garment Supply Chains, April 2020 (72% of buyers refused to pay for raw materials already purchased; 91% refused to pay cut-make-trim costs on cancelled orders). Complemented by: CGWR, Leveraging Desperation: Apparel Brands' Purchasing Practices during Covid-19, October 2020, finding that 61% of suppliers were paid 60–120 days after shipment and one in four buyers imposed terms of 120 days or longer, versus a pre-pandemic average of 43 days. For named brand examples: Business & Human Rights Resource Centre, "TK Maxx cancels orders and extends supplier payment terms from 30 to 120 days," April 2020.

1.3 Success Factors

Among the conditions that most consistently determine whether a brand can operationalise this principle effectively, three stand out.

> **Production flexibility:**

The ability to run small batches without unit cost penalties that make the model commercially unviable. This requires supplier relationships and manufacturing configurations that can accommodate high-frequency smaller orders rather than periodic large ones.

> **Network diversification:**

Sourcing spread across geographies and lead time profiles, enabling SKUs to be routed to the right production configuration based on their demand uncertainty. This requires SKUs to be developed at multiple geographies in anticipation of a potential reroute. Predictable basics can remain offshore on longer cycles.

> **Total landed cost visibility:**

The ability to model each production decision against its full cost including tariffs, transport mode, markdown risk, and inventory carrying cost. Incoming margin alone is an incomplete and often misleading basis for sourcing decisions.

1.4 Where Are You Today

Demand-driven operations describe a direction of travel, not a fixed destination. Where a brand sits today is shaped by its channel mix, supplier relationships, digital maturity, and the proportion of its range that carries genuine demand uncertainty. Where it needs to get to is a commercial decision, not a compliance exercise.

We have classified supply chain operations into three categories, which represent different points on the direction of travel.

Position 1: Forecast Driven

Production volume is determined months in advance across the full range. Seasonal buys are large, lead times (time from purchase order issuance to delivery) are long, and the system absorbs its errors through markdowns, clearance, excess inventory, and working-capital drag. Most of the industry operates here.

Position 2: Selectively Responsive

A portion of the range—typically the SKUs judged as high-risk or most trend-exposed—is carved out for smaller initial quantities and shorter lead times, while stable SKUs, e.g. core styles, colours and wholesale commitments, stay on the forecast model. It is worth noting that risk can fall in multiple dimensions: style, colour, or size - or even have regional dependencies. For example, one style might be stable in North America but not in Asia.

Moving to Position 2 captures part of the benefit of being demand-driven, but does not escape the challenges of forecasting. The decision of which SKUs belong on the responsive track is itself a forecast, made at the range-planning stage, with the same scarcity of evidence and the same exposure to error as the volume forecast it replaces. A “core basic” SKU that unexpectedly turns volatile stays trapped on the slow track while a “trend item” that settles into steady demand keeps incurring the cost premium of responsive production it no longer requires. Because the classification is set once and rarely revisited within the season, these misclassifications persist until the next planning cycle or beyond. Position 2's ceiling is set by the accuracy of a judgement made before the season starts.

Position 3: Demand-Driven at Scale

Position 3 does not permanently assign SKUs to fixed tracks; it dynamically allocates each commitment according to demand confidence, economics, and lead time. The governing principle is that every commitment—for a high-volume staple as much as an untested trend piece—is made as close to the demand signal as the product's economics and lead times permit, and sized to the confidence that the signal provides. An uncertain style, colour, or size launches small and scales only as sales velocity confirms it. A proven, high-volume seller still warrants a large production commitment, but that commitment is placed dynamically, at the point the signal justifies it, and structured to preserve flexibility, rather than fixed months ahead against a forecast. What unifies the range is therefore not a classification but a discipline: defer each decision to the latest responsible moment, and calibrate its size to the strength of the evidence rather than the calendar.

The Omnichannel Reality

A brand managing physical retail floor sets, wholesale commitments, and multiple consumer segments cannot move from Position 1 to Position 3 in a single planning cycle. It does not need to. Shifting 20 to 30 percent of the highest-uncertainty SKUs from large-batch forecast production to smaller responsive batches materially reduces inventory risk and markdown dependency without dismantling the operational infrastructure serving the rest of the business.

The lever does not need to be applied across the entire assortment; it is most useful when applied to the SKUs carrying the most uncertainty and risk.

1.5 Business Benefits

Moving the production decision from forecast to signal delivers three returns:

> **Capped downside:**

For novel, high-uncertainty SKUs, an 80–100 unit launch contains the cost of any single misjudgment. This can scale depending on whether there are signals from other similar styles or colours.

> **Captured upside:**

Every commitment—small or large—is timed to the demand signal and sized to the confidence it provides. A proven seller can be backed at full scale the moment the evidence justifies it. The brand captures demand that both forecast commitment and conservative ordering leave on the table.

> **Capital and margin benefits:**

Less capital trapped in stock; fewer units cleared at markdown. The durable inverse relationship between inventory velocity and operating margin documented earlier is the signature of this. Operating in a demand-driven fashion also requires greater communication and collaboration with suppliers. Over time, this shifts brand-supplier relationships from transactional to shared-risk and shared-value partnerships.

DIGITAL INFRASTRUCTURE, INTELLIGENCE & GOVERNANCE

2.1 From Latency to Live

Digital infrastructure is essential for demand-driven models, but technology on its own is not the solution.

Technology layered on top of a sound process compresses the time between a market signal and a production decision from months to minutes. Technology applied to a broken process has the potential to amplify the impact of bad decisions. As one operator running a 15-day concept-to-shelf cycle put it, technology layered on bad processes makes things worse¹⁹. For them, the process discipline came first, and the systems were built on top of it.

This sequencing is often overlooked. Before technology and AI can deliver operational improvements, the underlying process has to be re-engineered for responsiveness: data governance, granularity and frequency of data signals across the value chain, buying decision sequencing, defined performance standards, clear escalation when a supplier falls short, and a planning cadence that can act on a signal rather than file it. These are organisational disciplines, not software features. The shift this principle describes is a two-step move: redesign the process, then build the infrastructure that lets it operate at speed and across production tiers.

2.2 Why It Matters Now

The Brand Reality

The fashion industry has spent a decade investing in visibility: warehouse inventory dashboards, sell-through analytics, forecasting tools, and more recently, AI. The investment has been significant, but the operational returns have been inconsistent. The reason is often that the tools were built to help monitor the outcomes of current processes rather than coordinate decisions with the suppliers who can act on the information. This distinction matters: visibility is not coordination. A dashboard that shows a brand its own sell-through can inform its own processes but does not change what the supplier does next, because the supplier never sees it.

The information that would let a fabric mill or garment factory act earlier—future demand signals, real-time demand velocity, reorder probability—stops inside the brand.

The existing process means that the data that does reach suppliers is still transmitted through narrow, intermittent channels. The same phenomenon exists from the suppliers back to the brands. The result is that brands have improved their ability to observe the consequences of forecast error without fundamentally improving their ability to prevent it. Demand-driven operations require the opposite: a shared intelligence layer in which demand signals, predictive analytics, and supplier capacity flow in both directions, so that decisions are made where they have the most effect rather than only where they are most visible.

The Manufacturer Reality

From the factory floor, the information gap looks different. A garment manufacturer receives an order, produces against it, and ships. What happens to that product after it leaves is invisible. Whether it sells at full price, sits unsold, or triggers a reorder generates no signal that reaches production planning.

¹⁹ Primary interview conducted for this research, 2026.

The result is speculative behaviour compounding across tiers. Fabric mills plan production and hold stock based on judgement. Garment factories schedule capacity against orders or speculative forecasts they cannot anchor to actual demand. Risk accumulates silently until it surfaces as unsold inventory at the brand level and capacity loss and material waste at the supplier. When demand or a forecast changes, suppliers with limited technology integration struggle to replan against it. Equally, brands are unable to leverage any of the suppliers' data, e.g. real-time capacity, available material stock (yarns and fabrics) or day-to-day challenges such as quality issues to inform their decision-making. One such issue is lead times. Brands plan based on rule-of-thumb lead times provided by the supplier that are neither a reflection of reality on the production floor nor updated frequently enough to account for an influx of orders, delays at upstream producers, or geopolitical impacts such as port congestion.

When information flows both ways, and when suppliers have the tools to act on it, suppliers can pre-position fabric, adjust capacity, and reduce the speculative production that drives excess inventory across the chain. But the data alone is not enough. The supplier must have both the visibility and the internal planning capability to convert it into action, which is why this principle is as much about information management as it is about processes.

2.3 Success Factors

Four conditions most consistently determine whether a brand can operationalise this principle.

- > **Process definition before digitisation:**
The underlying workflow must be re-engineered for responsiveness before technology is applied. Defined planning cadences, single-source accountability per style, and clear performance escalation are the foundation. Technology amplifies the process it sits on.
- > **Data standardisation and integrity:**
Shared, standardised fields (and confidence in their accuracy) are what allow intelligence to move across tiers without distortion. Sophisticated analytics on inconsistent or unreliable data produce confident errors. Standardisation is usually the binding constraint.

> **Supplier onboarding and tooling:**

The intelligence is only as useful as the supplier's ability to act on it. This requires investment in supplier-side planning tools and structured onboarding, so that a demand signal received at Tiers 1 and 2 can be converted into a pre-positioning or capacity decision. Equally, the ability for suppliers to feed their physical and production constraints back to the brand enables better planning and a clearer view of what is feasible.

> **Governance:**

Clear rules on who can see what, who decides what, and how data is used across a multi-tier network. Without governance, bidirectional transparency stalls on legitimate concerns about commercial confidentiality and data misuse. With it, suppliers and brands can share the intelligence the model depends on.

2.4 Where Are You Today

Most brands are somewhere between holding data and using it to coordinate. The gap is rarely technological sophistication. It is a question of whether the process has been redesigned and whether the resulting intelligence reaches the people who act on it. The three positions below therefore move along two dimensions at once—process discipline and cross-tier integration—because advancing one without the other yields little.

Position 1: Siloed Systems, Traditional Processes

Brand and supplier systems do not communicate, and the underlying processes have not been re-engineered for responsiveness. Decisions are made on periodic reports. Suppliers receive orders without demand context. Information moves sequentially and slowly.

Position 2: Redesigned Internally, Connected to Tier 1

The brand has done the process work for its own operations. It has defined planning cadences, performance standards, single-source accountability, and has connected sell-through and forecast data to key Tier 1 suppliers. This position's ceiling is structural: Tier 2 and beyond remain outside the loop, and the connection usually rests on bilateral integrations rather than shared standards, so coordination degrades with each step up the chain. Many brands have run AI and demand-sensing pilots, which sit on unstandardised data and are unconnected to suppliers.

Position 3: Networked and Continuous

Redesigned processes are embedded across the network, and real-time signals reach at least Tier 1 and 2 suppliers on a common set of standardised fields. Suppliers transact directly within a shared platform or protocol and hold the planning tools to act on what they see. Fabric availability, production capacity, and sales velocity inform each other continuously. The decisive feature is not technological sophistication but the combination of process discipline, functional standardisation and multidirectional data flow, targeting zero information lead time. In practice, this may be a simple system, provided the process is disciplined, standardised, and fully connected. This extends to communication between Tier 1 and Tier 2. For example, where different Tier 1 suppliers pull fabric from a shared Tier 2, fabric availability is shared in real-time in order to ensure that sufficient materials are available in the necessary stages to proceed with production.

From Closed Networks to Open Standards

The gap between where most brands sit and where they need to get to is less of a technology investment problem than a data standardisation one. Systems that cannot speak to each other cannot share intelligence, regardless of how sophisticated they are individually.

The most integrated networks operating today demonstrate both the power and the limit of standardisation. Where a single orchestrator can impose common data definitions across its suppliers—a shared catalogue, standardised order fields, real-time sales visibility, direct cross-tier

transaction records—the coordination gains are substantial, and they can be achieved on technically modest infrastructure. Functional standardisation, not analytical sophistication, is what generates the advantage. But this form of standardisation is proprietary: it works across a brand and its supplier network, but does not extend beyond it.

This is a structural ceiling, and the cost surfaces on the supplier side. A single orchestrator can mandate its proprietary standard because it controls the relationship. But suppliers rarely serve a single brand. When each brand a supplier works with imposes its own proprietary system, the supplier must adopt and maintain several incompatible ones, duplicating cost, fragmenting their own data, and eroding the very efficiency that standardisation was meant to create. What is efficient within one network becomes inefficient across the industry, and the burden falls on manufacturers.

Open, interoperable standards resolve this because a supplier adopts once and connects with all. That is a materially harder problem to solve, because it depends on collective adoption rather than unilateral mandate—but the direction of travel is unambiguous: from proprietary standardisation inside closed networks toward shared standardisation across open ones.

2.5 Business Benefits

Three returns are specific to the infrastructure, not the intent:

- > **Speed:**
Signal-to-action latency falls from months to days. A sound process can decide correctly while the infrastructure enables that decision to be fast enough to capture full-price demand.
- > **Risk falls across the network, not just the brand or the manufacturer:**
Because intelligence flows both ways, suppliers stop producing against guesswork. Excess is prevented at source.
- > **One investment, two returns:**
The infrastructure built for commercial coordination is the same infrastructure that satisfies traceability and due-diligence regulations. Compliance becomes a by-product.

3. TRANSPARENCY, TRACEABILITY & COMPLIANCE

3.1 One Capability, Two Returns

Transparency and traceability are usually framed as a compliance burden. However, the capability that regulators now demand—verifiable knowledge of provenance from raw materials to finished products—is the same capability that determines whether a business can manage supply risk in an uncertain environment. A brand that can trace a fibre to source can both document compliance and respond when that source faces any issue or disruption. And as the EU and other authorities move to enforce existing regulation more forcefully²⁰ traceability becomes the means of coordinating that compliance across the supply chain, enabling the engagement of every supplier in delivering a compliant final product.

It is also inseparable from agility. A demand-driven supply chain runs on multi-tier information flows, and the same connectivity that lets a brand sense demand and reallocate production is what makes its chain traceable. Traceability is therefore both an enabler of agility—dynamic planning is only possible with suppliers you are connected to—and an outcome of it, since the digital infrastructure built for responsiveness generates the visibility compliance requires. The investment case does not rest on regulation alone. But regulation is making the cost of not having the capability more immediate.

3.2 Why It Matters Now

The Brand Reality

Across industries, multi-tier visibility remains limited. A McKinsey 2025 survey of global supply chain leaders found that while 95% have visibility into their Tier 1 risks, only 42% have visibility into Tier

2 or beyond²¹. A separate McKinsey 2021 survey found that only 2% of companies have visibility into Tier 3 suppliers²². KPMG's own research puts the Tier 1 problem in sharper relief: 43% of supply chain leaders report limited or no visibility even into their direct Tier 1 suppliers²³. Even where Tier 1 visibility has improved, McKinsey's 2024 supply-chain survey notes that visibility deeper in the chain has declined for the second consecutive year²⁴. In fashion, the gap matters more, because the highest-risk stages sit upstream: fabric mills and dye houses at Tier 2, fibre and yarn at Tier 3, raw material at the source. The distance between what a brand can disclose and what it can verify is widest exactly where its chemical, labour, and regulatory exposure is greatest.

That exposure is no longer hypothetical. The laws reaching the unverified tiers already exist, and their enforcement has escalated. REACH has restricted carcinogenic and reprotoxic substances in skin-contact textiles since 2020²⁵; the EU's Unfair Commercial Practices Directive has been enforceable since 2007²⁶. What changed is the appetite to apply them: France's Directorate-General for Competition, Consumer Affairs, and Fraud Control (DGCCRF) has pursued enforcement actions against multiple fashion brands for misleading pricing and environmental claims, with fines in individual cases reaching into the tens of millions of euros²⁷, and the UK's Digital Markets, Competition and Consumers Act now allows the Competition and Markets Authority to directly impose penalties not exceeding £300,000 or 10% of global turnover, whichever is higher²⁸. The exposure is independent of the progress around Ecodesign for Sustainable Products Regulation (ESPR) or other flagship EU Green Deal regulations. It only requires that existing law be enforced more forcefully, which is starting to happen.

20 For example, France's DGCCRF (Directorate-General for Competition, Consumer Affairs and Fraud Control) has increased scrutiny of existing consumer and environmental law compliance across sectors, including textiles. Under AGECLaw's Article 13 traceability obligations (mandatory geographic origin disclosure, phased in for textile firms from January 2023), the DGCCRF began active controls. In a 2024 anti-waste inspection campaign, the DGCCRF inspected over 110 textile companies. Across the full campaign covering food and textile businesses, it issued 21 injunctions and 2 criminal proceedings. Separately, in its 2023–2024 review of environmental claims covering over 3,000 businesses across textiles, furniture and other sectors, the DGCCRF found that more than 15% presented serious violations, resulting in over 430 compliance injunctions and 70 administrative fines and criminal reports. The number of controls increased by 38% in 2024 relative to 2023.

21 McKinsey 2025: Onspring, "The Tier 2 and Tier 3 Blind Spot," April 2026 (citing McKinsey 2025 Supply Chain Survey)

22 McKinsey 2021: Procurement Mag, December 2022, citing 2021 McKinsey global survey

23 Sedex, From Risk to Resilience: Build a Stronger Supply Chain, May 2026

24 McKinsey & Company, Global Supply Chain Leader Survey 2024

25 European Commission, "REACH Annex XVII Restrictions on CMRs in Textiles," October 2018 (effective October 2020).

26 Directive 2005/29/EC of the European Parliament and of the Council concerning unfair business-to-consumer commercial practices in the internal market, Official Journal of the European Union, L 149/22, June 11, 2005.

27 DGCCRF press releases; see specifically 2023–2024 enforcement actions against fashion brands for greenwashing under the UCPD.

28 Pinsent Masons, "DMCC Act overhauls UK consumer law enforcement," September 2024.

The Manufacturer Reality

For manufacturers, the visibility picture is similar but the economics are not. A garment factory knows its own direct suppliers—the mills and trim makers it buys from—but typically sees no further up the chain than the brand does, relying on the same attestations. The difference is in the incentive. A brand faces direct regulatory exposure and a clear resilience case: it is the party fined, and the one whose revenue depends on continuity of supply. For an upstream supplier, the return on investing in traceability is weaker and less direct, and the cost of documentation often outweighs any visible payoff.

This is where the mechanism a brand chooses matters. Responsibility can be passed down contractually: numerous brands have a delegation clause in their supplier code of conduct that makes Tier 1 suppliers accountable for upstream compliance. This transfers cost and liability without creating genuine economic upside incentives. Or responsibility can be shared, where brands invest in traceability and make it more economically neutral or rewarding for suppliers to join, enabling better data flow. The gap closes not by mandate alone but by aligning incentives with the actors who must make the investment.

3.3 Success Factors

Among the conditions that most consistently determine whether a brand can operationalise this principle, four stand out.

> **Designing with downstream awareness:**

The more fragmented a product's bill of materials, supplier base, and production route, the heavier the evidence burden becomes. This does not mean brands should avoid sourcing flexibility. It means complexity should be designed with visibility in mind from the outset. An agile supply chain depends on making decisions later and faster; as transparency and traceability requirements tighten, those decisions only create value if the underlying data can keep pace.

> **Incentive alignment across tiers:**

The binding constraint upstream is rarely just technical. It often lies in the fact that transparency, visibility, and compliance are viewed as a tax on suppliers. Brands rarely, if ever, factor in the suppliers' cost of investment to enable transparency. The return for suppliers therefore has to be concrete: faster onboarding, fewer duplicated audits, clearer and quicker production approvals, preferred-supplier status, better visibility into capacity requirements, and more predictable order allocation. In the operators we studied, orders are allocated based on demonstrated performance and transparency; the supplier that invests in visibility is therefore better positioned to win the next order. When provenance data is captured once and reused across onboarding, audit, approval, and allocation processes, compliance cost becomes an operating advantage.

> **Verification, not just declaration:**

Most traceability data today is self-reported. This creates the risk of missing subcontracting or other variations that may appear in the supply chain. The standardised data flows described in Principle 2 are the foundation, but the value lies in verifying what is declared. The gap between declared and verified data is where regulatory exposure sits.

> **Build for business resilience:**

A system designed only to populate a compliance document is a cost. One designed to monitor risk and inform sourcing decisions earns its keep operationally and satisfies compliance as a byproduct.

3.4 Where Are You Today

Traceability is a question of depth, verification, and purpose. Most brands have some visibility. Few have verified visibility built to act on.

Position 1: Tier 1 and partial Tier 2, compliance-driven.

Direct suppliers are known, and some fabric mills are mapped. Below that, the brand has limited visibility. The data is largely self-declared, and the purpose is to satisfy audit and reporting requirements, creating a static snapshot. Most of the industry sits here.

Position 2: Extended traceability, still declarative.

The brand has invested in a traceability platform and pushed mapping further toward raw material, often ahead of a regulatory deadline. Coverage is broader, but the data remains largely a digitised chain of attestations rather than verified facts, and it is built to report rather than to act on. This is where the most advanced brands sit today. The map is more extensive, but remains a snapshot.

Position 3: Verified traceability across strategic nodes, resilience-driven.

The brand identifies the nodes that matter most—high-risk materials, concentrated dependencies, the chemical and labour stages most exposed—and verifies and monitors those continuously. Dynamic data feeds live decisions on sourcing, planning, risk, and disruption. Compliance is satisfied as a byproduct of other operational requirements. Very few operate here at scale. Decision-making is unlocked by a live view of the supply chain. This transitions traceability from a tax to a source of value for the entire value chain.

Where to start.

Movement toward Position 2 does not require mapping everything at once. It starts with materials carrying the highest risk and on the actual regulatory requirements arriving the fastest. Movement toward Position 3 is less about collecting more data than about changing its purpose and raising its quality where it matters: verifying what you declare, and building the capability to act on it rather than merely report it. A system built for resilience delivers compliance as a byproduct. Position 3 requires a departure from the traditional top-down approach that the industry has taken towards a collaborative one. This means identifying opportunities for shared value and shared risk, so suppliers can unlock efficiencies while brands gain resilience and transparency.

CIRCULARITY & RESOURCE STEWARDSHIP

4.1 Designed Out, Not Managed Away

This paper's question is how to operate supply chains under uncertainty, so it approaches circularity from inside the supply chain rather than at the consumer interface. Circular business models like resale, rental, or repair are a vital part of the wider transition, but sit beyond this paper's boundary. The question here is narrower and operational: how can an agile supply chain enable the circular transition, and how can it benefit from doing it?

A circular economy maintains, restores, and circulates value. Its principles, driven by design, are to eliminate waste and pollution, circulate products and materials at their highest value, and regenerate nature²⁹. For a supply chain, a circular economy contributes in two key ways. The first is to avoid destroying value to begin with: every garment made in excess of the demand, and every kilogram of material wasted in production, is value lost before a product is ever used. The second is enabling recovered materials to circulate back into the chain as inputs, rather than being discarded at end of use.

Until recently, those losses were unaccounted for externalities: unsold goods were written off, incinerated, or shipped to secondary markets. These actions are increasingly outside what brands, and by extension their suppliers, are allowed to do. The EU's Ecodesign for Sustainable Products Regulation bans the destruction of unsold textiles by large companies from 19 July 2026³⁰, and Extended Producer Responsibility schemes make producers financially responsible for the collection, sorting, and recycling of what they place on the market. Those regulations are internalising value loss on companies' balance sheets, making it both more direct and quantifiable.

An agile supply chain can be an enabler of the circular transition. A demand-driven, digitally connected, traceable chain can reduce overproduction by ensuring the right products are made at the right time. It is also the beneficiary: in a period of climate-driven volatility in material price and availability, producing less waste and recirculating materials is not only an environmental gain but a hedge against the volatility of virgin inputs.

²⁹ Ellen MacArthur Foundation, Circular economy principles.

³⁰ European Commission, "New EU rules to stop destruction of unsold clothes and shoes," February 9, 2026.

4.2 Why It Matters Now

The Brand Reality

For brands, value leaks at three points.

The first, and largest, is overproduction. The industry makes around 25–40% more than it sells at full price. Every unsold unit is materials, energy, and transport consumed for no financial return. This is the supply chain's single biggest waste failure. With an estimated 4–9% of unsold textiles destroyed before being worn, it is also the value loss regulation now targets most directly³¹.

The second is returns. Why a customer returns a garment is partly a matter of consumer behaviour, which sits outside this paper's scope, and partly a matter of fit and sizing, which is a data problem brands can address. Better sizing and fit accuracy reduces the production of garments that were never going to be kept. A caution follows from the same logic: a low return rate is not automatically circular. What matters is whether the returned products are reused or their materials recycled.

The third is material not recovered at the end of use. When a product's composition is unknown, or its construction prevents the separation of fibres, its materials cannot be recovered and are either downcycled into lower-value uses or lost. This value could be retained through design, labelling, and traceability. It is also where resilience to raw material volatility could be unlocked.

The Manufacturer Reality

In current fashion sustainability and policy reporting, quantified waste metrics tend to focus on visible, product-level streams such as returned and unsold textiles and end-of-life disposal, while large volumes of pre-consumer process waste generated in factories (for example, cutting-room offcuts and rejected lots) remain poorly measured and rarely disclosed in a consistent way³².

When a factory lacks accurate visibility of current sales performance and future demand, it resorts to buffering based on unreliable historical patterns, holding excess fabric, running below optimal capacity to maintain flexibility, or changing over machines more frequently than necessary as schedules shift. Each decision wastes material, energy, and capacity that never reaches brands' circularity metrics because it happens on someone else's floor.


This effect is compounded by the bullwhip effect, where small fluctuations in consumer demand are progressively amplified as they travel upstream, so a minor dip at point of sale becomes a production emergency at Tier 2, triggering over-correction in both directions.

This waste is structural and invisible because it sits upstream. Sharing real-time demand signals, capacity data, and material pre-positioning across tiers collapses the information latency the bullwhip depends on, eliminating value loss that extends beyond material waste.

31 European Environment Agency (EEA), The destruction of returned and unsold textiles in Europe's circular economy, Briefing, 3 March 2024; based on ETC/CE Report 2024/4 Volumes and destruction of returned and unsold textiles in Europe's circular economy.

32 European Environment Agency (EEA), The destruction of returned and unsold textiles in Europe's circular economy, 2024, which documents policy and data efforts centred on returned and unsold textiles. Ahmed et al., Fashion Circularity: Potential of Reusing and Recycling Remnant Fabrics, Sustainability, 2025, estimating that apparel manufacturing generates about 60 billion m² of cut-and-sew fabric waste annually, most of which is landfilled or incinerated. European Parliamentary Research Service, Identifying data gaps in the textile industry and assessing current initiatives to address them, 2024, which highlights persistent data gaps on upstream production and waste flows in textile supply chains.

HYRAN TECHNOLOGIES—MAKING DEMAND-DRIVEN SUPPLY CHAINS EXECUTABLE ACROSS TIERS



A demand-driven model only works if demand signals can be converted into operational decisions across the supply chain: how much material to commit, what stage to hold it in, when to move from greige to colour, and how suppliers should plan production capacity. Hyran Technologies provides an example of how this can work in practice.

Hyran works with brands and suppliers to turn demand forecasts and live demand signals into material and production decisions. At the brand side, its planning model assesses demand confidence by style, colour, market and time horizon, then recommends how much material to pre-position and at what stage of completion. Lower-confidence demand can be held in lower-risk states, such as yarn or greige fabric. As confidence improves, the model supports the decision to move material into higher-commitment stages, such as dyed fabric, cut fabric and finished goods.

The same system creates value for suppliers. Rather than receiving demand changes as late-stage disruption, suppliers gain earlier visibility into expected material flows, capacity needs and reorder probability.

Hyran's supplier-side planning tools can then translate new demand signals into production plans, including which machines should run, which yarns should be used and how fabric orders should be sequenced. Planning that would otherwise take days can be updated in seconds.

The effect is to reduce speculation on both sides. Brands avoid over-committing to materials that may not be needed. Suppliers gain clearer visibility on yarn inflows, material positions and future capacity requirements. A shared ledger of material status—from greige to dyed fabric, across Tier 2 and Tier 1—allows both sides to delay irreversible decisions until the last responsible moment, while still protecting delivery timelines.

The reported outcomes are commercially material: Hyran's internal client data shows reductions in excess raw-material liability of over 45%, on-time delivery improvements of over 28%, and supplier machine-utilisation improvements of around 15%³³. The case illustrates the central argument of this report: agility is not created by smaller orders alone. It depends on the digital, financial and operational infrastructure that allows brands and suppliers to share risk, act on demand signals and prevent waste before it is created.

33 Hyran Technologies, internal client data (2023-2025)

4.3 Success Factors

Among the conditions that most consistently determine whether a brand can operationalise this principle, four stand out.

- > **Design for product and material circulation:**

Material choices and product architecture set at the design stage determine whether a garment's materials can later be identified, separated and recovered.

- > **Demand accuracy as waste prevention:**

The most effective intervention is producing closer to what is actually needed, including through better sizing and fit data. This is not a circularity tool bolted on at the end; it is the demand-signal integration of Principles 1 and 2 applied to the largest single source of value loss.

- > **Manufacturing planning visibility:**

Sharing current sales performance, future demand, capacity, and material pre-positioning across tiers reduces excess material stock and changeovers. The gain is structural, not incremental.

- > **Traceability to close the loop:**

Accurate labelling or digital product passport information are critical enablers of material recovery. Precise composition details that support faster sorting and re-processing can support price parity, critical to the adoption of recycled materials at scale. However, this needs to be a byproduct of the existing planning process within the supplier side instead of a parallel process that sinks time and resources without proper incentivisation for suppliers.

4.4 Where Are You Today

The earlier in the supply chain a brand can engineer out future inefficiencies, the less value is wasted across the system. The objective is to transition from a system that attempts to manage waste after it is created to designing a supply chain that eliminates it from the outset.

Position 1: End of Life Management

Production aligns with forecasts and waste is managed downstream through recycling partnerships, liquidation, and disposal. These programmes are visible and increasingly required by regulation, but they are the least efficient form of circularity, acting at the point where most of a product's value is being lost. Most of the industry sits here.

Position 2: Waste prevented within the chain

Brands share current demand performance and forward-looking demand signals with Tier 1 and Tier 2 suppliers at a regular cadence. Brands are supported by AI to automatically determine the last responsible moment to transition from a low-risk state (yarn, undyed fabric) to higher-risk states (dyed fabric, cut fabric), often called postponement. Suppliers share current material positions and changes in available capacity at a regular cadence. This in turn informs the brand's ongoing planning decisions. All of this is enabled through technology.

The continuous planning approach enabled by real-time visibility minimises overstocked finished goods at the brand level, but also cuts excess material stock, machine changeovers, and efficiency loss at the manufacturing level.

Less value is destroyed from the outset. This reduces the amount of waste that needs to be managed downstream. Very few brands operate here today. This capability is dependent on combining the demand-driven operations and digital infrastructure of Principles 1 and 2.

Position 3: A supply chain where waste is eliminated

On top of waste prevention, the chain is designed to trace and circulate materials back: products are designed and manufactured for recovery, composition is verifiable at end of use, and recovered material re-enters as feedstock. Textile-to-textile recycling is still maturing and lacks scaled commercial availability, although innovative solutions are starting to scale. This limits the ability of brands to operate in that position today.

5. FAIR VALUE DISTRIBUTION

5.1 Who Finances Whom

Fashion supply chains are built on an unexamined assumption: that the financial burden of production sits with the party best equipped to carry it. In practice, the opposite is frequently true. A garment factory in Bangladesh completes an order. It has already paid for the yarn, fabric, trims, dye chemicals, labour, energy, and logistics. The goods have shipped. The factory now waits—sixty, seventy, ninety, 120 days—for payment on work it has already completed. Through that wait, the supplier is financing the brand's working capital. In many manufacturing countries, interest rates run well into the double digits.

This is not a moral argument but a financial one. The parties with the strongest balance sheets—global fashion brands with investment-grade credit ratings and access to cheap capital—routinely receive interest-free financing from the parties with the weakest.

The asymmetry is not just financial. The same system that results in suppliers financing brands' working capital also requires suppliers to absorb brands' demand risk.

A cancelled or revised order leaves the factory holding goods for which material was already procured or for which the product was already produced. A volatile sourcing strategy leaves suppliers carrying capacity they cannot fill. The party least able to forecast or absorb demand volatility carries the most exposure to it. Fair value distribution addresses both questions at once: who finances production, and who bears the risk that demand does not materialise as planned.

The payment architecture was not designed with this asymmetry in mind. It was designed for a different production model—large batch, long lead time, predictable volumes, where ninety-day terms were a reasonable reflection of the cash conversion cycle. That model is changing. The payment architecture has not kept pace.

5.2 Why It Matters Now

The Brand Reality

Agility changes the shape of suppliers' financing needs. In absolute terms, the working capital tied up in a single large-batch order is far greater than that required for small, frequent runs: fronting a 10,000-piece order commits roughly a hundred times the capital of a 100-piece order. What agility asks of suppliers is different: not larger single commitments, but a faster operating cadence—smaller batches, quicker changeovers, shorter lead times, pre-positioned materials, reserved capacity, and the digital readiness to support them. That model only works when payment cycles shorten to match the speed of production. The on-demand operators examined in our research pair order sizes of around 100 pieces with weekly or 30 days payment, compared with conventional payment terms that often stretch to sixty, ninety or even 120-plus days after shipment. The incompatibility is clear.

A factory borrowing at double-digit interest rates while waiting months for payment cannot be expected to fund the operational, digital, and sustainability capabilities the brand is simultaneously demanding.

The payment architecture and the operational ambition are pulling in opposite directions.

This is no longer only a commercial question. Purchasing practices have moved inside the regulatory perimeter: the EU Corporate Sustainability Due Diligence Directive names them as part of a brand's due diligence obligation, requiring that procurement strategies do not contribute to adverse human-rights impacts. Payment terms that concentrate financial pressure on suppliers in ways that drive wage violations or unsafe conditions are no longer a private arrangement; they create compliance exposure.

Country/policy rate	2021	2022	2023	2024	2025	Latest (Jun 2026)
MANUFACTURING / SOURCING MARKETS						
India <i>Repo rate · RBI</i>	4.00%	6.25%	6.50%	6.50%	5.25%	5.25%
China <i>1-yr Loan Prime Rate · PBoC</i>	3.80%	3.65%	3.45%	3.10%	3.00%	3.00%
Bangladesh <i>Repo rate · Bangladesh Bank</i>	4.75%	5.50%	7.75%	10.00%	10.00%	10.00%
Vietnam <i>Refinancing rate · SBV</i>	4.00%	6.00%	4.50%	4.50%	4.50%	4.50%
Indonesia <i>BI-Rate (7-day RR) · BI</i>	3.50%	5.50%	6.00%	6.00%	4.75%	4.75%
CONSUMER / RETAIL MARKETS						
France (euro area) <i>ECB deposit facility rate</i>	-0.50%	2.00%	4.00%	3.00%	2.00%	2.25%
United Kingdom <i>Bank Rate · BoE</i>	0.25%	3.50%	5.25%	4.75%	3.75%	3.75%
United States <i>Fed funds target, upper · Fed</i>	0.25%	4.50%	5.50%	4.50%	3.75%	3.75%

Sources: Reserve Bank of India; People's Bank of China; Bangladesh Bank; State Bank of Vietnam; Bank Indonesia; European Central Bank; Bank of England; US Federal Reserve. Compiled from official monetary-policy statements; verified June 2026.

The Manufacturer Reality

The transmission from financial pressure to workers' wallets is direct. Deferred payments and unpredictable order volumes make it harder for a factory to commit to higher or stable wages. When an order is cancelled on goods already produced, the financial loss lands directly on the supplier's balance sheet and potentially on the income of its workers. Wage commitments become contingent on the brand paying on time.

One supplier interviewed for this research makes the mechanism concrete. The Guangzhou-based manufacturer, which runs 400–500 workers across seven workshops, reported moving most of its business onto a demand-driven partner after 2020. Management did not describe the draw as price. It was the payment architecture: small, frequent orders paired with payment on an approximately weekly cycle, compared with traditional brand contracts where terms had extended beyond 120 days.

Management described fast payment as the strongest driver of supplier loyalty and linked it to morale on the factory floor.

The example should not be read as universal, but it shows why payment timing can matter as much as unit price. For a supplier, liquidity is not an administrative detail; it is the condition that allows the factory to keep paying, planning, and investing.

It is estimated that fewer than 2% of global garment workers earn a living wage, a figure cited consistently by industry bodies and labour rights organisations³⁴. The gap between minimum wages and living wages across 31 key garment-producing countries sits at 41 percent³⁵. In Bangladesh alone, compliance violations related to wages and hours doubled between 2023 and 2024, according to Qima audit data³⁶.

These are symptoms of a financial architecture that imposes liquidity and demand risks on the actors of the supply chain least able to absorb them.

³⁴ Clean Clothes Campaign, "Why: The Issues," <https://cleanclothes.org/fag/why>. Also cited by multiple labour rights organisations.

³⁵ The Industry We Want / WageIndicator Foundation, "Wages," 2025.

³⁶ QIMA, "2024 Q4 Barometer," October 2024.

5.3 Success Factors

Among the conditions that most consistently determine whether a brand can operationalise this principle, four stand out.

- > **Payment term redesign as a deliberate decision, not a default:**

Most brands inherit payment terms from legacy contracts and procurement policies that were never designed with supplier financial health in mind. Changing them requires an active decision at the CFO and supply chain leadership level that most organisations have not yet made. Furthermore, while the first-order impacts of long payment terms can appear in the cost of goods, they are difficult to detect and even harder to disentangle. Second-order impacts of long payment terms can force suppliers to cut corners, or, worse, become insolvent, undermining the very supplier capabilities the brand is trying to build.

- > **Ensuring payment pass-through upstream:**

Improving payment terms to Tier 1 without mandating and monitoring pass-through to Tier 2 fabric mills and component suppliers moves the problem one tier downstream rather than resolving it. The financial architecture needs to extend to where the working capital constraint actually binds.

- > **Order stability and risk-sharing:**

Payment terms address liquidity, but they do not address the demand risk that small-batch agility can otherwise transfer to suppliers. In practice, that risk can be reduced through rolling, staggered purchase orders that replace one large commitment with a frequent cadence of smaller ones. It can also be reduced through committed base volumes that anchor part of a supplier's production load in stable, long-term orders. The operators we studied anchor roughly 40% of supplier load this way, with responsive demand flexing on top. Additional mechanisms include shared markdown or clearance liability, and inventory-sharing arrangements that allow unused materials to be redeployed rather than left on the supplier's balance sheet.

- > **Commercial incentive alignment:**

Fair value distribution only compounds when suppliers trust that investment will be rewarded. The most developed models make this explicit: supplier performance is rated, and higher-rated suppliers receive tangible advantages like faster payment, priority for additional orders, and preferred access to future programmes. This converts fair value from a static policy into a dynamic incentive, in which reliability and capability are continuously rewarded with the cash flow and order stability that let a supplier invest further.

5.4. Where Are You Today

The financial relationship between a brand and its supplier base sits somewhere on a spectrum between transaction and partnership. Most brands are closer to the transactional end than they acknowledge, not by intent but through inherited terms and risk allocation from a production model that no longer reflects what the industry is trying to become.

Position 1: Standard Terms, Unexamined

Payment terms of sixty to ninety days. Order cancellation provisions that protect the brand. Price pressure applied at renewal. No visibility into whether Tier 1 payment terms cascade to Tier 2. The financial architecture was set years ago and has not been revisited. Worker welfare monitoring exists on paper, but verification is limited.

Position 2: Selectively Redesigned

Key strategic suppliers receive improved terms, sometimes through dynamic discounting or supply chain finance. Order stability commitments exist for core relationships. Some brands mandate payment term pass-through to Tier 2. Worker welfare audits extend meaningfully below Tier 1 for high-risk categories.

Position 3: Integrated Financial Partnership

Payment terms are matched to the working capital cycle of the production model the brand is asking suppliers to run. Real-time demand visibility and fast payment combine to create a financial feedback loop, enabling suppliers to invest because they can both predict and are paid. Order stability, transparent pricing, and shared risk mechanisms are embedded in the commercial relationship. Worker welfare monitoring is continuous rather than periodic.

6. CLIMATE ADAPTATION & MITIGATION

6.1. Climate Is Already Here

Fashion supply chains are built across geographies with high physical climate and water vulnerability. More than half of global cotton production takes place in regions already facing high or extremely high water stress, including major producers such as China, India, Pakistan, and Turkey³⁷. The dyeing and finishing clusters that process a large share of the world's fabrics are concentrated in river basins in South and Southeast Asia—for example around Dhaka, Narayanganj, and Gazipur—where intensive groundwater extraction and untreated effluent discharges are driving both chronic water scarcity and severe river pollution, and where regulators are under growing pressure to tighten discharge controls³⁸. The manufacturing hubs that dominate global apparel exports—China, Bangladesh, Vietnam, Cambodia, India, and Pakistan—are simultaneously among the most climate-exposed production geographies worldwide, with China's Pearl River Delta identified as one of the regions most at risk from sea-level rise and storm-driven flooding, and Cornell/Schroders estimating that extreme heat and flooding could erase 22% of expected apparel export earnings (around \$65 billion) by 2030 across Bangladesh, Cambodia, Pakistan, and Vietnam alone³⁹.

In 2024, a combination of severe flooding and political disruption in Bangladesh caused cumulative losses to the garment export sector estimated at over \$400 million, compounding an already fragile operating environment⁴⁰. Heat stress in South Asian manufacturing facilities is already reducing productive working hours and increasing worker health costs. Drought-driven production shocks in key cotton regions are amplifying raw material price volatility to levels that stretch traditional procurement practices and risk models⁴¹, forcing brands and mills to rethink

how they hedge, source, and hold inventory⁴². The industry's bottom line is already affected by climate; it simply is not fully accounting for it.

Climate is material to supply chains in two ways. The first is adaptation: protecting production from a changing climate. The second is mitigation: reducing the supply chain's own contribution to it. They are usually run as separate agendas, but the central argument of this principle is that the supply-chain design choices addressing one frequently impact the other. A network diversifying geographies and freight modes can be both lower-emitting and less exposed to climate concentration risk. Adaptation and mitigation, designed together, can be reinforcing.

6.2 Why It Matters Now

The Brand Reality

Brands sit at very different stages of mitigation and adaptation. Mitigation has had over a decade of attention leading to increased measurement, science-based target-setting, and supplier engagement. Adaptation is emerging. A Cornell/Schroders analysis points out that the industry's effort has concentrated on mitigation while the physical risks of heat and flooding go largely unplanned for⁴³.

On mitigation, most large brands now measure Scope 3 and have set reduction targets. The more advanced run supplier decarbonisation programmes that include energy-efficiency support, emissions data collection, and renewable procurement. For many of these programmes, reach remains a key constraint. Engagement typically covers the largest Tier 1 suppliers and a small number of strategic Tier 2 suppliers, leaving much of the emissions base outside the programme.

37 World Resources Institute, "One-Quarter of the World's Agriculture Grows in Highly Water-Stressed Areas," 30 October 2013

38 United Nations University, Water Vulnerability in Dhaka, Narayanganj, and Gazipur Districts, Bangladesh: The Role of Textile Dye Production, 2025; World Bank, "How Dialogue is Shifting Bangladesh's Textile Industry from Pollution Problem to Cleaner Production," 2017.

39 Supply Chain Dive, "China manufacturing hubs face greatest risk from sea-level rise," 25 February 2020, citing Verisk Maplecroft; Cornell University Global Labor Institute / Schroders, Higher Ground? Fashion's Climate Breakdown, 2023/2024; Arab News / Reuters summary of the same estimate, 12 September 2023.

40 VOA News, "Post-revolution Bangladesh protests cost garment industry \$400m," October 2024.

41 International Cotton Advisory Committee (ICAC), "Flooding and Drought Are Having a Crushing Impact on the Global Cotton Industry," September 2022 / updated 2025. The ICAC notes that extreme flooding in Pakistan and severe drought in West Texas and Brazil have cut output by over 1 million tonnes in the U.S. alone, and warns that this confluence of weather events is driving continued volatility and potentially higher prices throughout the 2022/23 season, with an A Index forecast range of 99–157 cents per pound.

42 Market Prospects, "2025 Cotton Price Surge: Textile Strategies & Analysis," August 2025, which argues that recent cotton price spikes have "challenged traditional procurement strategies" and recommends expanded use of futures hedging, flexible procurement plans, diversified sourcing, and tactical inventory builds to manage commodity volatility.

43 Cornell University Global Labor Institute & Schroders, Higher Ground? Fashion's Climate Breakdown (Reports 1–2), 2023–2024.

On adaptation, most of the industry remains at early stages. At best, companies have run a double-materiality assessment for CSRD, identifying physical climate risk in broad terms, not yet mapped to specific facilities and responses. Meanwhile, the exposure is increasingly quantified: Cornell/Schroders estimate \$65 billion of export earnings at risk across Bangladesh, Cambodia, Pakistan, and Vietnam by 2030, and a diversification trap in which the geographical moves to adapt to changing tariff regimes concentrates production in some of the most climate-vulnerable geographies in the world⁴⁴. Adaptation is where the exposure grows the fastest, and the maturity is the lowest.

The Manufacturer Reality

The factories absorbing climate disruption are the same factories already absorbing the financial pressure of the forecast-driven model.

The Lancet Countdown estimates Bangladesh lost some 29 billion potential labour hours to heat exposure in 2024, a 92% rise on the 1990s average, with associated income losses near \$24 billion, around 5% of GDP⁴⁵.

Cornell's Global Labor Institute, Better Work, and field documentation by HeatWatch and others converge on the same reality at the facility level: rising indoor heat is materially cutting output, with meta-analyses and WHO estimates suggesting worker productivity declines by around 2–3% for every degree increase in wet-bulb temperature above thermal comfort thresholds⁴⁶, increasing illness and fainting episodes, and reducing the take-home pay of workers who lose shifts they cannot afford to lose.

On mitigation, manufacturers' constraints are structural. In several major producing geographies, the grid is unreliable and fossil-heavy: power cuts both worsen factory heat and make clean electrification difficult, since factories rely on diesel generators to buffer the grid's unreliability. And the cost of capital is high: a manufacturer borrowing at local double-digit interest rates cannot easily fund the energy-efficiency and renewable capex decarbonisation requires. In addition, investment is often directed toward the highest-return use of capital, which may be new machinery or capacity expansion rather than decarbonisation. Brand-backed sustainability-linked finance can bridge this gap by lending the brand's stronger credit to fund supplier decarbonisation, but such programmes remain limited in absolute scale.

Adaptation is currently driven by operational economics rather than by the people most exposed to the heat. Cornell's case studies of early-adapting factories in Vietnam and Cambodia find that managers frame heat abatement as an operational decision, weighing the cost of cooling against steadier output, fewer disruptions, and more reliable machine performance as automation deepens. Where that operational case is strong, factories invest. Where it is weak, worker exposure persists, because output and worker welfare do not always align. Workers bear the physical and wage cost directly through lost shifts, illness and fainting documented in recent field research.

⁴⁴ Cornell GLI & Schroders note that apparel brands and retailers that shift sourcing away from current climate-vulnerable hubs "would struggle to build the large-scale capacity they benefit from in South and Southeast Asia," warning that many production alternatives are themselves highly exposed to physical climate risks. Their investor toolkit also cautions investors against "cut-and-run" approaches that relocate production without addressing climate risk, as this risks simply redistributing exposure rather than reducing it.

⁴⁵ Lancet Countdown on Health and Climate Change, Health and Climate Change in Bangladesh—Data Sheet 2025, October 2025.

⁴⁶ Cornell University Global Labor Institute & International Finance Corporation (with input from ILO Better Work), The Heat is On: How Heat Stress Impacts the Apparel Industry, Jobs and Worker Wellbeing, 2025; Yeasmin, F. et al. (Griffith University / HeatWatch consortium), "Heat impacts on health and productivity: the case of two ready-made garment factories in tropical Bangladesh," 2025; World Health Organization, "Workplace heat stress—Questions and Answers," 21 August 2025, and WHO/WMO Technical Guidance on Workers and Heat Stress, 2025.

6.3 Success Factors

Among the conditions that most consistently determine whether a brand can operationalise this principle, three stand out.

> **Geographic and modal diversification as a compounding lever:**

Nearshore capacity for key markets and a managed modal split between sea and air freight reduce tariff exposure and transport emissions simultaneously, while improving climate resilience. Furthermore, nearshoring serves as a necessary risk mitigation strategy in scenarios where securing capacity or transporting goods from offshore locations proves challenging. The brands that treat nearshore and offshore as separate workstreams are missing the compounding return of solving them together.

> **Supplier investment support:**

For both mitigation and adaptation, including factory cooling, energy efficiency, and renewable energy procurement. They are the conditions that can unlock progress for the industry's decarbonisation while maintaining supplier productive capacity as climate conditions intensify.

> **Facility-level climate visibility, built into sourcing:**

Shifting from a one-off risk assessment to a facility-level, regularly updated view of exposure—heat, water, and flooding—that becomes a standing input to supplier selection and network design decisions alongside the traditional variables of cost, lead time, and quality. Without this, the supply chain will continue to be optimised without consideration for changing conditions.

6.4 Where Are You Today

Maturity runs on two axes—mitigation and adaptation. Most brands are further along on mitigation, where measurement and target-setting are well established, than on adaptation.

Position 1: Committed

Mitigation: ambitions and targets are set, and Scope 3 emissions are tracked, but supplier engagement is limited and little direct decarbonisation work is happening down the value chain. The measurement is in place, but the actions on the ground remain limited.

Adaptation: limited visibility into physical climate risks that have not been systematically assessed or mapped to the supply base.

Position 2: Engaged

Mitigation: direct engagement with key suppliers is underway, often as part of an industry coalition supporting energy-efficiency and decarbonisation work across factory networks.

Adaptation: a climate risk assessment or double-materiality assessment has been completed and key vulnerabilities identified through climate-scenario analysis. The company has a broad view of where its exposure concentrates.

Position 3: Dynamically Integrated

Mitigation: the brand has developed deep decarbonisation partnerships, including financing structures that lend its stronger credit rating to fund supplier decarbonisation while deploying multi-year agreements enabling the supplier to access better capital.

Adaptation: dynamic, live monitoring of supply-chain climate risk—integrated with facility-level analytics and exposure mapping—feeding sourcing and network decisions in real time.

Where to start

Most brands are further behind on adaptation than mitigation, and the exposure there is larger and less hedged. The practical first move is to take adaptation from a broad assessment to a facility-level view of the highest-risk nodes, then take action to reduce those risks, while supporting factories with adaptation planning and investments. On mitigation, the step that unlocks the most is moving from tracking to direct, financed engagement with the suppliers that carry the largest share of the footprint.

7. ECOSYSTEM BUILDING

7.1 From supply chain to ecosystem

A garment supplier examined in this paper's field research, working within SHEIN's demand-driven model, described growing fortyfold in five years⁴⁷. The growth was not driven by capital investment or technology alone, but by the nature of the relationship: payment within days, not months, freed working capital for reinvestment, demand visibility through a shared platform replaced speculative production with informed decisions, and a performance rating system tied order priority and payment speed to execution quality. The supplier knew what was expected, what it would receive in return, and had a strong commercial incentive to deliver.

A supply ecosystem operates like a network that compounds capability.

This distinction is not semantic. A supply chain moves goods from production to market. An ecosystem coordinates the people, capabilities, information, and incentives that determine how well that movement happens. Because it coordinates rather than standardises relationships, an ecosystem can contain suppliers with very different economics, some built for speed and some for scale, provided the system is designed to accommodate different models.

A supply chain can be outsourced. But an ecosystem has to be built, with governance, investment, and a clear understanding of what each participant needs to perform at the level the system requires.

Demand-driven operations, digital infrastructure, transparency, fair value, circularity, and climate resilience all depend on a supplier network capable of delivering on them. That capability does not emerge from procurement alone; it is built through the architecture of the ecosystem that surrounds it.

7.2 Why It Matters Now

The Brand Reality

Most brands manage suppliers transactionally: they place orders, set specifications, conduct audits, and measure on-time delivery. The relationship is structured around fulfilment—designed to ensure compliance with requirements rather than invest in the capability to exceed them.


The field research conducted for this paper points to a different model. Within the network examined, investment in the supplier base is treated as a strategic requirement, not a goodwill gesture: educational programmes run across thousands of factories, best practices are shared network-wide, and equipment and software pricing is negotiated on suppliers' behalf. But the point made most repeatedly concerned the order of operations. The ecosystem was not built by deploying a platform; it was built by establishing the operational foundations a platform could then accelerate.

The commercial return is not measured in goodwill. It is measured in a supplier base that performs at a speed and consistency that order placement alone cannot produce. The brand's task, though, is not to push a single operating model onto every supplier but to build a collaborative system—process, information flows, incentives—that enables the integration of suppliers with different economics, each performing according to SKUs' risk profile.

⁴⁷ Primary interview conducted for this paper, 2026.

Supply Chain Data Network

One brand | the platform | many Tier 1 (cut & sew) and Tier 2 (fabric) suppliers | data flows both ways



The Manufacturer Reality

From the supplier side, the ecosystem offers something that price alone cannot: a system that rewards execution with more orders, faster payment, and clearer incentives than a simple purchase order can provide.

But the same system that generates loyalty can generate risk. A supplier flagged single-brand dependency as the primary risk. A manufacturer that has restructured its entire operation around one brand's platform, payment cycle, and production requirements is both highly capable and highly exposed. The way to reduce that exposure is to serve more than one brand. However, this risk mitigation strategy only materialises if the platform has built-in interoperability. Enabling the supplier to diversify their portfolio without having to run a separate, incompatible information system and process for buyers.

Because SKU uncertainty is a continuum rather than a binary, the ecosystem has to optimise for two things at once: flexibility and fast reaction at one end, scale and efficiency at the other. The first favours suppliers whose advantage is responsiveness—typically smaller, more adaptable operations. The second favours larger suppliers, which compete on a different dimension, with margins derived from efficiency and dependent on long runs and stable, sizeable orders.

A collection spans both kinds of demand, so a mature supply base has to map to the range: responsive capacity at the uncertain end, efficient capacity at the stable end. This is the supply-side counterpart to the demand-side scoping in Principle 1: agility is not a silver bullet to be applied across an entire collection, and the ecosystem has to be designed to accommodate both kinds of manufacturer, not only the agile kind. However, it is worth noting that a large manufacturing group with multiple facilities can feasibly, with the right processes and technology, serve both ends of the uncertainty spectrum, albeit in separate facilities.

The same system can accommodate both, because even suppliers optimised for long, stable runs benefit from the ecosystem's information and data flows: demand sensing does not turn a large supplier into an agile one, but it lets them plan against a sharper signal, cut overproduction on the stable portion of the portfolio, and carry less of the inventory risk a forecast would otherwise leave on their books.

The ecosystem's performance is also inseparable from the market it sits in. The field research was explicit that the speed of the model rests on a fabric-access economy—immediate availability of materials at the right weight, finish and price, without minimum-order or lead-time constraints. That liquidity is what makes a fifteen-day concept-to-first-order timeline physically possible.

7.3 Success Factors

> **Incentive architecture designed, not inherited:**

The structure around payment terms, order allocation, performance ratings, and capability investment needs deliberate design. The brands operating at Position 3 did not arrive there by accepting the commercial terms their industry inherited. They made deliberate decisions about what the ecosystem needed each participant to do and built the incentive architecture that made taking those actions the most advantageous choice.

> **Geographic and market infrastructure awareness:**

An ecosystem does not operate in isolation from the market infrastructure around it. Fabric-access economics, cultural proximity, supplier mindset and regional manufacturing maturity all shape what an ecosystem can achieve in a given geography. Building nearshore ecosystems means understanding and investing in the market conditions that make ecosystem performance possible, beyond the relationships themselves. Ultimately, the rules that apply when entering a new consumer market apply equally when entering a manufacturing market.

> **Supplier-centric technology platform:**

A structural ceiling exists where brand-mandated platforms serve primarily to export administrative and data-entry burdens upstream. Genuine ecosystem performance requires a departure from this transactional logic toward a supplier-centric architecture. To flourish, the digital layer must transition from a reporting tax into a functional asset, delivering the planning tools and reciprocal visibility needed for manufacturers to optimise their own day-to-day operations rather than merely satisfying a buyer's requirement.

7.4 Where Are You Today

Most brands have supplier relationships. Far fewer have supplier ecosystems. The distance between those two things is not measured in technology investment or platform sophistication. It is measured in the deliberateness of the architecture that surrounds the relationship.

Position 1: Transactional and Fragmented

Suppliers are selected on price and capability, managed through purchase orders and audits and measured on delivery compliance. Payment terms are set across suppliers. Demand visibility does not flow beyond the brand's own planning systems, and capability investment is the supplier's responsibility. When performance falls short the response is a corrective action plan, not a structural question about what the relationship was designed to produce. Most brands sit here.

Position 2: Selectively Integrated

A defined subset of strategic suppliers operates within a more deliberate architecture: demand signals flow through a shared platform, payment terms have been redesigned for priority relationships, and capability programmes exist. The brand has begun treating supplier performance as a function of the ecosystem it has built around those suppliers rather than the suppliers' capability alone. This is where the most progressive brands sit today. The ecosystem logic is real within its boundaries.

Position 3: Governed and Compounding

The full supplier network operates within a deliberate and consistently applied architecture. Roles are defined by what each supplier is built to do—responsive and efficiency-led suppliers are held to different operating expectations rather than one—while demand signals and data flow to the entire base, so all suppliers plan against the same live picture. A performance rating system ties payment speed, order priority, and commercial incentives directly to execution quality. Capability investment flows across the network, not just to strategic relationships. Demand visibility moves continuously in both directions.

The ecosystem compounds because every element of the architecture is anchored in a supplier-centric technology platform. This platform transforms digital tools from a reporting burden into a functional asset, providing suppliers with the planning capability and real-time visibility they need to optimise their own day-to-day operations. This shifts the relationship from transactional compliance to functional partnership, where the digital infrastructure actively improves the supplier's efficiency rather than merely satisfying a brand's administrative and functional requirements.

Where to start.

The architecture that performs at speed in one geography was built over years, inside a mature fabric market and with a supplier base that grew within the ecosystem from the start. Replicating it elsewhere means building the underlying market infrastructure alongside the governance—a longer undertaking than deploying a platform or redesigning payment terms alone. The entry point is therefore not the technology but the foundations: define the roles and performance standards, align payment with operational reality, and choose a geography whose fabric market can support the speed the ecosystem is meant to deliver.

The Principles as a System

The seven principles described in this paper are not independent levers. Each one amplifies or constrains the others. Demand-driven operations require digital infrastructure to function. Transparency and traceability make climate risk visible and circularity possible. Fair value distribution is the financial condition under which suppliers can build the capabilities the other principles depend on. Ecosystem building is the architecture that allows these principles to compound across an industry rather than within a single brand's supply chain.

The following section maps those interdependencies.



INTERDEPENDENCIES BETWEEN THE PRINCIPLES

A Universal Supply Chain Illustration

In 2021, the global automotive industry lost an estimated 210 billion dollars in revenue and 7.7 million vehicles of planned production because it could not source sufficient volumes of semiconductor chips⁴⁸. Automakers had cancelled chip orders at the start of the pandemic anticipating falling demand. Chipmakers reallocated that capacity to consumer electronics where demand was surging. When automotive demand recovered, the capacity was gone. Chip lead times that had run at three to four months stretched to record

lengths, over twenty weeks at its peak⁴⁹. The failure was not forecasting alone. Automakers lacked visibility beyond Tier 1, had no shared early-warning mechanisms, and could not rapidly re-coordinate capacity across their supply network.

The principles this paper describes are not a fashion-industry-specific framework. They are a supply chain resilience framework, and the semiconductor shortage demonstrated that their absence is a structural exposure for any industry still built on the traditional model's founding assumption: the operating environment is stable and predictable.

The Five Interdependencies

1

The Digital-Agility Nexus

For demand signals to create value, they must reach the actors able to respond at the right time. Digital infrastructure without an operating model that acts on them produces dashboards and reports rather than decisions.

2

Transparency as the Prerequisite

Transparency and traceability form the principle that makes every other principle possible to verify and operationalise. Climate exposure, material provenance, supplier working conditions and constrained capacity cannot be managed if they remain invisible below Tier 1.

3

The Fair Value-Agility Loop

Fair value distribution and demand-driven operations create a mutually reinforcing dynamic: faster changeovers, reserved capacity and digital readiness require investment. Payment terms, order stability and access to capital determine whether suppliers can invest. In turn, better-capitalised suppliers are more able to deliver the responsiveness the model requires.

4

Circularity as Climate Enabler

Producing fewer unwanted products prevents both economic waste and the associated environmental footprint. Recovered materials can reduce exposure to volatile virgin inputs while lowering the environmental impact of new products.

5

The Regulatory Convergence

Regulatory compliance and operational transformation increasingly rely on the same underlying data: supplier identity and location, material composition and provenance, production status and cross-tier visibility. The regulatory convergence therefore does not need to sit as a separate burden on top of the operating model. Much of the infrastructure built to make the supply chain more responsive can also support traceability, due diligence and reporting requirements, allowing the two agendas to be pursued as one programme.

⁴⁸ AlixPartners, "Semiconductor shortages to cost the auto industry \$210 billion in revenues," September 23, 2021.

⁴⁹ Susquehanna Financial Group's monthly lead-time tracker, as reported by Bloomberg, shows that average chip lead times rose from pre-pandemic norms of around 10–15 weeks to 18 weeks in May 2021 and 20.2 weeks in July 2021, the longest since tracking began in 2017.



THE OPERATING DECISIONS LEADERS MUST OPTIMISE

Transitioning to a more agile, demand-driven model is not a single trade-off to weigh, but a set of operating decisions to optimise.

Each decision pulls on several axes at once. The financial axis spans working capital requirements, inventory outstanding, and the margins a business protects. The environmental axis spans the emissions avoided through lower overproduction and those incurred through freight. The operational axis spans where production sits, how capacity is reserved and how goods move.

These axes rarely move in isolation. A decision that improves one can pull against another, and the same choice can read differently depending on the product, segment, and supplier network behind it. What follows takes the principal decisions in turn: what each optimises across, where the tensions sit, and how to navigate them. The aim is not to treat agility as a blanket switch, but as a deliberate operating model that can be calibrated by product, geography, and risk.

A note on scope: This analysis focuses on upstream and outbound transport emissions, manufacturing emissions, and Tier 2 working capital dynamics. Reverse logistics is excluded because return rates are driven by business-model choices and consumer behaviour beyond this study's scope. In this section, "speed model" refers to all-air, direct-to-consumer responsiveness; "agile model" refers to differentiated SKU-level routing; and "traditional model" refers to the planned, forecast-driven model.


Optimising Inventory, Carbon and Capital

Producing close to demand in small, responsive batches can significantly reduce overproduction and the inventory that weighs on margins and immobilises capital. But when responsive production is combined with long-distance fulfilment, the carbon cost can shift from inventory to freight. Air freight is roughly 47 times more carbon-intensive per tonne-kilometre than sea freight, making it the clearest trade-off in this report between the business gain and environmental cost of responsiveness.

The point is not that agility inevitably increases transport emissions. It is that agility has to operate as a continuum of decisions rather than a binary model choice. A product's demand uncertainty should determine the production and logistics framework it uses. Stable, low-risk SKUs can be planned further ahead, produced in larger batches and shipped by sea. Higher-risk SKUs can use smaller initial batches, later commitments, and faster replenishment where the commercial case justifies it. The strategic question is therefore not whether to choose speed or efficiency, but how to optimise across inventory, margin and carbon at SKU level.

Agility and transport emissions

► The agile trade-off: leaner inventory, higher transport carbon



Evaluating operating models across inventory efficiency and transport carbon shows the trade-off clearly. Among established players, the relationship is broadly inverse: as inventory gets leaner, transport carbon tends to rise. The planned-mass model holds the most inventory and emits the least in transport. The mixed model holds materially less inventory while emitting several times more. The fully responsive model extends the line to its extreme: the lowest inventory of all and, by a wide margin, the highest transport intensity. A precise inventory figure is not disclosed by privately held agile operators, so this position is estimated from the expert interviews conducted for this paper.

The mixed model confirms the case for managing uncertainty through differentiated supply chain models. It holds near-lowest inventory while earning the strongest operating margin in the set, and emits a fraction of the fully responsive model's transport carbon.

Agility is therefore better read as an optimisation problem than as a compromise between two models. The inventory–margin relationship established earlier in this paper now carries a third dimension—carbon—and with it the real question:

not which model to adopt, but how to optimise across inventory, margin, and carbon as a shared objective.


Does reduced overproduction offset the air-freight penalty?

The optimisation question can invite a shortcut: if leaner inventory and less overproduction are the prize, why not move fully to a responsive, air-served model and assume that the increase in transport emissions will be offset by the inventory reduction it enables?

Our modelling suggests the answer is more complex. We map model-driven carbon—overproduction plus transport, with manufacturing of goods sold held constant—across a range of air-freight intensity and overproduction outcomes. The result is not a precise break-even line, but a navigator. Because the inputs carry real uncertainty, the model is best read through three zones: a zone of comparability, where carbon impact is not meaningfully distinguishable within a ± 40 tCO₂e/€m band; a clear-benefit zone, where responsiveness lowers carbon; and a higher-impact zone, where freight emissions outweigh the inventory benefit.

Shipping - unsold trade-off navigator

An orientation tool, not a precise benchmark—it shows roughly where a model sits on the air-freight and overproduction scales, and which way to move. The navy outline marks ± 40 tCO₂e/€m zone of comparability; colour deepens with magnitude beyond it.



Baseline = planned reference (transport 30—the sea-weighted average of listed planned operators—and 30% more produced than sold).
Cell = (transport - 30) - 240 × (0.30 - μ). Manufacturing intensity 240 tCO₂e/€m, the revenue-weighted average across the listed sample. Illustrative.

The shipping–unsold trade-off navigator: an orientation tool showing where operating models sit across air-freight intensity and overproduction, and which direction they need to move. Please refer to methodology in the Appendix.

Two features matter most. First, for carbon the freight choice is the larger lever: across the plausible range, transport intensity spans roughly two-and-a-half times the range attributable to overproduction. Second, a position that relies heavily on air freight sits firmly in the higher-impact zone, while a position that pairs moderate transport with materially lower overproduction sits within the zone of comparability on carbon—its advantages lying in margin, working capital, and waste reduction rather than in emissions.

A clear carbon benefit is unlocked by low overproduction, and the leaner a model runs the more freight it can carry and still net one: at around 5% overproduction a clear benefit holds even at a transport intensity of ~50 tCO₂e/€m—

well above a planned operator's—and closes out only above ~57. The benefit is therefore available to any model lean enough to earn the freight it uses, not only near-sea ones.

Narrowing the carbon gap without losing agility

Four levers compound to move a speed-driven approach down the carbon axis toward the efficient corner without surrendering all the benefits that responsiveness provides: modal shift changes how goods move, geographic diversification changes how far they move, volume consolidation improves the efficiency of each shipment, and sustainable aviation fuel reduces the intensity of the air freight that remains. Applied together their effect is multiplicative.

➤ **Lever 1—Modal split:** route stable, predictable SKUs via sea freight while reserving air freight for high-uncertainty items.

Modal split (sea / air)	Blended transport emissions intensity (tCO ₂ e/€m)	Reduction vs. all-air
0% / 100% ⁵⁰	243	—
30% sea / 70% air	176	–28%
40% sea / 60% air	153	–37%
50% sea / 50% air	131	–46%
80% sea / 20% air	~65	–74%

Source: Source: Hyran analysis. Blended intensity = (sea share × sea-freight EI) + (air share × air-freight EI), computed from the planned-mass (sea) and fully air-served (air) endpoints used in the navigator. Illustrative; actual reductions depend on route, load factor and product weight.

➤ **Lever 2—Geographic diversification:** producing closer to consumer markets can shorten the distance component of transport emissions regardless of mode. This is a distinct lever from modal split and compounds with it.

Strategy	Application	Transport effect
Nearshore for key markets	Turkey/Morocco → EU, Mexico/ Central America → US, SE Asia, and China → APAC	Can reduce distance 50–70% for served markets and make sea/rail viable where air was previously needed for speed, provided key upstream inputs are also regionally accessible
Offshore hub for core/stable	Retain offshore production for predictable, high-volume basics where unit economics dominate	Longer distance but sea freight; low emissions intensity (EI) per unit
Dual-source for high-uncertainty	Nearshore capacity for initial test runs + fast replenishment; offshore for scale-up once demand is confirmed	Air freight only for first batch (small); bulk replenishment via sea from offshore

⁵⁰ Modelled as fully air-served for illustration; in practice the leading on-demand operators already move some stable, predictable volume by sea and are actively expanding that share.

However, if fabric, trims, dyeing, or finishing still have to travel long distances, the lead-time gains associated with locating final assembly closer to market can be eroded. Responsiveness depends on sufficient supplier density and upstream capability around the regional manufacturing base. Onshore capability can play a targeted role at the front end: locating design, development, and small test runs close to the regional markets can shorten the time from concept to demand signal, before winning products are scaled through the production configuration that best fits their economics.

The compounding logic is significant: modal shift alone (40–60% sea) cuts transport EI by 37–46%. Geographic diversification independently reduces distance-driven emissions. Combined, a nearshore production base serving 30–40% of volume—with the remaining offshore volume on sea freight—could narrow the gap to mixed-modal peers substantially, without sacrificing speed for uncertain items. The exact reduction depends on the geographic mix, which requires modelling per market.

Our field research shows this lever is real but hard-won. The binding constraint is the agile ecosystem itself: replicating the supplier density and upstream capability the model depends on across materials, processing and manufacturing takes time to build in a new geography.

➤ **Lever 3—Volume consolidation and shipment efficiency:** within a DTC air freight model, emissions per unit can improve when more items are packed per parcel and when shipments are consolidated before dispatch.

Several mechanisms are available. At the order level, higher average order values (reportedly \$100–125 per transaction for some DTC agile operators, implying 8–12 items per parcel) mean

that emissions are shared across more units per shipment. A parcel containing 10 lightweight garments is not materially heavier than one containing 2–3. Packaging optimisation (lighter materials, higher packing density) further improves the weight efficiency of each shipment.

At the logistics level, consolidation at origin—batching individual orders into larger shipments before air dispatch—reduces the number of individual handling events and improves cargo load factors. And the emerging regional warehouse model (where proven, high-velocity SKUs are pre-positioned in regional hubs while new/uncertain styles ship from origin) creates a hybrid architecture that captures consolidation benefits for the most predictable portion of the range.

➤ **Lever 4—Sustainable Aviation Fuel (SAF):** for the air freight that is needed after modal shift and geographic diversification, SAF reduces the emissions intensity of each flight. Current SAF blends can reduce CO₂ emissions by 50–80% compared to conventional jet fuel.

Given its cost trajectory and emissions-reduction potential, SAF applies specifically to the share of volume that continues to move by air after modal split and nearshoring have been applied to reduce the per-flight penalty of remaining air shipments. It also represents a potential pre-competitive collaboration angle between brands sharing air freight corridors, which could collectively negotiate SAF procurement at more favourable terms while creating the demand signal that producers need to justify capacity investment. At current production volumes and cost levels, SAF is more credible as a medium-term lever (2028–2035 horizon for meaningful scale).

Factor	Current state (2025)
Global SAF production	~1.9 million tonnes (~0.3% of total jet fuel)
Cost premium	~2.8× conventional jet fuel (€2,085/t vs €734/t)
EU mandate (ReFuelEU)	2% blending from 2025, rising to 6% by 2030
Production trajectory	New facility announcements declined 50–70% in 2022–2023; IATA warns many airlines will revise 10%- by-2030 targets

Source: IATA, December 2025; EASA Annual Technical Report, 2025; BCG SAF Outlook, 2025.

Staple Products Versus Trend Products: Not All Uncertainty Is Equal

Demand uncertainty is not distributed equally across a product range. Core basics with established sales histories can often support longer lead times, larger batches, and sea freight. Trend-driven, seasonal or untested products carry greater forecast risk and benefit from small initial commitments, test-and-respond production, and faster replenishment. The operating model should therefore follow the uncertainty of the SKU, rather than a single brand-wide sourcing philosophy.

In practice, the balance will vary across business models. Illustratively, a basics-heavy retailer might route 60–70% of volume through the efficient track. An operator with high style turnover might reverse those proportions. The critical capability is the ability to classify and dynamically reclassify products as demand signals evolve: a product launched on the responsive track may migrate to the efficient track once its demand pattern stabilises.

This makes agility a portfolio-allocation discipline rather than a blanket speed strategy. Demand uncertainty determines three linked decisions: how much to commit, where to produce and how to move the goods.

Quality Ceiling and Capital Barriers at Tier 2

The speed model relies on Tier 2 suppliers (fabric mills) pre-positioning inventory so that Tier 1 garment factories can access materials rapidly when purchase orders or equivalent demand triggers materialise. Our research indicates that high-performing fabric suppliers in agile networks hold approximately 10 days of inventory for active materials. This works well for commodity and mid-range fabrics. It creates a structural barrier for premium materials.

An important caveat: this trade-off is segment-specific. At the mass and mid-market level, where the speed model is most widely deployed, the fabric quality differential between speed and traditional operators is minimal. A 180 GSM cotton jersey or a 40s cotton oxford shirt is the same fabric specification whether ordered by a €10 or a €25 brand. The working capital barrier becomes binding primarily when premium and luxury brands attempt to adopt speed principles.

The issue is straightforward: premium fabrics cost significantly more per metre than commodity fabrics. When a Tier 2 supplier is asked to pre-position 10 days of inventory, the working capital requirement scales linearly with fabric value.

It is worth noting that 10 days of fabric inventory can look drastically different depending on the types and volume of SKUs it serves. A SKU that is selling well may require 30,000 yards of fabric for 10 days versus a SKU that may be selling less well with a requirement of 4,000 yards.

► Tier 2 working capital burden by fabric quality

Fabric category	Relative cost per metre	Capital intensity
Commodity poly/cotton blend	1× (baseline)	Manageable
Mid-range woven	2×	Moderate
Premium wool/silk	4×	Significant strain
Technical performance fabric	5×	Challenging for most SMEs

Source: Hyran modelling. Assumes a Tier 2 supplier serving ~€50m in annual fabric demand. Working capital = 10 days × daily demand × relative cost multiplier. Note: premium and technical fabrics can reach 6–8× on specific programmes (e.g., high-spec outerwear, engineered knits), which would push 10-day inventory to €8–11m.

For an SME fabric mill, holding €5.5–6.8m in pre-positioned premium inventory represents a significant financing decision. At the upper end of the range, where specific programmes require fabrics costing 6–8× the commodity baseline, the exposure can reach €8–11m, potentially exceeding available working capital entirely. This can appear to create a quality ceiling for the speed model: the model operates well at commodity and mid-range quality levels, but encounters friction as fabric value increases into the premium and luxury range.

The burden also scales with the value of pre-positioned inventory. A 60- or 90-day payment term may be manageable on a €1m exposure, but becomes a severe strain on an €11m exposure; a mid-cycle change or cancellation transmits a correspondingly larger shock. The financing implication is the same as in the fair-value discussion above: the more volume and value a supplier is asked to pre-position, the more its liquidity depends on orders being honoured and payments arriving quickly. This is why a small-batch, fast-payment cadence can reduce supplier risk, provided it is paired with credible demand signals and clear commitments.

Mitigation strategies

Several approaches can lower this barrier, though none eliminates it entirely:

- > **Demand data sharing with Tier 2:** providing fabric mills with granular, forward-looking demand signals with required ex-mill dates for each tranche of fabric (as some agile operators already do) reduces the risk that pre-positioned inventory becomes obsolete. This lowers the effective cost of holding stock by reducing the required yarn inflow at any point in time and by minimising the write-off probability. But it does not reduce the upfront capital requirement.

- > **Fabric portfolio rationalisation:** concentrating speed-track production on a curated set of fabric qualities reduces the variety of fabrics a mill must hold, increasing the turnover velocity of each. This is the material equivalent of SKU rationalisation at the finished-goods level.
- > **Greige fabric postponement:** fabric mills typically use a form of postponement by holding undyed (greige) fabric and dyeing to order based on actual colour demand. This reduces colour-specific obsolescence risk significantly and is already widespread in the industry. For the speed model, formalising and extending this practice—with brands committing to greige fabric volumes while deferring high-risk colour decisions—can reduce the capital burden while preserving speed and responsiveness. Rise & Fall, profiled in the margin erosion section, uses precisely this strategy to manage premium materials within a speed framework.
- > **Tier 2 financing mechanisms:** brands or platform operators could facilitate supply chain financing that allows Tier 2 suppliers to access working capital at lower cost, using the brand's demand data and payment track record as collateral. This shifts the financing cost without shifting the operational model.

The practical implication is that speed models are not inherently limited to commodity products, but extending them to the premium segment requires deliberate investment in the financial and digital infrastructure of the Tier 2 supply base. For mass and mid-market brands, this trade-off is largely non-existent: fabric specifications are similar across models, and the working capital requirements appear manageable within existing structures.



CONCRETE APPLICATION OF THE PRINCIPLES

CONCRETE APPLICATION OF THE PRINCIPLES

The seven principles described above provide a framework for building more agile and adaptive supply chains. This section outlines how these principles translate into tangible business outcomes by exploring how an agile operating model can address three key challenges facing senior leaders today:

1. How can we reduce inventory risk and the associated working capital trap?
2. How can we protect gross margins from the pressures of markdown dependency?
3. How can we build a business model that is compliant with and benefits from the transition to a circular economy?

These are a reflection of financial markets' concerns, drawn from our analysis of investors' questions. The following section explores how the principles of agile supply chains can, and already do, help companies systematically solve these challenges while creating business opportunities. Each section is illustrated by case studies drawn from our research.

This section provides tangible elements that can be integrated into strategic planning and operationalised in the short term. Although transitioning to a fully agile supply chain can represent a daunting task, our aim is for the following section to give business leaders a practical view on where to start and what to monitor as the transition is implemented.

CROSS-CUTTING ENABLERS

Three capabilities are foundational to the applications that follow. First, multi-directional digital infrastructure: demand, inventory, and capacity data must move rapidly and seamlessly between the brand, Tier 1, and Tier 2 to inform live decisions. Second, supplier capability: responsiveness depends on suppliers having the processes, tools, and financial conditions to act on the signals they receive. Third, management incentives: buyers and planners need to optimise total system margin, balancing unit cost against markdowns, inventory, disruption risk, and working capital.

These capabilities can be built incrementally. The following sections show how they translate into practical decisions, where to start, and what to monitor.

AGILE SUPPLY CHAIN TO REDUCE INVENTORY

Excess inventory is one of the most visible and costly symptoms of a supply chain that is forecast-driven rather than agile. Across the industry, brands routinely produce 25–40% more than they sell at full price, absorbing the difference through markdowns, liquidation, or destruction. An analysis of over 36,036 investor questions from the past decade confirms that inventory overhang is the single most frequently raised concern, accounting for 12.9% of all problem-related topics, and its frequency has increased by 125% since 2020. The financial consequences are severe: excess inventory ties up working capital, generates warehousing costs and, increasingly, regulatory exposure as jurisdictions move to ban the destruction of unsold goods. The cost is also environmental. Every garment manufactured in excess of demand consumes raw materials, energy, water, and transport capacity that produce no return.

The root of the problem is a supply chain designed to absorb forecast error through overproduction rather than to limit it through agility. Traditional models, with their 6 to 12-month lead times⁵¹, force brands into a binary bet: over-order to avoid missing sales, or order conservatively and accept lost revenue. In practice, buyers rationally over-order because the penalty for a stockout is immediate and visible, while the cost of a markdown arrives months later and is diffused across the P&L. Historically, high intake margins and low production costs allowed brands to absorb that markdown risk. That buffer is narrowing as production, transport and tariff costs rise, and as geopolitical disruption makes landed cost less predictable. The result: the industry's aggregate exposure to excess stock has been estimated at \$70–140 billion annually⁵².

⁵¹ FashionUnited, "Everything about the (traditional) supply chain," October 2022

⁵² McKinsey & Company / Business of Fashion, State of Fashion 2025, 2025.

An agile supply chain addresses this at the root by structurally decoupling production volume from speculative forecasting and replacing it with a continuous cycle of small batches, real-time demand signals, and rapid replenishment. It reduces both sides of the inventory dilemma: the risk of overproduction and the fear of missing a sale. Rapid replenishment means brands can start small without permanently forfeiting upside. Even agile models have upper bounds to upside capture, however. Their ability to respond depends on material positioned at Tier 2 and capacity reserved at Tier 1. The aim is therefore not to forecast more accurately, but to reduce the consequences of being wrong.

How Agility Reduces Inventory: From Forecast to Signal

Agile supply chains reduce inventory exposure through four interconnected mechanisms. Each maps to the principles described earlier in this paper and represents a structural shift in how production decisions are made.

Traditional Model Challenge	Agile Enabler	Business Performance Outcome
Large upfront production commitments based on seasonal forecasts	Demand-Driven, Agile & Responsive Operations	Initial production runs as low as 80–100 units for unproven styles, with volume committed only after demand is confirmed.
Limited visibility into sell-through until end of season	Digital Infrastructure, Intelligence & Governance	Real-time point-of-sale data flows back through the supply chain, enabling automated reorder decisions for stable products and early identification of underperformers.
Siloed inventory data across brand and supplier	Transparency, Traceability & Compliance	Integrated information systems across the brand and its suppliers create a unified view of stock and demand, enabling more efficient allocation of production resources, reducing the risk of overproduction or lost sales.
Suppliers unable to respond to short-notice reorders	Fair Value Distribution & Building Ecosystem	Suppliers with financial stability (fast payment terms) and operational support (technical guidance, digital tools) can invest in the flexible capacity and rapid changeover capability that short-cycle replenishment demands.

The result is a portfolio-level, data-driven capital allocation, where the cost of being wrong is structurally minimised. Critically, however, a demand signal is only as valuable as a factory's ability to act on it. Receiving a reorder trigger means little if production lines and machines cannot be reorganised quickly enough to respond. In practice, this requires factories to be equipped with real-time production planning tools—such as Hyran Technologies—that can instantly replan machines in response to a new order signal, without the delays and manual coordination that characterise traditional scheduling.

SHEIN ON BUILDING AN ON-DEMAND INVENTORY SYSTEM ACROSS 7,500+ SUPPLIERS

The following case study is drawn from interviews conducted with different functional teams across SHEIN—supply chain operations, merchandising and planning, sustainability, supplier relationship management, fabric sourcing. Together, they provide a detailed view of how an agile inventory model operates in practice at scale.

The architecture: three tiers, one information system

The company operates a network of over 7,500 suppliers, producing across multiple product categories. At its core, the inventory model is built on the principle of producing only what is needed and scaling only when demand confirms it.

The architecture operates across three tiers, coordinated through a single digital platform:

- **Design and initial commitment.** New styles enter the market in initial batches of approximately 100 units. This quantity is deliberately small. It represents a test, not a forecast. For new colours of existing styles, the system automatically calculates the initial order quantity using sales performance from other styles which have the same colour, further reducing guesswork. The merchandising team establishes monthly development themes by combining historical sales data with trend forecasting, but the subsequent order quantity for the new style is deferred until the demand signal of the initial batch is captured.
- **Garment manufacturing (Tier 1).** Once an initial batch is placed with a supplier, all subsequent reorders for that style are automatically directed to the same factory, ensuring consistency. This is also an essential feature of the agile model. Since suppliers are leveraging sales signals to drive production decisions, if a single sales signal is distributed across multiple suppliers, each supplier may not know how the other suppliers are reacting to the changes in demand, resulting in disjointed production decisions. Approximately 50% of reorders are fully automated, triggered by real-time sales velocity and inventory data.
- **Fabric supply (Tier 2).** The company does not directly purchase fabric, but acts as a coordinator, while the actual transaction occurs between garment factories (Tier 1) and fabric mills (Tier 2). However, the company maintains full visibility of all such transactions on its platform in relation to which factory ordered from which mill, in what quantity, and the delivery status. Garment-level sales data is converted into fabric demand forecasts using the bill of materials and shared with fabric suppliers through a dedicated portal. Fabric suppliers reciprocally share their inventory levels with SHEIN and the Tier 1 suppliers, creating a multi-directional visibility loop. Delivery requirements for fabric from Tier 2 to Tier 1 suppliers are short: 24 hours within Guangzhou, 48 hours outside. Delivery performance directly determines supplier ranking and future business allocation, creating a self-reinforcing mechanism where good performers receive more orders.

» CASE STUDY

The inventory outcome

The system produces nearly ten times less unsold inventory than the reported industry average. This is an outcome of the small-batch, signal-driven architecture: when initial orders are around 100 units and reorders are triggered by confirmed demand, the structural opportunity for overproduction is dramatically reduced. When unsold inventory does occur, the company employs a collaborative approach to clearance. Suppliers have access to demand data, which allows them to anticipate orders and, in some cases, produce in advance. This is not mandated by SHEIN, but a practice that is often executed by suppliers to manage their own production capacity. When the advance production does not sell, the company assists in clearing the stock through its platform, with the financial impact of any markdowns shared between the company and its supplier.

Enablers

Three enablers emerged from our interviews as foundational to the system's inventory performance and were validated with external expert interviews:

- **Integrated data flows across tiers.** The company's order management system serves as the single source of truth for all product specifications. The same platform captures real-time sales velocity, production status, and inventory levels across the network. At the fabric level, the company's digitised bill of materials automatically calculates the exact yardage required per style and communicates this to fabric suppliers when orders are placed. This precision—from garment-level demand signal to fabric-level material requirement—is what allows the system to operate with minimal buffer stock and waste.
- **Financial architecture that enables supplier responsiveness.** The company operates payment cycles as fast as one week.

Critically, it also mandates and monitors prompt payment from Tier 1 garment factories to Tier 2 fabric suppliers, with fabric suppliers able to report non-payment through the system. This multi-tier payment discipline keeps inventory days low—often under 10 days for performing fabric suppliers—by ensuring that suppliers at every level have the working capital to hold minimal stock and respond quickly to new orders rather than building safety buffers against payment uncertainty.

- **Fabric postponement and consolidation.** Fabric suppliers naturally practise postponement: they prepare greige (undyed) fabric and commit to dyeing only when actual colour demand emerges. The company reinforces this through active fabric consolidation—preventing the proliferation of similar fabrics at the development stage and periodically evaluating existing fabrics for merger. This concentration of demand onto fewer fabric qualities simplifies upstream logistics, improves mill utilisation, and reduces the total volume of fabric inventory that needs to be held across the network.

Supplier capability investment

SHEIN has built a dedicated Supplier Relationship Management department to support suppliers. Their performance is measured on three KPIs: visit coverage rate, problem resolution rate, and supplier Net Promoter Score.

New suppliers require approximately six months of dedicated evaluation and onboarding before full integration. During this period, they are assessed on cost, capacity, delivery performance, and their ability to adapt to the small-batch workflow. Support includes a Supplier Learning Centre with online courses and live broadcasts, success story sharing from established suppliers, and on-site technical guidance from expert teams who help optimise factory layouts and equipment. An AI-powered chatbot ("Suppliers Ask") provides real-time digital support for operational questions.

Where to start

For brands beginning the transition, the following sequence provides a practical starting point:

- 1. Segment the product range by demand uncertainty.** Not all SKUs carry the same inventory risk. Core, replenishment-driven styles with stable demand can be managed through automated reorder triggers with relatively high confidence. New, trend-driven styles, where forecast error is highest and the cost of overproduction most acute, are where the agile model delivers its greatest value. Start by applying small-batch, test-and-respond disciplines to the latter category, where the return on reduced inventory risk is immediate and measurable. This is the same dual-track logic described in the operating-decisions section: differentiate supply chain strategy by demand uncertainty, not by brand-wide policy.
- 2. Build multi-directional data visibility with key suppliers.** The foundational step is creating a shared digital platform connecting the top 20% of suppliers (who likely represent 80% of volume). Prioritise data flows: frequent demand signal (real-time sales data or frequent demand forecast)—which our research found to be both a critical bottleneck for brands and a key enabler for suppliers, who use it to pre-position and postpone fabric, block and plan production capacity—and inventory levels across tiers i.e. greige fabric levels, dyed fabric levels, finished goods at Tier 1 and finished goods at the brand's warehouse. The demand signal can later be complemented by automated reorder triggers for replenishment of proven styles. A full collaborative system is not required on day one. More frequent demand forecasts, material pre-positioning with postponement, and inventory visibility across Tier 1 and Tier 2 can yield measurable improvements in lead time and sell-through rates within the first quarter.
- 3. Align payment terms with responsiveness expectations.** Asking suppliers to hold less inventory and respond faster while maintaining extended payment terms can be contradictory. Financial architecture and operational requirements must be aligned. This does not

necessarily mean paying suppliers in one week, but it does mean that payment terms should be structured to ensure suppliers do not have to bear both payment and order uncertainty.

- 4. Invest in supplier capability.** Factory-floor responsiveness is the rate-limiting step that determines whether a demand signal translates into replenished stock or a missed sales opportunity. Supporting suppliers to adapt to variable-order, fast-turnaround workflows is critical in unlocking agile models.

What to monitor:

As brands transition from forecast-driven to demand-driven inventory management, the following metrics track whether the model is delivering the working capital and inventory improvements it promises:

- > End-of-season sell-through rate:** the proportion of seasonal inventory sold at full price before the markdown window opens; the primary indicator of forecast-to-demand alignment.
- > Weeks of cover at point of reorder:** how much stock remains when a replenishment order is placed, which measures how tightly the system responds to real demand rather than safety stock assumptions.
- > Markdown rate as a percentage of gross margin:** captures the direct financial cost of overproduction and should decline as agile disciplines take hold (agile benchmarks in our research fall well below the industry average of 25–40% of inventory not sold at full price).
- > Reorder fulfilment lead time:** the elapsed time from a confirmed demand signal to goods available for sale, which directly measures factory-floor responsiveness and the effectiveness of production planning tools.
- > Initial batch-to-reorder ratio:** the share of styles that progress from a first small batch to at least one confirmed reorder, which measures how effectively the test-and-respond model is filtering demand signals.

From Inventory Reduction to Capital Redeployment

The traditional inventory model is a capital trap: working capital is locked in unsold goods, warehouse space is occupied by slow-moving stock, and management attention is absorbed by markdown decisions and clearance logistics.

The agile model does not eliminate inventory, but it structurally limits the conditions under which excess inventory accumulates. By replacing large upfront commitments with small test batches, seasonal forecasts with real-time demand signals, and siloed data with integrated multi-tier visibility, it converts inventory from a speculative bet into a responsive system. This is the working capital dimension of the total system margin shift described in the cross-cutting enablers section. It is not achieved through better forecasting. It is achieved through a supply chain architecture that makes the consequences of forecast error structurally small.

AGILE SUPPLY CHAIN TO REDUCE MARGIN EROSION

Most brands operate on a cycle of seasonal forecasting and bulk production that leads to inevitable end-of-season markdowns. This cycle is a primary driver of brand value and margin erosion. An analysis of 36,036 investor questions from the past decade reveals that gross margin compression and markdown dependency are the second and third most pressing concerns expressed by investors to senior leaders, collectively accounting for over 20% of all problem-related topics discussed.

The reliance on markdowns is not a minor cost of doing business; it is a structural flaw that can erode 10–15% of potential revenue⁵³, depress margins and signal a fundamental disconnect between what is produced and what customers actually want and are willing to pay.

Traditional supply chains, with their 6–12-month lead times, are designed to push products into the market based on forecasts. When those forecasts are wrong—as they increasingly are in a volatile market—the primary recourse is the markdown cycle. An agile supply chain can flip this model. By creating a tight feedback loop between real-time demand signals and flexible production, it shifts the focus from managing forecasting errors to avoiding them altogether. This doesn't just optimise markdowns; it can make them largely unnecessary.

A note on scope: clearing products that did not sell according to forecast is only one of the motivations behind markdowns. Others include deliberate pricing strategy to drive sales through perceived value and promotional urgency, which are independent of the supply chain model. This section focuses on markdowns driven by overproduction and forecast error—the component that an agile supply chain can structurally address.

The Agile Advantage: From Markdown Optimisation to Markdown Elimination

An agile supply chain protects margins by design, replacing the traditional “make-then-sell” model with a “sense-and-respond” capability. This structurally reduces the risk of overproduction and the subsequent need for margin-eroding discounts.

Traditional Model Challenge	Agile Enabler	Business Performance Outcome
Forecast inaccuracy	Demand-Driven, Agile & Responsive Operations	Production is initiated by real-time sales data or direct customer pre-orders, not long-range forecasts, minimising the risk of producing unwanted goods.
Long production lead times	Building Ecosystems	A digitally connected network of suppliers enables rapid, small-batch production runs, allowing brands to react to trends within the season rather than committing months ahead.
High minimum order quantities (MOQs)	Fair Value Distribution	By offering suppliers better terms and operational support, brands can negotiate lower MOQs, enabling a test-and-learn approach that reduces initial inventory risk.
Lack of in-season visibility	Digital Infrastructure, Intelligence & Governance	Real-time data on sales and inventory levels allows for precise, data-driven replenishment of winning styles, maximising full-price sell-through.

⁵³ Coresight Research and Celect, as reported in Retail Dive, estimate that non-grocery retailers missed out on roughly \$300 billion in revenues in 2018—about 12% of total sales—due to markdowns.

MARGIN PROTECTION IN PRACTICE

A growing number of brands across segments are using agile models to protect margins, enhance pricing power, and strengthen brand equity. The following examples illustrate how these principles are being applied at different price points and scales.

RISE & FALL: AGILE PRODUCTION WITH LUXURY-GRADE MATERIALS

UK-based Rise & Fall demonstrates how the core mechanics of the agile, test-and-respond model can be applied to the accessible luxury segment—using luxury-grade materials and factory partnerships while maintaining the demand-driven discipline that protects margins.

The brand operates a just-in-time model, using a prediction engine to test demand for small quantities of new products before committing to larger orders. Initial production runs are modest; successful styles are replenished rapidly through factory partnerships that can turn around repeat orders within 48 hours for some categories. When demand outstrips supply, the brand uses pre-orders to capture sales with zero inventory risk.

A knitted hoodie illustrates the model's power. The brand launched it with an initial order of approximately 100 units across a few colours. As real-time demand signals proved positive, rapid repeat orders followed. Within weeks, thousands of units had sold, with at least half fulfilled via pre-order. The item grew to represent 40% of the brand's orders in the final weeks before Christmas, precisely because the agile supply chain allowed production to chase demand rather than forecast it.

Three elements of Rise & Fall's approach are particularly instructive for the principles outlined in this paper:

- **Fabric palette rationalisation:** the brand works with a deliberately streamlined fibre basket—cashmere, merino, cotton, silk, leather, linen, and organic denim—and specifies technical

requirements (micron count, fibre length) within each. This concentration of demand onto a curated material palette reduces supply chain complexity, enabling faster response from mills.

- **Greige fabric postponement:** Rise & Fall holds undyed fabric in stock, committing to colour only after demand signals emerge. This is a practical solution to the Tier 2 working capital barrier identified in our limitations analysis: by holding inventory in a flexible, uncommitted form, the brand reduces both obsolescence risk and the capital cost of pre-positioning, while maintaining the speed to respond when demand confirms.
- **Geographic proximity and supplier relationships:** the brand sources directly from quality-grade factories, investing in long-term relationships that enable the small MOQs and rapid turnaround the model requires. As discussed in the cross-cutting enablers section, building this supplier capability takes sustained effort but it is what makes the agile model commercially viable.

Rise & Fall's commercial results provide evidence that protecting margin and building brand equity are not in tension. Since 2021, the company has more than doubled revenue year-on-year, sustaining a compound annual growth rate exceeding 100% over a three-year period. In an interview with Vogue, the founders explained that the brand found that when it raised prices on certain products to better match customer quality expectations, conversion rates increased—a counter-intuitive result that illustrates the pricing power preserved by avoiding markdowns⁵⁴.

54 Vogue (US), "How Rise & Fall Is Making Luxury Fashion Affordable," 12 March 2026

» CASE STUDY

ASPHALTE: PRE-ORDER AS A STRUCTURAL MARKDOWN ELIMINATOR

French B Corp-certified brand Asphalte has scaled a model that takes the agile principle one step further: producing only what has already been sold⁵⁵. The company co-creates its collection with its community, then launches products for pre-order, manufacturing only against confirmed demand.

This zero-inventory approach has three consequences for margin protection:

- No overproduction: the model is designed to avoid overproduction and unsold stock, a primary driver of forced markdowns.
- Improved cash flow: under their pre-order system, customers pay for garments up to three months in advance, financing the production run and significantly reducing the need for working capital tied up in inventory.
- Strong commercial validation: According to a case study published by Shopify, the model has delivered €2.4 million in revenue from its first physical store in the seven months following

its September 2024 Paris opening, with a Net Promoter Score of 75—considered excellent for retail⁵⁶—suggesting that customers accept the trade-off of longer delivery times when the value proposition is clear.

Asphalte demonstrates that a pre-order model can be a powerful engine for profitable, sustainable growth in the mid-market, and that avoiding the traditional markdown liability generates structural competitive advantage.

MARIA GRACHVOGEL: AGILITY IN LUXURY

The principles extend to luxury. London-based designer Maria Grachvogel operates a tiered system that combines replenishment of core bestsellers based on sales velocity, limited-run testing for new styles, and made-to-order for archive and occasion wear. By producing locally and maintaining a strict no-discounting policy, she eliminates the markdown cycle entirely while preserving brand exclusivity. Her model demonstrates that for luxury brands, agility does not mean speed for its own sake. It means the strategic precision required to match supply exactly to demand, protecting both margin and brand integrity.

Where to Start

The three case studies above span different price points, scales, and customer relationships, but they follow a common sequence that any brand can adapt:

1. Start with a core, replenishment-driven category. Grachvogel's "iconic pieces" tier and Rise & Fall's streamlined fibre basket both concentrate the agile model on products with strong, repeatable demand. Piloting test-and-respond on a core category—where you have established supplier relationships and reliable sales data—minimises risk while building the operational muscle for more volatile categories. This is the same sequencing logic recommended in the circularity section: start where you have the strongest foundation.
2. Introduce a test-and-learn mechanism for new styles. Whether through pre-order (Asphalte), limited-run testing (Grachvogel), or small-batch initial orders with rapid replenishment (Rise & Fall, SHEIN), the principle is the same: decouple the decision to develop a product from the decision to produce it at scale. Commit to design; defer the volume commitment until demand confirms. In practice, this means initial production runs as low as 80–100 units for unproven styles, with reorder quantities calculated dynamically based on actual sales velocity.
3. Invest in supplier relationships that enable flexibility. None of these models work without suppliers willing and able to accept small, variable orders with rapid turnaround. This requires the sustained investment in supplier capability building described in the cross-cutting enablers section through dedicated relationship management, technical support,

55 Candiani Denim, "The French brand Asphalte takes a new production approach," March 2023

56 Shopify, "How Asphalte developed its retail strategy," shopify.com/no-en/case-studies/asphalte (verified June 2026).

and financial terms that make the agile model viable for the supplier, not just the brand. Rise & Fall's founders are explicit that they are "selling suppliers the dream of where the business will be in five years, not the first purchase order."

4. Rationalise the material palette. A curated fabric library, whether Rise & Fall's streamlined fibre basket or the centralised fabric library of SHEIN, concentrates demand onto fewer fabric qualities. This simplifies upstream logistics, improves mill utilisation, and critically, enables the greige fabric postponement strategy that reduces the Tier 2 working capital barrier for higher-value materials.

What to Monitor

As brands transition from forecast-driven to demand-driven production, the following metrics track whether the model is delivering the margin protection it promises:

- > **Full-price sell-through rate:** the proportion of units sold without markdown. This is the single most important indicator of demand-supply alignment and the primary metric that distinguishes agile from traditional models. Industry averages for full-price sell-through typically range from 60–70%; the agile models described above target 90%+ by structurally reducing the volume of goods that require discounting.
- > **Pre-order and made-to-order share:** the percentage of total revenue generated through pre-order or made-to-order channels, where inventory risk is extremely low. Asphalte's model sits near 100%; Rise & Fall uses pre-order selectively when demand exceeds available stock. Tracking this metric over time shows how much of the business has been structurally de-risked from the markdown cycle.
- > **Replenishment velocity:** how quickly a proven style can be restocked after selling through. This is the operational metric that enables everything else: if replenishment is fast enough, brands can afford to start with small initial quantities, confident that they can scale up within the season. Rise & Fall's or SHEIN's 48-hour turnaround for some categories represents the frontier we observed in our research; most brands will start with a longer cycle and work to compress it.

From Margin Protection to Margin Creation

The traditional markdown cycle is not just a cost. It is the signal of a supply chain operating with a fundamental mismatch between what is produced and what is wanted. The brands profiled above have demonstrated that this mismatch is not inevitable. By replacing forecasts with demand signals, long lead times with rapid replenishment, and large upfront commitments with test-and-learn approaches, they have structurally removed the conditions that make markdowns necessary.

The result goes beyond margin protection: capital that was previously tied up in inventory that would eventually be discounted is instead deployed against confirmed demand at full price. This is the total system margin shift described in the cross-cutting enablers section, applied to one of the industry's most persistent and costly problems.

AGILE SUPPLY CHAIN TO DRIVE CIRCULAR TRANSITION

Shifting from Regulatory Burden to Business Opportunity

Several pieces of circular economy legislation are reshaping the competitive landscape for the fashion industry. In the EU, the ESPR will, from July 2026, ban the destruction of unsold apparel and footwear for large companies, while the Digital Product Passport (DPP) will soon mandate unprecedented product-level transparency. In the United States, California's Responsible Textile Recovery Act (SB 707) establishes the nation's first statewide textile EPR framework. Covered producers must join CalRecycle's approved Producer Responsibility Organisation (PRO) by 1 July 2026. The scheme's wider operational and financial end-of-life-management requirements will take effect progressively, with full implementation expected by 2030. These regulations can represent a significant new layer of complexity and cost for companies operating with traditional linear supply chains. However, for businesses that have invested in agile, demand-driven models they represent a source of competitive edge. The core capabilities required to thrive in an agile system—data integration, multi-tier traceability, and a strong connection between design and manufacturing—are the same capabilities required to meet new circularity standards efficiently and effectively. This alignment can turn a compliance mandate into a business performance opportunity.

How Agility Enables Circular Operations

Agile supply chains provide the foundational infrastructure to comply with new circularity regulations and create business value. The principles of agile supply chains directly map to the emerging regulatory requirements and circular economy principles.

Regulatory Requirement	Principle	Business Performance Outcome
ESPR: Digital Product Passport (DPP)	Digital Infrastructure, Intelligence & Governance Transparency, Traceability & Compliance	Companies with centralised order-management systems and multi-tier supplier data could already hold information likely to be relevant to future Digital Product Passport requirements. For these operators, generating DPPs is a low marginal cost extension of existing data flows—not a new compliance build. This also creates a customer engagement tool and a supplier accountability mechanism.
EPR & Eco-Modulation Minimum Design Criteria	Transparency, Traceability & Compliance Circularity & Resource Stewardship	The ability to track material composition from Tier 3 and to source safe, renewable, and recyclable materials enables accurate EPR reporting and allows for design adjustments to meet minimum design requirements and qualify for eco-modulated fee reductions.
Ban on Destruction of Unsold Goods	Demand-Driven, Agile & Responsive Operations	The most efficient way to comply with a ban on destroying unsold goods is to not create them in the first place. An agile model that starts with small initial batches (as low as 80–100 units for unproven styles) and scales production only in response to confirmed demand structurally minimises the volume of unsold goods that would fall under this regulation.
Circular Design Implementation	Building Ecosystems	An orchestrated network with strong feedback loops allows new circular design standards (e.g., mono-materiality, rivet removal) to be rapidly implemented by tapping into supplier manufacturing knowledge. Where siloed supply chains force top-down mandates, agile networks enable co-development.

THE ELLEN MACARTHUR FOUNDATION'S JEANS REDESIGN

The Ellen MacArthur Foundation's (EMF) Jeans Redesign project (2019-2023) provides a concrete illustration of how to implement circular design principles at scale. By focusing on an iconic, high-volume garment, the project demonstrated that designing for circularity is a practical, achievable goal. The project brought together over 100 organisations—including major brands like H&M, Gap, and Zara, as well as manufacturers and fabric mills—around a common set of guidelines. These guidelines established a minimum bar for circular design, focusing on tangible, measurable criteria.

Guideline	Key Requirement	Circular Rationale
Durability	Garments must withstand a minimum of 30 home laundries.	Extends the useful life of the product, reducing the need for new production.
Material Purity	Fabric must be composed of at least 98% cellulose-based fibres.	Enables the garment to be effectively recycled back into new fibres without contamination.
Recyclability by Design	Metal rivets must be eliminated or designed for easy removal.	Simplifies the disassembly and recycling process, increasing the yield of recovered material.
Safe Materials	Adherence to a ZDHC Manufacturing Restricted Substances List (MRSL).	Ensures that materials circulating within the system are non-toxic and safe.

The Results:

By 2023, participants had brought over 1.5 million pairs of redesigned jeans to market. What made this project possible was the involvement of the entire supply chain from the outset in the design of requirements. By bringing fabric mills and manufacturers into the room alongside design teams, the project was able to pool expertise and resolve what some stakeholders initially perceived as technical barriers. For many of the criteria, the project's conclusion was striking: success was, in the words of the EMF, "no longer a question of technical capability, but a design choice."

Where to start

Leaders looking to operationalise circularity within their own supply chains can adapt the deliberate sequence that EMF followed for the Jeans Redesign:

1. **Start with one high-volume category, not the entire range.** EMF chose denim precisely because it is a high-volume garment with a well-understood supply chain. This made the guidelines tangible and testable.

For a business leader, the equivalent is to select a single product category where you have strong supplier relationships and good data visibility—likely a core, replenishment-driven line rather than a trend driven item.

2. **Define a small number of measurable criteria before engaging suppliers.** EMF did not ask participants to "be more circular." They defined four specific, binary criteria (material purity $\geq 98\%$, 30 washes, no non-removable rivets, ZDHC compliance). This clarity is essential because it allows suppliers to assess feasibility and cost impact before committing, and to propose practical solutions for how to achieve those ambitions, including where to be more ambitious. In an agile system with centralised product specifications, these criteria can be embedded directly into tech packs and informed directly by suppliers' manufacturing expertise, enhancing the likelihood of commercial viability.

3. **Involve Tier 2 partners from the outset.** Many of the Jeans Redesign's criteria—particularly material purity and chemical safety—are determined not at the garment factory (Tier 1) but at the fabric mill and dyehouse (Tier 2). EMF's inclusion of mills alongside brands and manufacturers was a critical success factor. For companies building agile supply chains, this reinforces the importance of extending digital connectivity and relationship management beyond Tier 1. A curated fabric library, for instance, can serve as a powerful lever: by restricting the available material palette to circular-compatible options, it embeds compliance from the material selection stage.

What to monitor

As companies move from pilot to scale, tracking the right metrics is essential to demonstrate progress, justify continued investment, and identify bottlenecks early.

1. **SKU compliance rate:** the percentage of active SKUs that meet defined circular design criteria (e.g., material purity thresholds, durability standards). This is the primary output metric and should be tracked at the category level.
2. **Total margin impact of circular specifications:** rather than isolating the per-unit cost difference between circular and conventional specifications, track the full margin picture: any change in production cost alongside the reduction in markdowns, unsold inventory write-offs, and—in EPR jurisdictions—eco-modulated fee savings that circular designs deliver. The Jeans Redesign demonstrated that many criteria (e.g., rivet removal) are cost-neutral or cost-reducing at the unit level; when avoided overproduction and end-of-life liabilities are factored in, the total system margin case is substantially stronger.
3. **Regulatory readiness score:** the percentage of product volume for which the company can already generate the data required by incoming regulations (DPP fields populated, EPR material declarations complete, minimised inventory). This metric turns compliance into a visible, trackable gradient and makes visible the gap between current capability and regulatory deadline.

From project to operating system

The Jeans Redesign demonstrated what is possible when circular criteria are clearly defined, and the full supply chain is engaged to create products that are not only compliant with future regulations but also commercially successful. But it was bounded in time, scope, and category. The challenge for the industry is to move from project-based circularity to circularity as a default mode of operation.

This is where agile supply chains offer a structural advantage. The same capabilities that power demand-driven production—centralised product specifications, digital connectivity across tiers, curated material libraries, and collaborative supplier relationships—are precisely the capabilities needed to embed circular design criteria at scale, collaboratively deploy them through automated sampling and quality gates, and track compliance in real time. Agile supply chains do not merely make circularity compliance easier; they provide the operating system to make it a source of competitive differentiation.



APPENDIX

APPENDIX A—INVESTOR QUESTION ANALYSIS: METHODOLOGY & TOPIC REFERENCE

The Scene Setting section draws on two complementary datasets: a financial benchmark tracking the operating performance of 47 listed fashion companies, and a structured analysis of 36,036 analyst questions from their earnings calls. This appendix documents how each was constructed, which companies were covered, and the full taxonomy of topics surfaced—so the figures in the main text can be read with their methodology in view.

A.1—The two datasets

Two separate analyses underpin the Scene Setting narrative. They should not be conflated: one measures financial outcomes, the other measures investor attention.

DATASET 1—FINANCIAL BENCHMARK

A panel of listed fashion and apparel companies tracked on revenue growth, operating profit growth, and operating margin movement between 2014 and 2025. This is the source of the segment-level margin table in Scene Setting (Premium, Sportswear, Mass, Luxury). The data comes from audited company disclosures.

DATASET 2—INVESTOR QUESTION CORPUS

A structured analysis of 36,036 analyst questions extracted from 1,000 fashion and luxury earnings call transcripts across 39 companies (8 of the 47 did not have available transcripts), spanning 2014 to 2025. This is the source of all signal-frequency figures.

A.2—Financial benchmark: segment definitions

The four segments in the Scene Setting margin table group companies by business model and price positioning. Because these segments operate through materially different channels—DTC, omnichannel, and wholesale-dependent—a brand's segment indicates its structural exposure. Representative constituents are listed below to illustrate the segments to readers.

Segment	Defining characteristics	Representative constituents
Luxury	High price, using mostly natural materials and Tier 1 manufacturing mostly located in Western Europe. Vertically controlled distribution.	<i>LVMH, Kering, Hermès, Richemont, Prada, Moncler, Burberry, Zegna</i>
Premium	Mid-tier pricing, mix of natural and synthetic materials, Tier 1 manufacturing spread between Western Europe, Eastern Europe and Southeast Asia. Wholesale-dependent.	<i>PVH, Capri, Tapestry, Ralph Lauren, Hugo Boss, Levi's, Gap</i>
Sportswear	Pricing similar to premium, activewear and shoe driven. Mostly synthetic materials, manufacturing concentrated in China and Southeast Asia. Wholesale and direct distribution.	<i>Nike, adidas, Puma, Lululemon, Deckers, On, Amer Sports, VF Corp</i>
Mass	High-volume, value-led, mostly vertically controlled distribution. Heavy reliance on synthetic materials, Tier 1 manufacturing mostly in Southeast Asia.	<i>Inditex, H&M, ASOS, American Eagle, Abercrombie</i>

A.3—Investor question analysis: methodology

Data collection

- Source: earnings call transcripts (2014–2025) sourced from financial data providers.
- Coverage: 1,000 earnings call transcripts across 39 fashion and luxury companies (47-company target universe; 8 without available transcripts).
- Total analyst questions extracted: 36,036, after excluding management speech via speaker-role detection. Coverage period: 2014 to 2025.

Signal detection

- Parsing: speaker-role detection isolates the Q&A portion of each call and separates analyst questions from management responses using speaker-affiliation tags, so only analyst-initiated questions are counted. Calendar years are extracted directly from filenames to ensure US and European reporting formats are bucketed consistently.
- Classification Layer 1—broad topics: a 10-category keyword taxonomy provides backward-compatible bucketing of every question.
- Classification Layer 2—signals: a purpose-built 25-signal taxonomy (15 operational problems and 10 strategic aspirations), with keyword patterns mapped to the white paper’s seven principles.
- Signal intensity is measured as mentions per 1,000 analyst Q&A blocks. A single question can trigger multiple signals where problems overlap, so frequencies count attention rather than mutually exclusive buckets; percentages reflect the share of total questions, and small-base signals are rounded.
- Principle mapping is auto-derived from signal tags rather than coded separately, ensuring consistency between signal counts and principle-level totals.

Coverage by sector

The transcript corpus spans six fashion and luxury sectors. Sector-level transcript volumes vary with reporting cadence: European-listed houses report semi-annually and generate fewer transcripts than US-listed quarterly reporters.

Sector	Reporting cadence	Coverage period
Luxury	Mostly semi-annual	2014–2025
Mass	Quarterly	2014–2025
Premium	Quarterly	2014–2025
Sportswear	Quarterly	2014–2025
Fashion retailer	Quarterly	2014–2025

A.4—Topics surfaced: full problem-signal taxonomy

The fifteen problem signals below are the complete set of operational concerns the corpus tracks, ranked by total frequency across all years. The “% of questions” column is each signal’s share of all 36,036 analyst questions; “per 1,000” expresses the same frequency as mentions per 1,000 Q&A blocks. Together the three operational symptoms at the top—inventory overhang, gross margin compression, and markdown dependency—account for 47% of all problem-signal mentions.

Alongside the fifteen problems, the corpus tracks ten aspiration signals—the outcomes investors push management toward. The most frequent, by a wide margin, is investment in data and digital infrastructure, followed by speed-to-market and channel control. These map directly onto the white paper’s core principles and argument.

Problem signal	Mentions	% of Qs	Per 1,000
Inventory overhang & working capital trap	4,766	13.2%	132.3
Gross margin compression	3,744	10.4%	103.9
Markdown dependency & full-price erosion	3,654	10.1%	101.4
Digital channel fragmentation & platform dependency	3,568	9.9%	99.0
Wholesale channel loss of control	3,398	9.4%	94.3
Currency & FX margin exposure	1,851	5.1%	51.4
Supply chain disruption & lead-time volatility	1,089	3.0%	30.2
Geographic concentration & geopolitical overexposure	1,001	2.8%	27.8
Pricing power erosion & consumer pushback	672	1.9%	18.7
Demand unpredictability & forecast failure	647	1.8%	18.0
Return rates & reverse logistics cost	619	1.7%	17.2
Brand dilution from overexposure or inconsistency	425	1.2%	11.8
Regulatory & compliance burden	144	0.4%	4.0
Carbon & climate physical risk to operations	95	0.3%	2.6
Supplier financial fragility & payment terms	58	0.2%	1.6

A.5—How investor attention shifted, 2020–2025

The pandemic and its aftermath produced a step-change in operational scrutiny, with most signals peaking in the 2020–23 inventory crisis before easing. The exception is geographic and geopolitical exposure, which stays subdued until a sharp 2025 spike driven by tariff escalation.

The table below traces signal intensity (mentions per 1,000 Q&A blocks) through the period. It is the data behind the compounding forces described in Scene Setting, and is offered as the raw material for a timeline visual should the team wish to give that analysis its own figure.

Problem signal	2020	2021	2022	2023	2024	2025
Inventory overhang	161	149	192	170	123	110
Gross margin compression	88	110	112	118	128	115
Markdown dependency	101	94	125	124	100	92
Wholesale loss of control	110	108	102	109	106	98
Digital fragmentation	138	125	103	86	83	73
Supply chain disruption	29	58	61	39	27	23
Demand unpredictability	19	21	31	27	21	22
Geographic / geopolitical	30	10	7	13	29	129
Carbon & climate physical risk	4.6	4.6	5.2	3.4	3.8	1.4

APPENDIX B—METHODOLOGY: OPTIMISING INVENTORY AND CARBON

This model is an orientation tool, not a precise benchmark. It maps the direction and order of magnitude of the trade-off between freight intensity and overproduction. Its inputs are archetype emission-intensity values informed by research across several listed players, not calibrated company figures. It is presented as a navigator—a surface a brand can locate itself on and read which way to move—rather than as a precise calculation or a comparison of named operators. Intensities are expressed per € million of revenue (tCO₂e/€m); figures are illustrative and rounded.

A. Objectives of the model

The model isolates the carbon that differs between operating models for the same hypothetical brand: overproduction, defined as the manufacturing carbon embodied in goods made but not sold, and transport. Manufacturing of the goods actually sold is identical across models for a given revenue base and is therefore excluded from the comparison.

Rather than compare two fixed points, the model maps this carbon across a plausible range of air-freight intensity and overproduction outcomes—a surface, not two points. Every real operating model becomes a region a brand can locate itself within, and the trade-off can be read without relying on, or revealing, any single company's figures.

B. Scope and boundaries

- **Variable:** transport emissions and the manufacturing emissions embodied in unsold goods.
Held constant: manufacturing of goods sold, which is identical across models for a given revenue base and therefore excluded from the comparison.
Out of scope: returns, the use phase and end-of-life.
- **Basis:** FY2024. Intensities are expressed per € million of revenue, the standard basis where unit volumes are not disclosed. The workbook holds revenue constant at €1bn for scenario calculation, but the navigator reports results per €m to reduce the influence of scale.

C. Key figures

Inputs are archetype values informed by research across several listed players spanning the responsive-to-planned spectrum. They are not intended to represent any single operator.

Input	Value	Source
Manufacturing EI (baseline)	~240 tCO ₂ e/€m	Mass-market manufacturing intensity; average across several listed players. See §G.
Transport EI—y-axis	15—225 tCO ₂ e/€m	Planned / sea-shipped archetypes sit around ~45–60; air-served cross-border DTC archetypes sit around ~200–225. Values are derived from several players' 2024 transport Scope 3 disclosures divided by revenue.
Unsold rate—x-axis	2—35%	Defined as production above what is sold. Fully responsive models sit around ~2%; mixed models around ~10%; planned / traditional models around 25–30%, against a wider industry range of 25–40%.

EI = emission intensity. Manufacturing = Scope 3 “purchased goods & services”; transport = Scope 3 upstream & downstream transportation. Companies anonymised; emissions from public 2024 disclosures.

D. Key calculations

μ = overproduction rate, defined as units produced above what is sold. Revenue captures sold units only.

- > **Overproduction EI = Manufacturing EI $\times \mu$**
- > **Model-driven EI = Transport EI + Overproduction EI**
- > Net trade-off = (Transport EI – baseline transport EI) – Manufacturing EI $\times$ (baseline μ – μ)
- > A positive result is a net increase in emissions intensity. A negative result is a net reduction.
- > **Important basis note:** μ is expressed as production above sold units. If a source reports unsold goods as a share of total production, it should be converted before being used in this formula.
- > The baseline against which the net is measured is set at transport 30 tCO₂e/€m, the median transport emission intensity of the operators in our sample between 2023 and 2035, and overproduction 30%.

E. Reading the navigator—model output

Because the inputs carry real uncertainty, the navigator does not draw a single break-even line. It treats differences within ± 40 tCO₂e/€m as a zone of comparability: configurations whose carbon impact is not meaningfully distinguishable given the precision of the calculation. Below -40 is a clear benefit, or net reduction. Above $+40$ is a higher impact, or net increase. Colour deepens with magnitude beyond the band.

Two readings follow. First, transport is the larger lever. Across the surface, transport intensity spans roughly 210 tCO₂e/€m, while the overproduction effect spans about 80 tCO₂e/€m. Total carbon therefore moves around two-and-a-half times more with the freight choice than with the inventory choice. Second, the zones matter. A fully air-served model sits firmly in higher impact. A clear carbon benefit is unlocked by low overproduction, with the transport ceiling rising as overproduction falls: it holds up to a transport intensity of ~ 38 tCO₂e/€m at 10% overproduction, ~ 50 at 5%, and ~ 57 at 2%. Above ~ 57 , no overproduction level nets a clear benefit.

F. Assumptions and limitations

Orientation tool, not benchmark. Archetype values are informed by several listed players, not any single operator. Results should be read for direction and order of magnitude.

Manufacturing baseline. The manufacturing baseline of ~ 240 tCO₂e/€m is the revenue-weighted average purchased-goods intensity (Scope 3, Category 1) across the listed sample over 2023–2025, from company disclosures (sample range ~ 150 – 460).

Linear overproduction. Overproduction carbon scales as Manufacturing EI $\times \mu$, consistent with the “produces $\mu\%$ more than sold” basis used in the model. Unsold units consume manufacturing emissions but earn no revenue.

Per-revenue intensities. Per-€ intensities embed a price-point and volume effect. Holding revenue constant at €1bn in the workbook limits that interference, while reporting the output per €m keeps the result comparable.

Model-level overproduction rates. Overproduction rates are model-level estimates: $\sim 2\%$ for fully responsive models, based on first-party research; $\sim 10\%$ for mixed models; and 25–40% for planned models, based on external benchmarks. The output shows the shape and order of magnitude of the trade-off, not a calibrated company-specific figure.



 **hyran**



↓ SWANSTANT ↑