GREEN FUTURE TECH
LIMITED / HONG KONG SAR

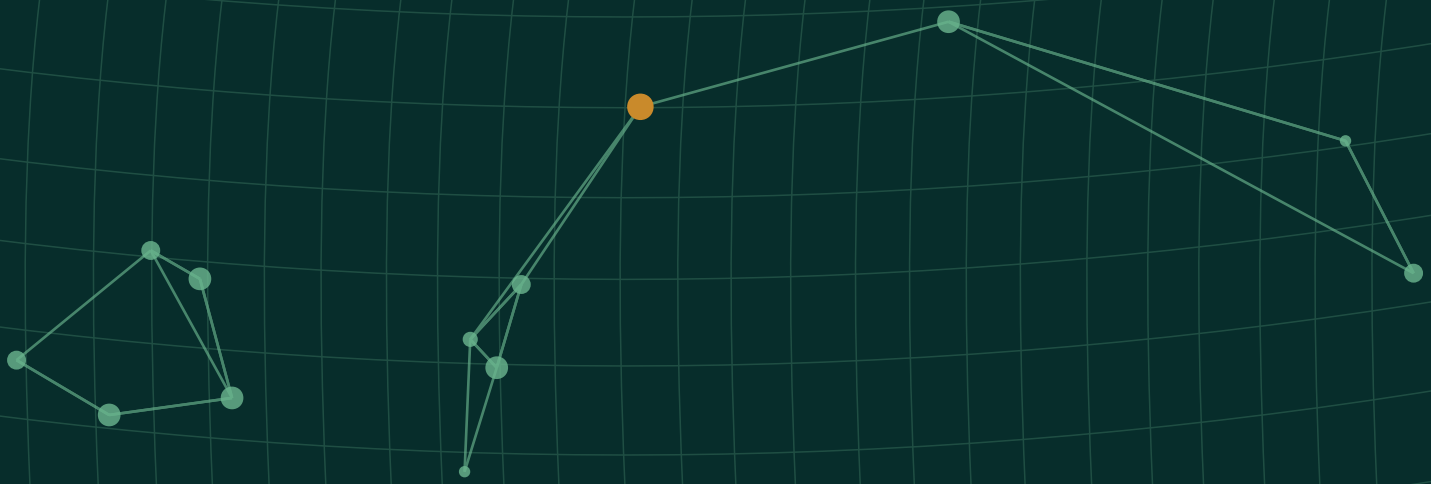
GFT RESEARCH / WARP AND SIGNAL

Inventory as a Decision

Inventory and supply-chain optimisation for emerging-market apparel makers

wp12
RESEARCH EDITION

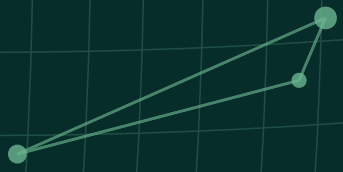
EDITION 2026-10-30
HONG KONG SAR



THE GFT THESIS

Value is made where it is measured.

GFT engineers sustainable factories that capture their own data from the first shift, so raw-material economies can keep the value they grow.







Read the evidence before acting on the recommendation.

ISSUER	Green Future Tech Limited, Hong Kong SAR
DOCUMENT	wp12 / Inventory as a Decision
EDITION DATE	2026-10-30
AUDIENCE	Apparel factory owners and general managers, Production and supply-chain planners, Procurement and warehouse leads, Finance and working-capital managers, Industrial-development partners evaluating operating readiness
STATUS	Research edition. Generated locally. Publication and operating-performance acceptance are separate decisions.
PROGRAMME	GFT's Cameroon textile programme: feasibility, performance-management system design and digital strategy. Prepared 31 January 2024; source-refresh annotations retained in 2026.

IMPORTANT NOTICE

This document is not investment, legal or tax advice. Proposed controls and illustrative calculations require project-specific validation. Targets and announced third-party projects are not evidence of delivered operating performance.

The confidential programme PDFs are not reproduced or published. No client identity or engagement code is included.

EVIDENCE REMAINS GRADED

A: official. B: institutional or major research. C: trade or local press and market-research vendors. D: unverified; excluded from use. A grade describes the supplied register and does not certify accuracy.

Inherited publication dates unrecorded: R01, R02, R03, R06, R07, R08, R09, R10, R11, R12, R14, R16, R27, R29, R32. Temporal eligibility is unverified; these are inherited source-text checks, not independent source validation.

The operating case runs from fibre to feedback.

01 / Inventory earns its place only when it protects an explicit promise 14

1.1 / A processing ambition does not determine a purchase order 15

1.2 / A missing trim can immobilise an otherwise complete garment 17

1.3 / The inventory objective must include cash and delivery 19

1.4 / The base case must fit a garment maker rather than an integrated mill 21

1.5 / Evidence determines how strongly a recommendation can be stated 23

1.6 / The paper begins where a forecast becomes an action 24

1.7 / The decision record makes the thesis testable 25

02 / Only identifiable, usable and reconciled stock can support a promise 28

2.1 / An item code must describe the material that production can substitute 29

2.2 / Unit conversions must preserve physical quantities 32

2.3 / Stock states prevent quality holds becoming delivery promises 34

2.4 / Receiving closes a chain of evidence 37

2.5 / Cycle counts diagnose recurring transaction errors 39

2.6 / RFID is a candidate capture method that requires a local trial 42

2.7 / Capture can continue during a network outage without corrupting stock 44

03 / Stock policy must follow commercial purpose and supply criticality 48

3.1 / Value classification is useful but insufficient 49

3.4 / Demand variability and intermittency require separate descriptions 51

3.7 / The customer promise determines the stock object 53

3.10 / Service measures must name the event they count 56

3.12 / Seasonal sell-through needs a different risk appetite 58

3.15 / Cash and physical risk can override an ABC label 59

3.17 / A policy matrix must prescribe an action 60

04 / Measure the full replenishment clock and its delays before choosing a buffer 63

4.1 / The usable-receipt boundary includes approval and inspection 64

4.5 / Historical corridor evidence is a diagnostic prompt 67

4.8 / Incomplete shipments must remain visible in the sample 69

4.11 / Tail measures and sample size constrain confidence 72

4.14 / Demand and lead time may worsen together 74

4.16 / Incoterms allocate responsibilities but do not erase elapsed time 75

4.17 / Reducing uncertainty can be preferable to buying more buffer 77

05 / Safety stock works only when its assumptions match the replenishment process 79

5.1 / Inventory position and available stock answer different questions 80

5.2 / The reorder point protects demand over lead time 81

5.3 / Random lead time adds a different variance term 83

5.4 / Periodic review protects the review gap as well 85

5.5 / An empirical quantile can be safer than an unsupported bell curve 87

5.6 / Fill rate requires expected shortages and order quantity 89

5.7 / The calculator must reject misleading inputs 90

06 / A cheaper purchase price can produce a more expensive inventory policy 93

6.1 / EOQ is a benchmark with restrictive conditions 94

6.2 / Holding cost includes opportunity cost without duplicating charges 96

6.3 / A discount must beat its extra stock and cash cost 97

6.4 / Time-phased demand needs time-phased lots 99

6.5 / Capacity and pack sizes make feasible quantities discrete 101

6.6 / Fabric choices create hidden obsolete stock 102

6.7 / Negotiating a smaller lot can change more than the price 104

07 / A seasonal buy needs a recovery plan before the first unit is ordered 107

7.1 / The single-period model prices both mistakes 108

7.2 / Demand evidence must include the days when stock was unavailable 110

7.3 / A second buy is valuable only if usable stock arrives in time 112

7.4 / Sizes and colours create non-interchangeable residuals 114

7.5 / Postponed manufacturing and postponed planning are different 116

7.6 / Markdowns and returns change the residual-stock calculation 117

7.7 / A seasonal stop rule protects cash from wishful demand 119

08 / A material plan is feasible only when every component and process can meet the date 121

8.1 / The bill of materials converts orders into dependent demand 122

8.2 / Net requirements must recognise what can actually be consumed 124

8.3 / Back-scheduling reveals late material before the line is loaded 126

8.4 / Minimum lots and capacity can make an MRP plan infeasible 128

8.5 / Work in progress is time as well as money 129

8.6 / A small buffer may protect a bottleneck while a large queue hides failure 131

8.7 / A complete-kit gate can improve the delivery promise 132

09 / Pool material and demand only where specifications and risks permit it 135

9.1 / A stock location is a service decision 136

9.2 / Demand correlation determines pooling potential 138

9.3 / Common fabric can pool risk before differentiation 140

9.4 / Multiple echelons can count the same uncertainty more than once 143

9.5 / Transport consolidation trades freight cost against stock and time 144

9.6 / Two suppliers may still share one disruption 146

9.7 / The network decision must state who carries residual stock 147

10 / Supplier reliability and fair contracts change how much stock is needed 149

10.1 / An on-time dispatch can still be an unusable delivery 150

10.3 / Supplier scoring needs visible trade-offs 151

10.5 / Order amplification is a system outcome 154

10.7 / Shared demand information needs an agreed meaning 156

10.9 / Flexibility contracts allocate rather than remove risk 158

10.12 / A late change must not be financed by harmful working practices 160

10.14 / Alternate sourcing needs qualification before an emergency 162

11 / The feasible stock plan is bounded by cash, not just by statistical service 165

11.1 / Map cash dates across the material-to-payment cycle 166

11.2 / Stock value, cash released and cost saved are different outcomes 168

11.3 / A cash-constrained allocation must reveal unmet demand 170

11.4 / Currency exposure belongs in scenario costs 172

11.5 / Receivables finance does not automatically fund the first fabric buy 173
11.6 / Ageing should trigger recovery decisions before value disappears 174
11.7 / Finance and operations must approve one feasible policy 175

12 / Optimisation earns authority through reproducible decisions and safe fallback 178

12.1 / The baseline should be executable without machine learning 179
12.2 / Optimisation must expose its objective and constraints 181
12.3 / Uncertainty belongs in scenarios that a planner can inspect 182
12.4 / An input error should stop an order recommendation 184
12.5 / Integration needs an idempotent decision and event log 185
12.6 / Automation authority should remain narrower than model coverage 187
12.7 / The safety-stock tool must show what it does not calculate 189

13 / A policy improves the factory only when measured service and cash improve together 192

13.1 / A replay must use only information available at the decision date 193
13.2 / Observed sales do not reveal every missed request 195
13.3 / Simulation isolates policy logic but does not validate local inputs 196
13.4 / Stress tests should include combined rather than isolated failures 198
13.5 / The field pilot needs a credible comparator 199
13.6 / The scorecard must reveal who paid for the improvement 201
13.7 / A failed gate blocks expansion until its cause is resolved 202

14 / Build stock truth first, prove replenishment second and scale only after the gate holds 205

14.1 / Set up the stock boundary before committing capital 206
14.4 / Launch with usable-stock capture and simple review rules 208
14.7 / Scale only after policy comparison survives normal exceptions 210
14.10 / Optimise network and contractual choices after the stock record is dependable 211
14.13 / Assign decisions to accountable roles rather than to a dashboard 212
14.16 / Measure the promise, the buffer and the cash on a common boundary 214
14.19 / Close with explicit conditions for changing the recommendation 217

Executive summary 12
Evidence register 222
Glossary 239
Term index 242
Method and imprint 244

Every exhibit states a testable claim.

EXHIBIT 01 / stacked
The processing gap frames the opportunity, not the stock plan 16

EXHIBIT 02 / flow
One missing component can stop a complete order 18

EXHIBIT 03 / matrix
Every stock decision needs a purpose and an owner 21

EXHIBIT 04 / table
Different evidence permits different claims 24

EXHIBIT 05 / layered-architecture
Identity, status and ownership form separate layers 31

EXHIBIT 06 / table
Unverified unit conversions can turn a physical shortage into a book surplus 34

EXHIBIT 07 / waterfall
Book stock becomes promiseable stock only after exclusions 37

EXHIBIT 08 / flow
A discrepancy stays open until its cause and correction are approved 41

EXHIBIT 09 / matrix
RFID evidence answers different questions in each study 44

EXHIBIT 10 / heatmap
A policy matrix separates value, variability and criticality 51

EXHIBIT 11 / small-multiples
Equal annual demand can produce different ordering risks 52

EXHIBIT 12 / matrix
Commercial archetypes map to different stock positions 55

EXHIBIT 13 / table
Service definitions cannot be substituted for one another 57

EXHIBIT 14 / gantt
The usable-material clock extends beyond shipment 66

EXHIBIT 15 / table
Historical port evidence cannot calibrate a current buffer 69

EXHIBIT 16 / dumbbell
Omitting open orders and partial lots understates the upper tail 72

EXHIBIT 17 / range
A mean hides the tail that exhausts stock 73

EXHIBIT 18 / matrix
Route choice changes the measurement requirement 76

EXHIBIT 19 / line
Higher cycle service buys diminishing reductions in stockout probability 82

EXHIBIT 20 / table
The safety-stock equation is conditional on explicit assumptions 86

EXHIBIT 21 / dumbbell
Adding a review gap changes required protection 87

EXHIBIT 22 / small-multiples
The same mean and variance need not imply the same tail risk 88

EXHIBIT 23 / line
The EOQ minimum is a benchmark, not a purchase instruction 95

EXHIBIT 24 / waterfall
A bulk discount must pay for the stock it creates 98

EXHIBIT 25 / table
Time-phased methods produce different terminal stock 100

EXHIBIT 26 / heatmap
Cash and storage limits can rule out the unconstrained lot 101

EXHIBIT 27 / line
The critical fractile changes when residual recovery changes 109

EXHIBIT 28 / gantt
A second order creates value only before the last useful receipt 113

EXHIBIT 29 / heatmap
A style can be overstocked while demanded sizes are missing 115

EXHIBIT 30 / flow
Residual goods require an approved recovery route 118

EXHIBIT 31 / flow
Dependent demand links one garment order to every required component 123

EXHIBIT 32 / table
Usable supply, not book stock, determines the additional requirement 125

EXHIBIT 33 / gantt
Material-ready dates must precede their consuming operations 127

EXHIBIT 34 / line
A growing queue consumes delivery time 130

EXHIBIT 35 / heatmap
A material plan can be complete while capacity is infeasible 132

EXHIBIT 36 / layered-architecture
Inventory boundaries differ from ownership boundaries 137

EXHIBIT 37 / line
Pooling benefits shrink as demands move together 139

EXHIBIT 38 / matrix
Postponement is feasible only when conversion and promise times fit 142

EXHIBIT 39 / waterfall
Freight savings can be consumed by slower replenishment 145

EXHIBIT 40 / matrix
A supplier score cannot conceal a disqualifying adverse impact 152

EXHIBIT 41 / waterfall
Scrap, sorting and buffer costs can erase a nominal price discount 154

EXHIBIT 42 / small-multiples
Batching and delayed information can amplify ordering swings 155

EXHIBIT 43 / flow
A forecast becomes a commitment only at a defined contractual event 157

EXHIBIT 44 / table
Purchasing evidence supports safeguards without proving local prevalence 161

EXHIBIT 45 / timeline
Cash leaves before the garment can generate a receivable 167

EXHIBIT 46 / waterfall
A stock reduction is not equal to a recurring profit gain 169

EXHIBIT 47 / line
A tighter cash ceiling can force a service trade-off 171

EXHIBIT 48 / table
Finance products cover different points in the cash cycle 174

EXHIBIT 49 / layered-architecture
A recommendation needs an inspectable path to approval 180

EXHIBIT 50 / table
Solver status determines the permitted next action 182

EXHIBIT 51 / flow
A retry must preserve the original business identifier 186

EXHIBIT 52 / matrix
Authority expands only within an approved envelope 188

EXHIBIT 53 / table
A longer assumed lead time raises the illustrative trigger 190

EXHIBIT 54 / flow
Each evidence stage permits a narrower claim than factory-wide success 194

EXHIBIT 55 / small-multiples
Identical assumed requests can produce different closing stock paths 197

EXHIBIT 56 / heatmap
Combined stresses change demand, release and funding together 198

EXHIBIT 57 / table
Service and cash gains remain conditional on supplier and worker safeguards 202

EXHIBIT 58 / gantt
The implementation roadmap advances through evidence gates 208

EXHIBIT 59 / matrix
Every recurring decision has an accountable role 213

EXHIBIT 60 / table
GFT would measure a balanced inventory outcome 215

EXHIBIT 61 / waterfall
State exclusions and safety buffers reconcile book value to free stock 217

EXHIBIT 62 / flow
Evidence can require a simpler policy or a changed promise 219

Build the factory and its evidence as one system.

Inventory is a commitment of cash to a particular material, location and customer promise. For an emerging-market apparel maker, the relevant question is therefore not how to minimise stock in isolation, but which stock is worth holding, for how long and under whose authority. A missing trim can stop a complete garment order while unwanted sizes occupy the warehouse. Average lead times can hide the delays that exhaust a buffer, and apparently cheaper purchasing can consume the cash needed to finish production. This paper argues for a measured decision system that begins with usable-stock truth, segments commercial promises and then selects replenishment rules suited to each material and selling window. It connects continuous and periodic review, seasonal ordering, material requirements planning, postponement and network choices to cash, capacity and supplier constraints. Foundational methods are distinguished from external field evidence and research manuscripts. Worked cases are explicitly illustrative; no operating savings or deployed GFT factory system are claimed. The Cameroon textile programme provides the design context, with its connected feasibility, performance-management and digital-strategy work. The paper closes with a proposed staged pilot, explicit stop conditions and a measurement contract covering service, stock accuracy, cash exposure and supplier and worker safeguards. It is a research and operating-method discussion, not investment, legal or tax advice.

01 Usable-stock accuracy is a prerequisite for credible replenishment recommendations. [S15] [S16] [S27] [S31]

02 Segmentation must account for service purpose and line-stop criticality as well as usage value. [S01] [S05] [S07]

03 Lead-time variability and incomplete observations can make an average-based buffer misleading. [S12] [S22] [S24] [S30]

04 Cycle service and unit fill rate require distinct definitions and calculations. [S05]

05 Quantity discounts can increase total cost once cash, holding and terminal stock are included. [S02] [S03] [S06] [S09]

06 Seasonal buying needs explicit overage, underage and recovery assumptions. [S06] [S17] [S21] [S26]

07 A material plan must reconcile component completeness with finite conversion capacity. [S07] [S19] [S28]

08 Supplier flexibility and pooling redistribute risk and can transfer hidden costs. [S08] [S10] [S11] [S17] [S20] [S25]

09 Accepted-receivable finance does not by itself establish funding for the initial materials commitment. [S18]

10 Policy value requires separate simulation, replay and field-pilot evidence, measured against service, cash and harm safeguards. [S19] [S22] [S23]

GREEN FUTURE TECH

LIMITED / HONG KONG SAR

01

PART I: DEFINE THE STOCK DECISION BEFORE BUYING THE STOCK / THE POSITION

Inventory earns its place only when it protects an explicit promise

GFT argues for measured, cash-constrained inventory decisions that protect usable material and customer commitments without manufacturing surplus.

CHAPTER 01
DECISION TEST

GFT argues for measured, cash-constrained inventory decisions that protect usable material and customer commitments without manufacturing surplus.

Inventory is a conditional commitment of cash to a particular material, form, place and date. It earns its place only when it protects an explicit customer or production promise more effectively than the available alternatives, after usable-material, capacity, labour, space and cash constraints have been respected. The correct objective is therefore not the lowest possible stock balance. It is the least costly feasible commitment that supports a stated service outcome, with every exception visible and every residual exposure assigned to an owner. [S01] [S05] [S07]

Green Future Tech Limited (GFT), Hong Kong SAR, develops this position within the design context of its three-part Cameroon textile programme: a feasibility study, performance-management system design and digital strategy, prepared 31 January 2024 with source-refresh notes retained in 2026. This chapter proposes a decision framework; it does not claim an installed factory system, a completed inventory pilot or realised GFT savings. All worked quantities and commercial terms below are illustrative. Accounting, contractual and financing choices require appropriate professional review, and this research is not investment, legal or tax advice. [S09] [S18]

1.1

A processing ambition does not determine a purchase order

Cameroon's inherited sector baseline records about 95 percent of cotton fibre exported raw and about 5 percent processed locally. The national ambition restated in SND30 is 50 percent local processing by 2030. [R06] [R07] These figures establish a strategic processing gap, not a factory's stock requirement. They do not identify which garments customers will order, which fabric constructions those garments require, which suppliers can deliver accepted material, or how much cash a particular maker can commit. A policy ambition can justify investigating industrial capacity. It cannot, by itself, authorise a purchase order.

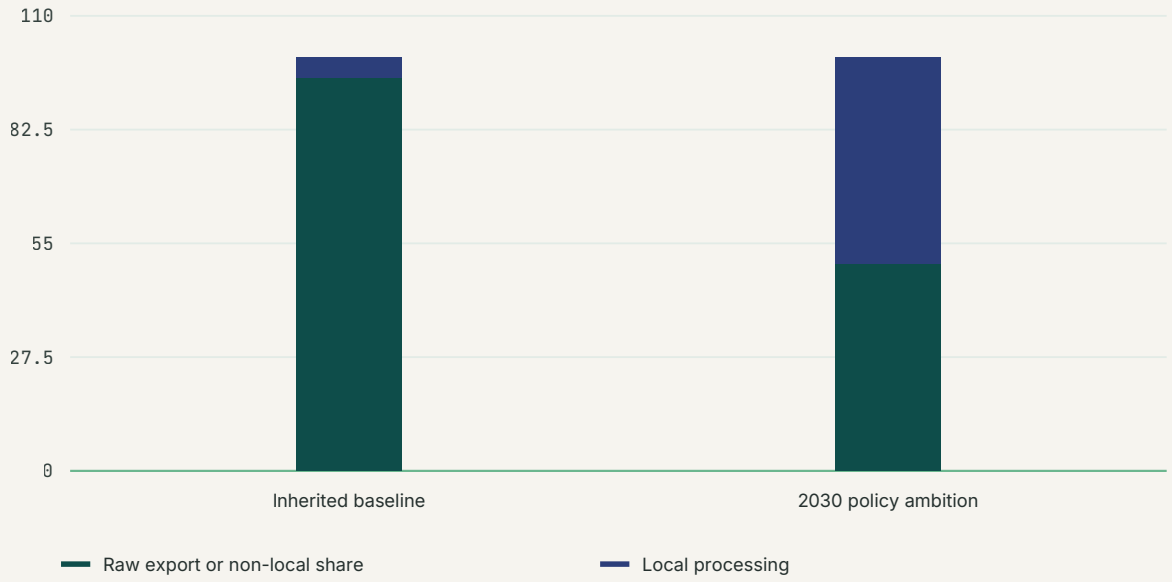
The physical objects also change as value is created. Seed cotton is an agricultural input; ginning produces fibre; spinning produces yarn; knitting or weaving produces fabric; dyeing and finishing alter performance and appearance; cutting and sewing convert fabric, thread, trims and packaging into garments. The inherited register's seed-cotton production series belongs to the first of these boundaries, not to a garment maker's fabric ledger. [R01] Each transformation has its own accepted input, output, losses, quality release and unit of measure. Material requirements planning begins with a defined parent item and its bill of materials, not with an undifferentiated crop total. [S07]

A tonne of seed cotton therefore cannot be divided by an assumed garment mass to produce a trustworthy buying plan. That shortcut would conceal ginning turnout, spinning and fabric-forming losses, finishing shrinkage, cutting utilisation, rejected material, size and colour mix, and the commercial question of whether the resulting garment has a buyer. WP06 owns the measurement of process yields. This paper accepts only quantities that have been reconciled at the relevant stocking boundary and refuses to carry an upstream total across several transformations without measured conversion evidence. [R01] [S07]

Exhibit 1, p. 16 keeps the strategic context visible while separating observation from ambition. The target-year bar is not a forecast of realised processing or raw exports. It is a policy reference plotted beside the inherited baseline.

EXHIBIT 01 / PERCENT OF COTTON FIBRE

The processing gap frames the opportunity, not the stock plan



DERIVED [R06] [R07] Inherited source-annex context, not a current factory stock requirement. Baseline: about 95 percent raw fibre export and about 5 percent local processing (R06). Policy ambition: 50 percent local processing by 2030 (R07). The 50 percent non-local-processing share in the ambition bar is the arithmetic complement 100 minus 50, included only to complete the two-category comparison; it is not a forecast of raw exports.

The exhibit cannot establish current national throughput, future demand or the quantity of inventory that a garment maker should hold. Its operational meaning is narrower: an industrial sponsor must translate ambition into identifiable promises before committing cash. For a factory, the chain runs from customer need or an explicitly approved speculative programme, to product specification, required delivery date, bill of materials, usable supply and feasible conversion. A purchase recommendation becomes defensible only when those links are recorded. [S01] [S07]

> **GFT NOTE**
>
> Assign every proposed buffer four attributes before approval: a service purpose, a usable-stock boundary, a cash owner and a review rule. If the purpose is only "support industrialisation", the buffer is not yet an inventory policy. Return it to the commercial and production owners until they can name the promise it protects and the condition that would release or remove it.

This requirement does not prohibit buying ahead of a confirmed order. A factory may rationally hold generic fabric for a repeat programme, reserve capacity against a tender, or create a speculative seasonal range. The distinction is that each commitment needs an explicit hypothesis: who is expected to buy, in what form, by when, under which selling window, with what residual recovery and within what cash limit.

Segment-specific service rules and seasonal exit decisions are developed later in the paper. At this stage, the essential control is to call an option on future demand an option, rather than presenting it as a stock requirement mechanically derived from national policy. [S01] [S06]

1.2

A missing trim can immobilise an otherwise complete garment

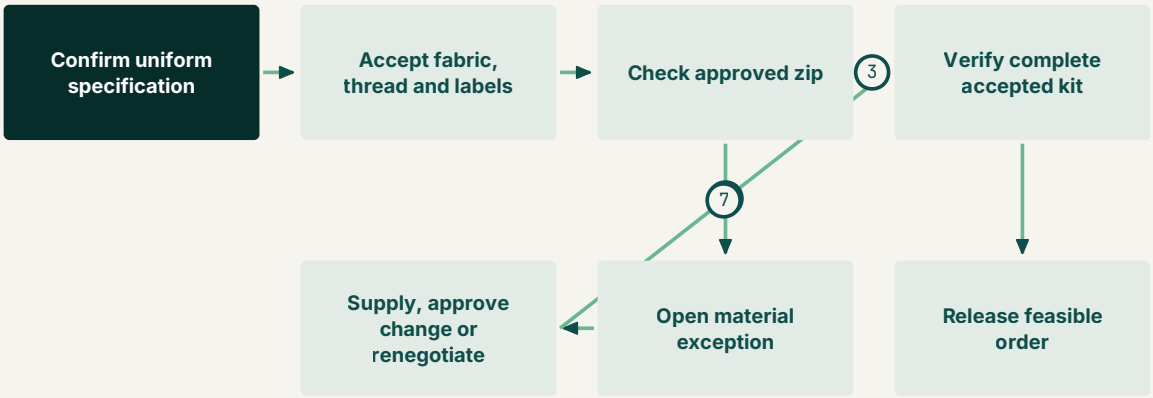
A customer buys a complete garment, not an aggregate value of warehouse contents. Fabric, thread, labels, packaging and approved trims have dependent demand because the parent order consumes them together through a versioned bill of materials. [S07] A warehouse can therefore contain ample cloth and valuable finished goods while remaining unable to complete a particular uniform order. The operational test is not whether total stock exceeds total order value. It is whether every required component, in its accepted specification and required quantity, can be complete at the consuming operation before the promised date.

Consider an illustrative institutional uniform order whose shirts require approved fabric, matching thread, woven size labels, branded buttons and a specified collar zip. The first four component families may be accepted and staged, yet a missing zip prevents completion of the garment design. Purchasing a different zip from an unapproved supplier is not a neutral warehouse substitution. It changes the approved product and may alter function, appearance or customer compliance. The planner must either obtain the specified zip in time, secure an authorised specification change, negotiate a revised delivery promise, or decline the order. More fabric cannot repair that constraint. [S07] [S20]

Exhibit 2, p. 18 reduces the decision to its necessary branches. It is a proposed control flow, not a depiction of an existing GFT system. Its central gate asks whether the complete, accepted kit exists before the factory gives a firm production and delivery commitment.

EXHIBIT 02 / PROPOSED ORDER-RELEASE DECISION

One missing component can stop a complete order



3. Available
4. Missing
7. Resolved

ILLUSTRATIVE [S07] Illustrative uniform-order workflow based on dependent requirements and bill-of-material completeness (S07). Fabric, thread and labels are assumed accepted; the approved zip is unavailable. Release waits for an accepted receipt, an authorised specification change or a revised customer promise. No automatic substitution or installed GFT workflow is claimed.

The gate also prevents a misleading form of efficiency. Loading the line with incomplete orders can create work in progress that waits for a minor component, occupies space and complicates sequencing without increasing completed delivery. Releasing only feasible kits is not a universal rule for every production system, because controlled partial work may be sensible where later operations can be sequenced around the shortage. The promise owner must nevertheless see the exception. An apparently busy factory is not necessarily a factory fulfilling orders. [S07]

This completeness test depends on clear stock categories. **Raw materials** have not entered the defined conversion process. **Work in progress** has entered but not completed that process. **Finished goods** have completed production and await the relevant release, sale or delivery event. These categories describe process position; they do not by themselves establish quality acceptance, legal ownership or availability to promise. [S07] [S09]

Three further categories defeat a simple warehouse-total view. **Supplier-owned stock** may be physically present but usable only under agreed call-off, payment and title conditions. **Goods in transit** may already consume cash while remaining unavailable to the line. **Customer returns** may require inspection, repair or disposition before any quantity can re-enter saleable stock. Contract terms and accounting treatment must be evaluated separately from physical location. [S08] [S09]

Customer-supplied material adds another boundary. In an assembly arrangement, the garment maker may hold the customer's fabric and trims for a defined order. Diverting those inputs to another buyer would not be made acceptable by an urgent delivery problem or by a positive book balance. Conversely, a supplier's consignment quantity should not be treated as free factory stock merely because it sits in the same rack. Every reservation must respect ownership, specification and permitted use as well as quantity. [S07] [S08]

Returns deserve the same discipline. A returned garment is not automatically equivalent to a newly finished first-quality garment. Until its condition and resale eligibility are established, it is an uncertain recovery asset with handling cost, not a unit that can silently offset a new order. IAS 2 provides an accounting measurement principle, lower of cost and net realisable value, but it does not decide physical usability or the commercial promise. [S09] The decision system needs both perspectives without confusing them.

1.3

The inventory objective must include cash and delivery

An inventory policy is an economic choice under constraints. Its relevant costs can include purchase price, ordering and approval effort, transport and receipt, holding, shortage, expiry or obsolescence, and expediting. Its constraints include the promised service outcome, available cash, safe labour practice, storage and conversion capacity. The optimisation question is which feasible combination of quantity, timing, location and material form best protects the promise, not which isolated unit quote is cheapest. [S02] [S03] [S05] [S20]

Service belongs inside that formulation as an explicit event. A promise might require a complete order by a confirmed date, immediate availability of a requested variant, or continuity of a critical component during a replenishment cycle. These events are not interchangeable. Probabilistic inventory methods distinguish cycle service from unit fill rate for precisely this reason. [S05] A factory that reports only a high aggregate availability percentage may still miss complete orders if shortages concentrate on one required trim or one demanded size.

Cash is equally specific. The relevant exposure is not merely the accounting value shown at month end. It includes deposits, supplier balances, freight and other committed payments, material in transit, quality-held goods and the delay before a customer accepts and pays an invoice. A policy can reduce expected shortages and still be infeasible because it consumes cash needed for wages, conversion or another order. The planner and finance owner must therefore review one common commitment set rather than separate, incompatible stock and cash totals. [S02] [S09] [S18]

Accounting valuation answers a different question. IAS 2 states that inventories are measured at the lower of cost and net realisable value. [S09] That principle informs reporting and write-down exposure. It does not calculate the economic cost of placing the next order, the cash already committed, or the operational loss from an incomplete kit. Decision models must also avoid double counting: a financing charge

included in holding cost should not be added again as an unexplained penalty, and an accounting write-down should not be presented as a new cash payment when the cash left the business earlier.

Quantity discounts make the distinction tangible. Total-relevant-cost analysis compares a price reduction with the ordering, holding and feasibility consequences of the quantity required to obtain it. [S03] A lower unit price can increase the initial commitment, create non-interchangeable residual stock and delay recovery. Accepted-receivable finance does not automatically remove that problem. The opened IFC programme description concerns post-shipment finance based on accepted receivables, which is a different point in the cash cycle from the initial fabric or trim purchase; it is not evidence of eligibility or funding for GFT or a proposed venture. [S18]

> **WORKED EXAMPLE: A LOWER ZIP PRICE RAISES THE COMMITMENT**
>
> An illustrative uniform order requires 1,000 approved zips. Offer A sells 1,000 zips at XAF 500 each, a purchase commitment of XAF 500,000. Offer B charges XAF 400 each but requires a minimum lot of 3,000 zips, a commitment of XAF 1,200,000. All prices, quantities and terms are teaching assumptions, not supplier quotations. [S03]
>
> Offer B has a 20 percent lower unit price, calculated as (500 minus 400) divided by 500, but requires XAF 700,000 more initial cash and leaves 2,000 zips beyond the order requirement. If those zips have no approved alternative use, their recovery is uncertain rather than free profit. Under an illustrative later recovery of XAF 100 each, the residual proceeds would be XAF 200,000 and the net purchase outlay after that recovery would be XAF 1,000,000, still XAF 500,000 above Offer A before holding, handling or collection costs. [S03] [S09]
>
> The contrary error is to buy too few zips in the name of cash discipline. An illustrative purchase of 950 zips would leave the final 50 uniforms incomplete. If the customer requires one complete delivery, the whole order can miss the agreed order-level promise even though 950 garments are finished. The correct comparison therefore includes both the XAF 475,000 small-lot cash saving in this constructed case and the service consequence of the missing 50 components; neither stock minimisation nor price minimisation is sufficient. [S05] [S07]

This example does not imply that small lots are always preferable. A larger lot can be rational when the component is genuinely common, future demand is sufficiently evidenced, supplier terms are fair, cash is available and the residual has a credible recovery route. Nor does it prove that a discount is worth a fixed maximum premium. The decision changes with demand, substitution, lead time, storage, supplier reliability and cash opportunity cost. The worked arithmetic simply prevents the word "discount" from concealing the full commitment. [S03] [S08]

Exhibit 3, p. 21 shows why the first screening conversation should combine delivery consequence with cash exposure. The coordinates are deliberately ordinal teaching scores, not measurements. They organise questions; they do not calculate a reorder point or rank actual GFT materials.

EXHIBIT 03 / ILLUSTRATIVE ORDINAL SCREENING SCORE

Every stock decision needs a purpose and an owner



ILLUSTRATIVE [S01] [S05] [S27] Illustrative GFT decision charter informed by segmentation (S01), explicit service definitions (S05) and accountable discrepancy review (S27). Both axes run from 0 to 10 and are invented ordinal scores, not measured probabilities or monetary values. Each real item-location decision must separately name its service owner, cash owner, usable-stock boundary and next review; a matrix position does not assign those authorities.

A low-cost zip can sit beside a high-cost fabric on the delivery axis because either can block completion. Speculative finished garments may combine cash exposure with uncertain recovery, while generic packaging may tolerate simpler review if approved substitutes exist. The proper response is not automatically more stock for every high point. Alternatives include supplier call-offs, an approved substitute, a changed promise, a smaller production release or a controlled purchase exception. The matrix makes the owner of that choice visible. [S01] [S08]

1.4

The base case must fit a garment maker rather than an integrated mill

The programme design base for this paper is cut, make, trim and finishing. Spinning and knitting are conditional partnered stages, not installed GFT operations. This boundary matters because the garment maker's immediate stock decisions concern approved fabric, trims, work in progress and finished garments, supported by dependable specifications and receipts. It does not begin by deciding how much seed cotton or spinning capacity to own. The three-part programme connects feasibility, performance management and digital strategy at this operating boundary without asserting deployment.

An integrated mill would face a different set of decoupling points and physical conversions. It would need to coordinate fibre, yarn, greige or finished fabric, wet-processing queues and downstream garment requirements. That may be a valid partnered configuration after its technical and commercial case is proved, but import-

ing its stock architecture into the base case would obscure the first decisions a garment maker must control. A model should fit the operation being governed, not the most expansive possible value chain. [S07]

Within the garment-making boundary, three illustrative contract archetypes create different promises:

- **Make-to-stock local apparel** begins with an approved demand hypothesis. The factory commits material and conversion before individual buyers are known, then holds finished variants for rapid fulfilment. Its central exposures are variant mix, selling-window expiry and residual recovery. The service promise concerns immediate availability of defined products, not merely completion of production. [S01] [S06]
- **Make-to-order uniforms** begins with a confirmed specification, quantity and delivery requirement. Fabric and trims may be order-specific and approved substitutions limited. The central exposures are incomplete kits, usable-receipt timing, customer changes and cancellation terms. Finished stock is not the intended buffer if the order has a known recipient. [S07] [S08]
- **Customer-supplied-material assembly** places some purchasing and ownership decisions outside the factory while leaving receipt verification, custody, yield reconciliation and delivery execution inside the operating system. The factory must not treat customer-owned inputs as free supply for another contract. Its cash exposure and its authority to reject or replace defective material differ from the other archetypes. [S07] [S08]

These archetypes can coexist in one factory. That does not permit one undifferentiated stock pool. The same nominal cloth may be forecast-owned for a local range, reserved for a uniform order, or supplied under custody for an assembly customer. Planning may identify technical commonality, but commercial release requires compatible ownership and approval rights. Later chapters examine segmentation, pooling and postponement in detail; the opening rule is simply that physical similarity does not erase a contractual boundary. [S01] [S08]

The archetype also changes what counts as a shortage. In make-to-stock apparel, an unavailable requested variant can become a lost sale. In make-to-order production, a missing component can delay a confirmed order and trigger expediting or renegotiation. In customer-supplied assembly, an unusable receipt can create a custody discrepancy and a demand for replacement outside the factory's ordinary purchasing authority. A single service formula applied without the contract context would report different failures as if they were the same event. [S05] [S07]

> **METHOD**
>
> Start every item-location policy with four boundary statements: the commercial archetype, the process stage at which stock is recognised, the party that owns the quantity, and the event that makes it usable for the named promise. Add the bill-of-material and specification version before calculating need. This is a proposed GFT design discipline based on dependent-requirement and contract-allocation methods, not a claim that one software product is required. [S07] [S08]

1.5

Evidence determines how strongly a recommendation can be stated

A decision paper must distinguish a method from an observed result. An accounting standard can define measurement, a teaching text can establish a conditional model, a field study can report what happened in its sample, and a research manuscript can examine an algorithm under stated assumptions. None of these, alone, proves that a proposed policy will save money in a Cameroon garment factory. The wording of a recommendation should carry the strength and boundary of its evidence. [S01] [S09] [S22] [S23]

Official and foundational sources are indispensable but limited. IAS 2 supports the distinction between accounting valuation and operational decision cost; it does not prescribe a stock quantity. [S09] MIT OpenCourseWare provides methods for segmentation, deterministic and probabilistic inventory models, dependent requirements and contracts. Those methods allow transparent calculations, but their assumptions must be checked against local demand, supply and contract conditions before their outputs become orders. [S01] [S02] [S05] [S07] [S08]

Field evidence answers a different question. The ILO's 2017 purchasing-practices survey reports 1,454 supplier responses across 87 countries. [S10] That sample supports attention to the relationship between purchasing and working conditions, not an estimate of current prevalence in Cameroon. Likewise, a fashion-retail conference study comparing an RFID cycle count with a physical count supports validating a capture method against physical truth. It does not establish a savings rate or read accuracy for a garment factory with different packaging, layout and processes. [S15]

Historical place-specific evidence needs equal care. The World Bank's Douala publication provides a historical diagnosis of port efficiency and investment constraints. [S14] It justifies measuring clearance, transport and receipt as separate stages. It does not provide a current lead-time parameter for a new purchase policy. Using an old corridor observation as if it were a present supplier distribution would convert a useful diagnostic prompt into an unsupported buffer calculation. Current quotations, route observations and usable-receipt timestamps must replace inherited assumptions. [S14] [S30]

Research manuscripts help formalise failure modes without supplying local operating proof. Work on lost sales and uncertain supplies identifies censoring: observed demand and supply can be truncated by availability. [S22] A simulation-optimisation manuscript illustrates repeated stochastic comparison of inventory policies. [S23] These contributions inform what an evaluation should measure. They do not make unobserved demand known, validate a factory's input distribution or establish that a simulated ranking will survive real supplier and customer behaviour.

Exhibit 4, p. 24 sets these evidence classes beside their permitted use. The table is a source-bound reading guide, not a quality ranking. A carefully executed simulation can be more informative than a poorly matched field sample, but only for the question its design actually addresses.

EXHIBIT 04 / SOURCE METHOD AND PERMITTED INFERENCE

Different evidence permits different claims

Evidence	Method	Permitted use	Boundary
IFRS, IAS 2 (S09)	Standard overview	Separate valuation from ordering	No optimal stock rule
ILO survey, 2017 (S10)	Supplier questionnaire	Motivate purchasing safeguards	No current Cameroon prevalence
World Bank Douala, 2015 (S14)	Historical diagnostic	Prompt staged time measurement	No current route parameter
RFID cycle-count study, 2013 (S15)	Fashion-retail comparison	Test capture against physical counts	No transferable factory saving
Lost-sales manuscript, 2023 (S22)	Theoretical inventory research	Retain censored-data uncertainty	No local demand observation
Simulation manuscript, 2024 (S23)	Stochastic policy comparison	Design repeatable evaluations	No field-performance guarantee

SOURCED [S09] [S10] [S14] [S15] [S22] [S23] Sourced methodological comparison using the opened register entries. Dates and study scopes remain attached to their sources. No source establishes a GFT factory saving or a deployed inventory system.

This discipline also determines what must be absent. Missing annex references are excluded rather than cited, and unsupported tariff or training claims are not carried into the stock argument. The source register's verification limitations remain visible instead of being converted into present-tense certainty. Where no current local observation exists, the chapter uses a labelled illustrative case or states the input required; it does not fill the gap with a precise number.

> **GFT NOTE**
>
> Label each material claim in the decision record as **sourced observation**, **derived arithmetic**, **external study finding**, **research proposition** or **illustrative assumption**. Attach the observation date and denominator. When evidence changes, change the strength of the sentence as well as the calculation. A confident interface should never make an uncertain input look observed.

1.6

The paper begins where a forecast becomes an action

A forecast estimates uncertain demand; an inventory policy converts that estimate into a commitment. WP04 owns demand distributions and forecast residuals. This paper begins at the decision interface and chooses quantity, timing, location, material form and approval conditions. It also decides what not to buy, which uncertainty to carry, and when a planner must return to the commercial owner. [S05] [S06] [S07]

That separation prevents two common errors. The first is to treat a central forecast as a deterministic order, ignoring lead-time variation, usable stock and lot constraints. The second is to add an unexplained buffer to every forecast in the name of caution. A high expected demand does not justify a purchase if the required material cannot arrive in time or the cash limit is binding. A low expected demand does not imply zero stock when a confirmed promise, a long protection period or a line-stopping component changes the consequence of shortage. [S03] [S05] [S07]

The decision interface must therefore carry more than a demand number. It needs the forecast vintage, requested service event, stock and reservation boundary, open supply, usable-receipt distribution, approved substitutions, order bounds, cash limit and expiry of the recommendation. Later chapters derive specific replenishment and seasonal rules. Their common obligation is to expose the assumptions that transform an estimate into an action. [S05] [S22] [S30]

Other papers retain their own subjects. WP05 owns detailed traceability qualification, WP09 labour and supervisory capability, WP07 energy and infrastructure architecture, and WP13 financing structures. Their constraints remain present here: stock cannot be promised without usable identity, service cannot be protected by assuming harmful working practices, and a materials commitment is not funded merely because a future receivable is imaginable. [S10] [S18] [S20]

1.7

The decision record makes the thesis testable

The thesis becomes operational only when each recurring decision leaves an inspectable record. GFT proposes an item-location policy record with a named owner, input vintage, approved service definition, order bounds, exception route and next review. This can begin as a controlled document and structured register; it does not require a sophisticated platform. The essential feature is that another reviewer can reconstruct why a quantity was recommended, which commitments it protected and what would invalidate it. [S27] [S30]

The record should answer seven questions before approval:

1. **Promise:** which customer, production or selling-window event is protected, and how is success counted? [S05]
2. **Object:** which item, specification, material form, ownership state and location are being committed? [S07] [S08]
3. **Evidence:** which demand, stock, lead-time, price and contract inputs were available at the decision date? [S22] [S30]
4. **Authority:** who may approve the service target, cash exposure, substitution and any override? [S27]
5. **Bounds:** which minimum lot, pack, capacity, space, labour and cash constraints make the order feasible? [S03] [S20]
6. **Exception:** which missing or contradictory input blocks automatic recommendation, and who resolves it? [S27]

7. **Review:** when does the policy expire, and which observed changes require earlier reassessment? [S05] [S30]

A decision record also makes overrides useful evidence. If a planner repeatedly increases a zip order because approved receipts arrive short, the correct response is not to hide the manual adjustment. It is to investigate supplier fulfilment, specification approval and the receipt boundary, then revise or suspend the policy. If finance repeatedly rejects an otherwise feasible recommendation, the business must reconcile its service promise with its cash constraint rather than allow two authorities to issue incompatible instructions. [S03] [S27]

The falsifiable criterion follows directly. Release less cash to inventory only if measured service and material completeness remain within agreed limits on a common comparison boundary. The evaluation must include order fulfilment, complete-kit failures, aged or residual stock, expediting and cash-at-risk, with supplier-payment and labour safeguards so that an apparent factory improvement does not conceal transferred harm. [S05] [S09] [S10] [S20] No universal saving is prescribed, because the allowable reduction depends on the baseline, demand, supply, contract and uncertainty actually observed.

A failed condition blocks expansion. Unresolved stock identity, repeated incomplete kits, a breached cash ceiling, stale inputs or an overdue policy review should return the recommendation to human reassessment and, where necessary, a simpler manual rule. Chapter 13 defines the replay, simulation and field-pilot evidence needed to compare policies; chapter 14 sets the staged expansion gates. This opening chapter supplies the proposition those later tests must try to disprove.

Stock is justified by the promise it protects, the cash it can safely commit and the evidence that remains when the forecast is wrong.

CHAPTER SOURCE NOTES

S01 MIT OpenCourseWare. 2006.

S05 MIT OpenCourseWare. 2006.

S07 MIT OpenCourseWare. 2006.

S09 IFRS Foundation. n.d..

S18 International Finance Corporation. 2026.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R01 Business in Cameroon. Publication date unrecorded; inherited source-text check.

S06 MIT OpenCourseWare. 2006.

S20 OECD. 2018-03-07.

S08 MIT OpenCourseWare. 2006.

S02 MIT OpenCourseWare. 2006.

S03 MIT OpenCourseWare. 2006.

S27 North Carolina Office of the State Controller. 2023-10-24.

S22 arXiv, Chen, Jiang, Zhang and Zhou. 2023-10-31.

S23 arXiv, Sarit Maitra. 2024-06-27.

S10 International Labour Organization. 2017-06-09.

S15 Association for Information Systems Electronic Library. 2013.

S14 World Bank. 2015-02-12.

S30 World Customs Organization. 2025.

02

PART I: DEFINE THE STOCK DECISION BEFORE BUYING THE STOCK / STOCK TRUTH

Only identifiable, usable and reconciled stock can support a promise

An inventory policy fails when it orders against quantities that are wrongly identified, reserved twice or unavailable for use.

CHAPTER 02
DECISION TEST

An inventory policy fails when it orders against quantities that are wrongly identified, reserved twice or unavailable for use.

An inventory policy can only protect a customer promise when the stock record behind it identifies the exact material that production can substitute, converts physical units without distortion, separates quality-held or reserved goods from free supply, and reconciles physical movements before new commitments are made. A reorder point, a material requirements plan or an order confirmation is an arithmetic operation applied to a ledger balance. If that balance includes rolls from an incompatible dye bath, fabric awaiting incoming inspection, or quantities already promised to another work order, the formula will suppress replenishment while the sewing line starves. [S01] [S05] [S07] Green Future Tech Limited (GFT), Hong Kong SAR, therefore treats usable-stock truth as the first operational gate in any inventory system: before choosing a statistical buffer or an optimisation rule, the factory must establish an event-linked ledger in which identity, unit conversion, stock state, ownership and discrepancy resolution are governed by explicit engineering and accounting controls. [S27] [S31] [S0203]

This chapter specifies that ledger architecture for an emerging-market garment maker within the design context of GFT's three-part Cameroon textile programme (feasibility study, performance-management system design and digital strategy, prepared 31 January 2024 with source-refresh notes retained in 2026). Whereas the national baseline records about 95 percent of cotton fibre exported raw and about 5 percent processed locally alongside a policy target of 50 percent local processing by 2030, a garment factory operating at the cut, make, trim and finishing stage cannot manage stock as undifferentiated fibre tonnage. [R06] [R07] Every numerical case, roll conversion and stock balance in this chapter is illustrative and constructed solely to demonstrate ledger mechanics. No live factory ledger, warehouse audit or deployed GFT software performance is claimed, and nothing here constitutes investment, legal or tax advice. [S09] [S27]

2.1

An item code must describe the material that production can substitute

An item code in a garment factory is an engineering equivalence rule, not a purchasing shorthand. Two materials belong under the same stock-keeping unit (SKU) only if the cutting table and sewing line can substitute one for the other without altering the customer specification, the garment fit, the visual appearance or the approved bill of materials. [S01] [S07] When a warehouse groups distinct materials under a broad description such as "navy cotton twill" or "white core-spun thread", the ledger creates an illusion of pooled availability. A planner who sees an illustrative book balance of 2,000 metres of navy fabric and releases a 1,200-garment institutional uniform order will cause a line stop if those 2,000 metres consist of four incompatible shade lots, two different cuttable widths and an unapproved weave construction. [S07]

For finished and semi-finished garments, item identity requires style code, garment revision, colourway, size and pack configuration. For fabrics, item identity requires fibre composition, yarn count, weave or knit construction, nominal areal weight in grams per square metre (g/m2), nominal cuttable width between selvedges, finish treatment, colour reference and the list of qualified suppliers. [S07] For trims and packaging, identity must capture functional dimensions and compliance markings: a

zip requires teeth material, gauge, exact length, slider lock type, tape colour and approved vendor; a care or size label requires language version, artwork revision and customer brand authorisation. [S07] [S20] If any of these attributes changes in a way that prevents free mixing on the production line, the material requires either a distinct item code or a mandatory lot-level restriction that the planning engine respects during netting. [S07]

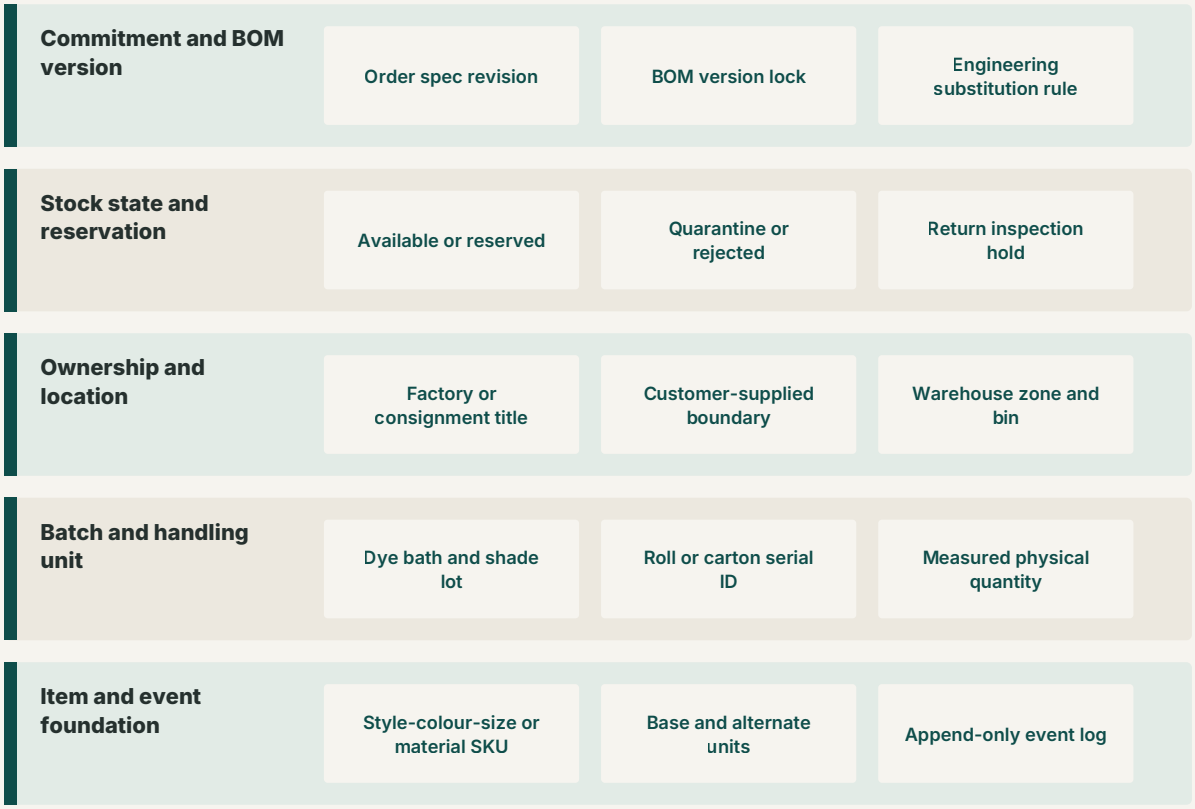
Even when two rolls share an identical item code and supplier specification, textile dyeing and finishing introduce batch-to-batch variation that prevents arbitrary mixing within a single garment or across matching pieces of a uniform set. A ledger that tracks only item code and aggregate quantity cannot prevent a cutter from laying rolls of noticeably different shade cast into the same bundle. GFT therefore separates five ledger dimensions that must never be collapsed into a single text field:

1. **Item specification identity**, which defines the engineering and commercial equivalence class (SKU) together with its active specification version. [S07]
2. **Lot and shade identity**, which records the supplier production batch, dye-bath number and inspected shade band (A1, A2, B1, for example, in an illustrative shade-sorting scheme) so that cutting markers draw from compatible shade groups. [S07] [R29]
3. **Handling-unit serial identity**, which assigns a unique barcoded or tagged identifier to each physical roll, cone carton or garment tote so that individual roll lengths, splice marks and bin movements remain traceable. [S15] [S31]
4. **Location identity**, which maps the warehouse zone, rack, bin or production staging lane where the handling unit physically rests. [S27]
5. **Ownership and encumbrance**, which distinguishes factory-owned stock from supplier-consigned stock and customer-supplied material (CMT or assembly contracts) that cannot legally be diverted to another buyer's order. [S07] [S09]

Exhibit 5, p. 31 illustrates how these dimensions sit in distinct architectural layers linked by immutable transaction events. The layered design is a proposed GFT control schema informed by material requirements planning (S07) and supply-chain event reconciliation (S31), not a depiction of an installed warehouse system.

EXHIBIT 05 / PROPOSED LEDGER SCHEMA LAYERS

Identity, status and ownership form separate layers



ILLUSTRATIVE [S07] [S31] Illustrative ledger schema proposed by GFT. Separates master specification and version control, batch and shade-lot attributes, handling-unit location and legal ownership, stock usability state, and event-linked transaction logs. Informed by MRP bill-of-material control (S07) and ship-receive event reconciliation (S31); no installed GFT factory database is claimed.

Separating these layers protects the factory when specifications evolve mid-season. Suppose a buyer updates a uniform shirt specification from revision v1 to revision v2 by changing the pocket button diameter and the care-label composition text. If the warehouse ledger ties reservations only to the parent style name, the planning run will net the new v2 order against obsolete v1 labels and buttons sitting in the bin. Retaining both the item specification version and the bill-of-material (BOM) version on every purchase order, quality release and production reservation forces an explicit engineering decision: either the technical lead signs a written concession allowing the v1 trim to be consumed on the v2 order, or the v1 trim remains excluded from the v2 net requirement calculation. [S07] [S0203] Substitution is therefore recorded as a dated, auditable engineering approval with a named role, never as an informal bin swap on the shop floor. [S01] [S07]

2.2

Unit conversions must preserve physical quantities

A stock ledger fails silently when purchasing, warehousing and production measure the same material in incompatible units and convert between them using unverified averages. In apparel manufacturing, unit changes occur at almost every physical boundary. Knitted cotton jersey is frequently purchased from spinning and knitting mills by conditioned mass in kilograms; woven shirting and trouser twill are purchased by linear metre or roll; sewing thread is purchased by cone with a nominal length in metres; buttons and zips are purchased by gross or carton; and the cutting and sewing rooms schedule bundles in garment pieces. [S07] If the conversion factors linking these units omit width variations, areal-weight drift, splice allowances or roll-end remnants, the ledger will report positive book metres after the physical racks are empty. [S07] [S0203]

The physical relationship between fabric mass and linear length shows why a fixed conversion factor is hazardous. For a roll of fabric with conditioned mass M_{kg} (in kilograms), actual areal weight w (in grams per square metre, g/m^2) and actual cuttable width b (in metres), the physical linear length L_{m} (in metres) is: [S07]

$$L_{\text{m}} = \frac{1{,}000 \times M_{\text{kg}}}{w \times b}$$ [S07]

If a planner converts an invoiced delivery of knitted or woven fabric from kilograms to linear metres using only the nominal catalogue values w_0 and b_0, two physical errors enter the stock record simultaneously. [S07] First, textile finishing tolerances allow areal weight w to vary by several percent around the nominal specification; a heavier-than-nominal fabric yields fewer linear metres per kilogram. Second, if the cuttable width b between pinhole selvedges is narrower than the nominal marker width b_0, the cutting room must either re-lay a narrower marker that consumes more linear metres per garment or reject the narrow roll. Furthermore, aggregate linear metres across multiple rolls cannot be divided directly by average garment consumption if individual roll lengths leave unusable end remnants shorter than a single marker lay. [S07]

Just as a fabric roll cannot be converted by a single unverified multiplier, upstream agricultural figures must never be converted directly into garment stock plans. The inherited evidence register records Cameroon seed-cotton production at 330,000 tonnes in 2022, 314,455 tonnes in 2023 and about 340,000 tonnes in 2024, with a 2025 projection above 350,000 tonnes. [R01] [R02] Those figures describe harvested seed cotton entering ginning, not spun yarn, finished fabric or cuttable garment kits. Ginning turnout from seed cotton to lint fibre, carding and combing waste in spinning, knitting or weaving loss, wet-processing shrinkage in dyeing and finishing, and marker utilisation in cutting each occur at separate physical boundaries with distinct yield distributions. In the GFT white-paper library, WP06 owns the measurement of process yields and conversion losses across those stages; within WP12, the inventory ledger records only the measured physical input and accepted output at each stocking boundary, never a universal cotton-to-garment conversion constant. [R01] [S07]

WORKED EXAMPLE: RECONCILING INVOICED MASS, ROLL LENGTH AND MARKER YIELD

Consider an illustrative delivery of cotton twill ordered for a workwear trouser batch requiring 1,080 garments. All figures are teaching assumptions constructed to demonstrate unit-conversion reconciliation, not measured factory parameters. [S07]

1. **Nominal purchasing conversion**: The supplier invoice states a net mass of $M_{\text{kg}} = 500\text{ kg}$ at a nominal specification of $w_0 = 200\text{ g/m}^2$ and nominal cuttable width $b_0 = 1.50\text{ m}$. Using the nominal formula, an unverified book ledger posts $L_{\text{nom}} = (1{,}000 \times 500) / (200 \times 1.50) = 1{,}666.67\text{ m}$ across 17 rolls. At a nominal marker consumption of 1.50 m per trouser, the book ledger reports enough fabric for $\lfloor 1{,}666.67 / 1.50 \rfloor = 1{,}111$ trousers, comfortably above the 1,080-garment order. [S07]
2. **Physical roll-by-roll inspection**: At incoming inspection, the warehouse weighs and measures the rolls. Actual average areal weight is $w = 205\text{ g/m}^2$ (within a 3 percent commercial weight tolerance) and actual cuttable width is $b = 1.48\text{ m}$. Gross physical length across the 17 rolls is therefore $L_{\text{gross}} = (1{,}000 \times 500) / (205 \times 1.48) = 1{,}648.00\text{ m}$, which is 18.67 m shorter than the nominal conversion. [S07]
3. **Defect and remnant deductions**: Inspection logs 12 major weave defects across the rolls (manual inspection typically achieves 60 to 70 percent accuracy while computer-vision systems reach above 90 percent across more than 40 defect types in published literature, which is why the inspection method itself must be recorded). [R29] Each spliced or cut-out defect requires an illustrative 1.25 m splice loss ($12 \times 1.25 = 15.00\text{ m}$). In addition, because cutting lays require 9.12 m per 6-garment marker (1.52 m per garment at the narrower 1.48 m width), roll-end remnants shorter than 9.12 m across the 17 rolls total 29.04 m that cannot form a complete marker lay. [S07]
4. **Net usable kit capacity**: Net usable length in full marker lays is $1{,}648.00 - 15.00 - 29.04 = 1{,}603.96\text{ m}$, which supports exactly 175 lays of 6 garments, or $1{,}050$ trousers ($1{,}050 \times 1.5276\text{ m}$ gross equivalent, or $175 \times 9.12 = 1{,}596.00\text{ m}$ with 7.96 m residual). What looked like a 31-garment surplus under nominal kilogram-to-metre conversion is actually a 30-garment shortage against the 1,080-garment promise. [S07]

Exhibit 6, p. 34 summarises the illustrative reconciliation from invoiced weight to cuttable garment capacity. The table shows why the ledger must store both the purchase unit and the verified production unit on every handling unit.

EXHIBIT 06 / ILLUSTRATIVE FABRIC METRES AND TROUSER PIECES

Unverified unit conversions can turn a physical shortage into a book surplus

Measurement stage	Basis and assumptions	Fabric (m)	Garment capacity (pieces)	Variance vs 1,080 order
Nominal invoice conversion	500 kg at 200 g/m2, 1.50 m width, 1.50 m/pc	1.67e+03	1,111	31
Measured gross roll length	500 kg at 205 g/m2, 1.48 m actual width	1,648	1,084	4
Less defect splice allowance	12 logged defects at 1.25 m splice loss each	1,633	1,074	-6
Less unusable roll-end remnants	17 roll tails shorter than a 9.12 m 6-pc lay	1,596	1,050	-30

ILLUSTRATIVE [S07] [R29] Illustrative reconciliation of a 500 kg fabric receipt across 17 rolls against a 1,080-garment order. Nominal conversion uses 200 g/m2, 1.50 m width and 1.50 m/pc (1,666.67 m; 1,111 pcs). Inspection measures 205 g/m2 and 1.48 m width (1,648.00 m gross), deducts 15.00 m for 12 splices (1.25 m each) and 29.04 m in roll remnants, and applies 1.52 m/pc in 9.12 m 6-pc lays to yield 1,596.00 m (175 lays; 1,050 pcs). Context: S07, R29. All figures are teaching assumptions.

To prevent this failure mode, GFT proposes a strict dual-unit ledger rule: every roll, cone or carton is received in its commercial invoice unit (`kilograms`, `cones` or `cartons`) for accounts-payable matching, but cannot enter `available` production stock until its actual production unit (`inspected usable linear metres`, `cuttable width` and `shade lot`) is recorded at the handling-unit level. [S07] [S31] Any discrepancy between ticketed roll length and measured roll length outside the supplier's contractual tolerance remains attached to the supplier lot record for commercial claim or replenishment adjustment rather than vanishing into cutting-floor scrap. [S27] [S31]

2.3

Stock states prevent quality holds becoming delivery promises

Even when item identity and physical units are exact, a stock ledger will mislead planners if it treats every physical unit inside the warehouse fence as available to promise. At any moment, a factory warehouse contains materials in fundamentally different operational and legal states: cartons just unloaded on the receiving apron, rolls awaiting laboratory shrinkage or shade testing, accepted materials already allocated to a confirmed work order, defective lots awaiting return to a vendor, and customer-returned garments awaiting inspection. [S07] [S09] [S26] Collapsing those quantities into a single on-hand figure causes the planning system to double-promise scarce components and delay reorders until the cutting room physically discovers that the bin holds only quarantined stock. [S05] [S07]

GFT's ledger specification enforces six mutually exclusive status states for physical inventory on site, alongside a separate pipeline state for goods in transit:

- **In-transit (`Q_trans`)**: Material dispatched by the supplier against an open purchase order and supported by shipment documentation, but not yet physically received at the factory gate. [S13] [S31] In-transit stock is part of pipeline supply and cash exposure, never part of on-hand stock. [S09] [S31]
- **Quarantine (`Q_quar`)**: Material physically received at the dock or returned from the line that is blocked from production commitment pending count verification, incoming quality inspection, shade-band classification or technical disposition. [S07] [S31]
- **Available (`Q_avail`)**: Material that has passed incoming inspection, carries a verified unit quantity and shade lot, and is free to be allocated to a production order or customer promise. [S07]
- **Reserved (`Q_res`)**: Material that has passed inspection and is bound by an active, dated reservation to a specific confirmed customer order or released work order. [S07]
- **Rejected (`Q_rej`)**: Material that has failed specification or quality tolerances and is physically segregated pending supplier return, authorised rework or accounting write-down under IAS 2 net realisable value rules. [S09] [S27]
- **Return-inspection (`Q_ret`)**: Customer returns or post-packing rejects held in a segregated inspection zone until grading determines whether they can be restocked as first-quality goods, repaired, downgraded to seconds or lawfully disposed of. [S09] [S26]

RECONCILING BOOK STOCK, ACCEPTED USABLE STOCK AND NET REQUIREMENTS

Let total physical book stock on the premises for a given item-lot-location be B:

$$B = Q_{\text{avail}} + Q_{\text{res}} + Q_{\text{quar}} + Q_{\text{rej}} + Q_{\text{ret}}$$

Only two of those five physical buckets have passed technical acceptance for standard production. Define **accepted usable stock** (U_{acc}) as:

$$U_{\text{acc}} = B - \left(Q_{\text{quar}} + Q_{\text{rej}} + Q_{\text{ret}}\right) = Q_{\text{avail}} + Q_{\text{res}}$$

Define **unreserved usable stock** (U_{free}), which is the quantity available to support a new, unallocated promise without displacing an existing commitment, as:

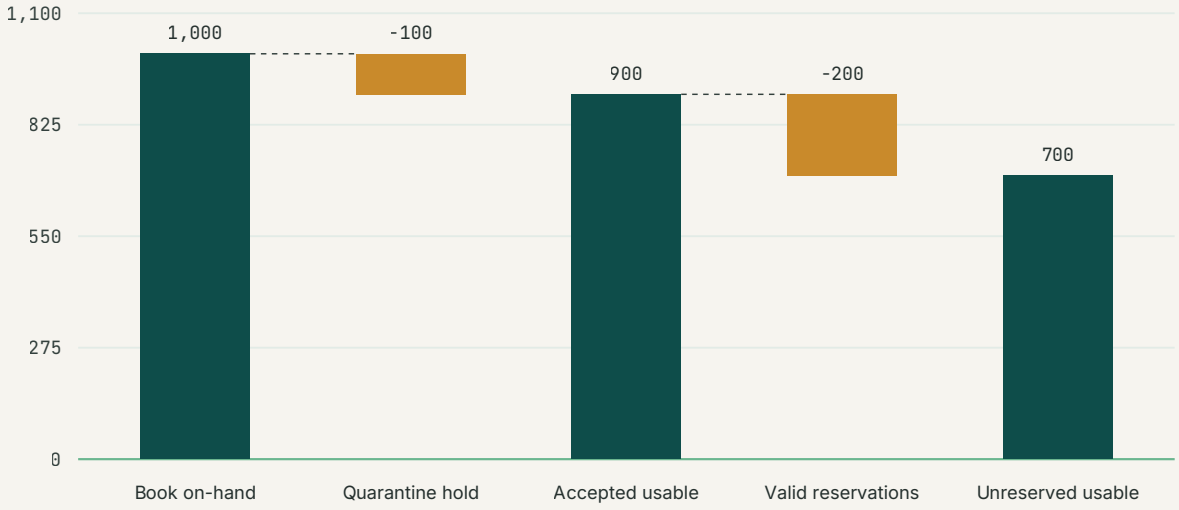
$$U_{\text{free}} = U_{\text{acc}} - Q_{\text{res}} = Q_{\text{avail}}$$

When a planner or material requirements planning (MRP) run calculates net requirements, the reservation boundary must obey a strict no-double-netting invariant. [S05] [S07] If the gross requirement vector G_{new} includes only **new, unallocated orders** (because previously reserved orders have already been removed from the demand schedule), then the net requirement is $\max(0, G_{\text{new}} - U_{\text{free}} - F_{\text{free}})$, where F_{free} is unallocated qualifying firm receipts. Conversely, if the gross requirement vector $G_{\text{all}} = G_{\text{new}} + Q_{\text{res}}$ includes **all open orders** (both reserved and unreserved), then the calculation must net against total accepted usable stock U_{acc}: $\max(0, G_{\text{all}} - U_{\text{acc}} - F_{\text{all}})$. Subtracting Q_{res} from U_{acc} to obtain U_{free} and then subtracting Q_{res} a second time inside the demand schedule would deduct the same reservation twice and inflate purchase orders by exactly Q_{res}. [S05] [S07]

Exhibit 7, p. 37 shows how an illustrative physical book balance of 1,000 pieces steps down to 700 unreserved usable pieces once quarantined holds and valid reservations are excluded. Every figure in the chart is an independent teaching assumption used to make the state transitions inspectable. [S07] [S15] [S31]

EXHIBIT 07 / ILLUSTRATIVE COMPONENT OR GARMENT PIECES

Book stock becomes promiseable stock only after exclusions



ILLUSTRATIVE [S07] [S15] [S31] Illustrative arithmetic waterfall. Starts with 1,000 physical pieces on hand in the book ledger, subtracts 100 quarantined pieces awaiting quality disposition to reach 900 accepted usable pieces, and subtracts 200 validly reserved pieces committed to active orders to leave 700 unreserved usable pieces available for new promises. Reservations must not be subtracted a second time if demand has already been netted. S07, S15 and S31 inform the separation of physical, inspected and committed states; all quantities are teaching assumptions.

The waterfall also highlights why reservations require an expiry and reconciliation rule. In many factories, planners place manual reservations on fabric or trims for tentative customer enquiries that are later cancelled or modified, leaving "ghost reservations" (`Q_res`) locked in the ledger. Those phantom reservations depress `U_free`, trigger unnecessary replenishment orders and block other feasible work orders from releasing. [S07] [S0203] GFT would attach a mandatory parent-order reference and an expiry timestamp to every reservation record: if a tentative customer quote is not converted into a confirmed order by its expiry date, or if a work order is closed short, the ledger automatically releases the residual reservation back to `Q_avail` and logs the state transition. [S07] [S27]

2.4

Receiving closes a chain of evidence

The moment external supply crosses the factory gate is where most ledger inaccuracies originate. When a busy receiving dock posts a goods receipt by simply clicking "receive in full" against the open purchase order screen, any supplier short-shipment, carton substitution, transit damage or unapproved shade split is silently absorbed into the factory's book stock. [S15] [S31] Months later, that unrecorded receiving discrepancy resurfaces as an unexplained stockout on the sewing line or an inventory write-off at year-end. [S09] [S27] Receiving must therefore operate as a four-way chain of evidence that links commercial authorisation, supplier dispatch, physical dock tally and technical quality release. [S07] [S31]

Evidence from retail and apparel supply chains confirms that order, shipment and receipt records diverge frequently even among established trading partners. In Project Zipper, a one-year study conducted from June 2017 to July 2018 by GS1 US and the Auburn University RFID Lab, researchers examined EPC/RFID and transaction data exchanged between eight brand owners and five retailers. [S31] While the numerical error percentages reported in that North American study are specific to its retail sample and must never be cited as guaranteed operating savings or a baseline for a Cameroon factory, the study establishes a vital structural principle: purchase orders, advance shipping notices (ASN), physical receipts and inventory records are distinct claims about reality that require explicit event-by-event reconciliation. [S31]

In GFT's proposed receiving ledger, every inbound delivery progresses through four linked event records, each carrying both an **event timestamp** (t_event, when the physical action occurred) and a **recorded timestamp** (t_rec, when the entry was posted to the system): [S30] [S31]

1. **Approved purchase order (PO)**: Records the authorised item code, specification and BOM version, ordered quantity, purchase unit, agreed price, Incoterms rule and named place, and promised delivery window. [S07] [S29]
2. **Supplier shipment notice (ASN or packing list)**: Records the quantities, roll or carton serial identifiers, dye-bath lots and carrier references that the supplier claims to have dispatched on t_ship. [S31] Until the physical goods cross the factory dock, this quantity sits in in-transit (Q_trans) rather than warehouse stock.
3. **Physical goods receipt (GRN)**: Recorded at the dock gate on t_dock via a blind or unit-verified physical count of cartons, rolls or cones before the operator sees the supplier's packing-list subtotals. [S15] [S27] The GRN moves the physically verified quantity from in-transit into quarantine (Q_quar), not directly into available stock. Any gap between the ASN quantity and the GRN tally immediately opens a **shipment discrepancy record** (short receipt, over-receipt, damaged container or wrong item label). [S27] [S31]
4. **Incoming quality and shade release (IQR)**: Recorded by the quality role on t_release after physical width and weight checks, shade-band continuity grouping and defect inspection. [R29] [S07] The IQR transitions accepted handling units from quarantine (Q_quar) to available (Q_avail) (or directly to reserved, Q_res, if pegged to a waiting order) and moves failed units to rejected (Q_rej) with a mandatory exception reason code. [S07] [S27]

Separating t_event from t_rec across these four steps serves two diagnostic purposes. First, inspired by the World Customs Organization's Time Release Study methodology (S30), which uses process mapping and segmented timestamps to isolate clearance bottlenecks and inspect outliers, recording both dock arrival (t_dock) and quality release (t_release) reveals how much replenishment time is consumed inside the factory's own receiving and inspection queue before material becomes usable. [S30] Second, comparing t_event with t_rec exposes delayed clerical posting: if rolls are physically unloaded on Monday morning (t_dock) but keyed into the ledger on Wednesday afternoon (t_rec), planners making decisions on Tuesday are looking at stale supply data. [S27] [S31]

Crucially, a short delivery or a rejected roll must also be preserved in the supplier performance and replenishment dataset. As research on inventory systems with uncertain supplies notes, both demand and supply data can be censored when only truncated or completed quantities are recorded. [S22] If a storekeeper overwrites the original ordered quantity of 1,000 metres down to the 850 metres actually delivered so that the purchase order can be administratively closed without a variance flag, future lead-time and supplier-fill calculations will see a 100 percent complete delivery of 850 metres instead of a 150-metre supply shortfall. [S22] [S31] The ledger must retain Ordered (`1,000`), Shipped (`900`), Received (`880`) and Accepted (`850`) as separate immutable fields on the receipt line. [S31]

2.5

Cycle counts diagnose recurring transaction errors

Even with disciplined receiving, physical stock and book records drift apart during daily operations. Operators issue an extra roll of interlining without logging a pick ticket, return a partial cone of thread to the wrong bin, misread a size label during bundle kitting or fail to record cutting-table scrap. [S15] [S27] Waiting for an annual financial stocktake to discover those errors leaves the replenishment engine flying blind for fifty-one weeks of the year. Cycle counting replaces the once-a-year shutdown with a continuous, stratified audit process whose primary purpose is not merely to adjust the ledger balance, but to diagnose and eliminate the transaction failures that caused the drift. [S01] [S27] [S0203]

As documented in the North Carolina Office of the State Controller's cycle-counting procedures (`S27`, `INV-44`), the cycle-counting process assesses stock levels within an inventory organisation, links count frequency to ABC classifications and enforces structured discrepancy correction. [S27] For an apparel manufacturer, however, stratifying count frequency by annual monetary usage (`ABC` value) alone leaves a dangerous operational gap: a custom woven size label, a specific Pantone sewing thread or a uniform collar zip may represent less than 1 percent of annual purchasing spend (`C` value) while possessing 100 percent line-stop criticality if its book balance is overstated. [S01] [S07] GFT therefore recommends a two-dimensional count-priority rule that assigns the highest cycle-count frequency to items that are **either** high in annual usage value (`Class A`) **or** flagged as single-source, long-lead-time or line-stopping components in the bill of materials. [S01] [S07] [S27]

To prevent a cycle count from becoming a rubber-stamp exercise where warehouse staff simply copy the computer balance onto the count sheet, the workflow must enforce strict role separation and blind counting, drawing on the procedural controls detailed in `S27`, `S0201` (`INV-15 Enter Cycle Count Results`) and `S0202` (`INV-28 Approve Cycle Counts`): [S27] [S0201] [S0202]

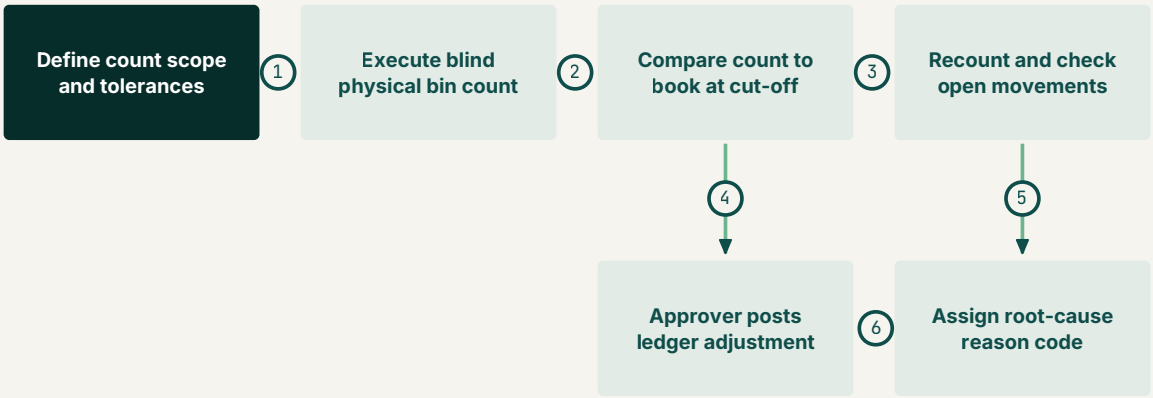
1. **Schedule and scope definition (`Warehouse Manager` role)**: Defines the count schedule, ABC assignment groups, item-location scope and acceptable quantity and value variance tolerances before generating count sequences. [S27]

2. **Blind physical count entry (`Warehouse Counter` role)**: The counter receives a count sequence listing the warehouse organisation, zone, bin, item code, lot and unit of measure, but **omitting** the current book quantity. The counter physically measures or counts the handling units in the bin, enters the count quantity and submits the sequence. [S0201]
3. **Automated variance check and recount trigger**: Once the count sequence is submitted, the system compares the blind physical count against the book balance at the count timestamp and notifies the Cycle Count Approver to approve, reject or request a recount. [S0201] If the discrepancy exceeds the item's quantity or value tolerance, an independent second counter performs a blind recount while verifying whether in-flight picks or unposted dock receipts occurred during the count window. [S27] [S0201]
4. **Variance review, reason coding and approval (`Cycle Count Approver` role)**: Before any ledger adjustment posts, the approver must review the cycle-count variance report, compare the quantities, assign a mandatory root-cause reason code (such as `unrecorded issue`, `misplaced bin location`, `unit-of-measure conversion error`, `unlogged scrap` or `receiving tally error`) and formally approve, reject or return the sequence for recount. [S0202]

Exhibit 8, p. 41 maps this segregated workflow. No single operator can both count a bin and approve an inventory write-off, and no discrepancy closes without a classified cause that feeds operational remediation. [S27] [S0201] [S0202]

EXHIBIT 08 / PROPOSED WORKFLOW STAGES

A discrepancy stays open until its cause and correction are approved



1. hide book qty
2. submit result
3. outside tolerance
4. within tolerance
5. variance confirmed
6. remediation logged

ILLUSTRATIVE [S27] [S0201] [S0202] Illustrative discrepancy-resolution workflow proposed by GFT, adapted from public-sector cycle-counting control procedures (S27, S0201, S0202). Separates count definition, blind physical entry, variance comparison, independent recount and transaction cut-off check, root-cause classification and approver sign-off before ledger adjustment.

Reporting cycle-count results also requires mathematical care. A single "99 percent net financial inventory accuracy" headline can conceal catastrophic operational errors if positive and negative bin discrepancies cancel each other out. Empirical studies of inventory record accuracy in fashion and apparel (S15, S16) emphasise evaluating physical count deviations at the granular record level rather than relying on netted monetary totals. [S15] [S16] GFT therefore requires every cycle-count report to publish three separate metrics across the N counted item-lot-bin records:

- **Exact-match (or tolerance-match) record rate**: The proportion of counted item-lot-bin records where $|Q_{\text{phys},i} - Q_{\text{book},i}| \le \tau_i$, where τ_i is zero for discrete piece/carton items and a stated measurement tolerance (such as 0.5 percent of roll length) for continuous roll goods. [S15] [S27]
- **Gross absolute unit discrepancy (GAUD)**: Defined for each unit-of-measure group as $\sum_{i=1}^{N} |Q_{\text{phys},i} - Q_{\text{book},i}| / \sum_{i=1}^{N} Q_{\text{book},i}$. Because the numerator uses absolute values, a $+50$ piece error in Size M and a -50 piece error in Size L adds to a 100-piece absolute error rather than netting to zero. [S15] [S16]

- **Gross and net financial value discrepancy**: Computed using unit cost c_i as $ \sum_{i=1}^{N} c_i |Q_{\text{phys},i} - Q_{\text{book},i}|$ (gross absolute value variance) alongside the signed net accounting adjustment $\sum_{i=1}^{N} c_i (Q_{\text{phys},i} - Q_{\text{book},i})$ required for IAS 2 financial reporting. [S09] [S27]

When a bin count reveals a missing roll (`-1` in Bin A) and a simultaneous surplus of the exact same roll serial ID in an adjacent rack (`+1` in Bin B), the approver should classify the event as a **location transfer error** rather than a stock loss, correct the bin location and retrain the put-away team. Treating location errors and physical losses separately prevents false shrinkage alarms while still fixing the bin records that pickers rely on. [S27] [S0202]

2.6

RFID is a candidate capture method that requires a local trial

When managers confront recurring scan omissions or slow cycle counts, vendors frequently propose replacing barcode labels with radio-frequency identification (`RFID`) as an automatic cure for inventory inaccuracy. A disciplined engineering evaluation must resist that shortcut. Both barcode scanning and ultra-high-frequency (`UHF`) passive RFID are physical signal-capture mechanisms subject to distinct error modes, environmental interference, unit economics and reconciliation overhead. [S15] [S16] [S31] Neither technology guarantees stock truth by itself, and published studies from retail stores or brand-to-retailer distribution centres answer different questions that must not be conflated into a factory purchase justification. [S15] [S16] [S31]

A printed 1D or 2D barcode on a fabric roll ticket, thread carton or garment bundle is an optical line-of-sight token. Its strengths are low unit media cost, unambiguous one-to-one targeting (the operator knows exactly which roll or carton the laser or imager just read), and immunity to electromagnetic reflection from metal racking or adjacent bins. Its operational weaknesses are human and mechanical: every roll or carton must be physically oriented so the label is visible, high-speed counting requires manual handling of each unit, thermal labels can scuff or peel in humid dye-house or cutting environments, and operators under time pressure may scan one carton label and manually type "$\times 20$" on the keypad rather than scanning all twenty cartons. [S15] [S27]

Passive UHF RFID eliminates line-of-sight presentation by reading tag antennas in bulk via radio waves, reducing the physical time needed to sweep a rack of tagged garments or pass a carton pallet through a portal. [S15] [S16] However, deploying UHF RFID inside a garment factory or fabric store introduces four engineering challenges that require empirical measurement:

1. **Read completeness (`false negatives`)**: Dense stacks of folded garments, high-moisture cotton bales, metallic buttons or zips, foil packaging and metal storage racks absorb, detune or reflect UHF signals. A wand sweep that moves too rapidly past a dense rack will miss tags buried in the centre of a stack, causing the cycle count to understate physical stock and create false negative variances. [S15] [S16]

2. **Stray and cross-zone reads (`false positives`)**: Because radio waves propagate through cardboard and across aisle boundaries, a handheld reader or dock portal tuned at high power can inadvertently read tags sitting in an adjacent bin, a quarantine cage two metres away or a forklift passing behind the operator. In a ledger that separates `quarantine`, `available` and `reserved` bins, stray reads corrupt location and status boundaries. [S15]
3. **Duplicate reads and tag-commissioning errors**: While EPC/RFID middleware deduplicates identical electronic product codes (`EPCs`) within a single scan window, printing a duplicate tag or attaching a replacement tag without decommissioning the damaged original causes one physical roll or garment to be counted twice. [S15] [S31]
4. **Reconciliation labour and unit tag cost**: Attaching and encoding an RFID tag on low-margin basic garments or disposable trim cartons adds recurring unit material and labour cost on every shift. Moreover, a fast RFID wand sweep does not eliminate the time required to investigate variances when the RF count disagrees with the book ledger. [S16] [S27]

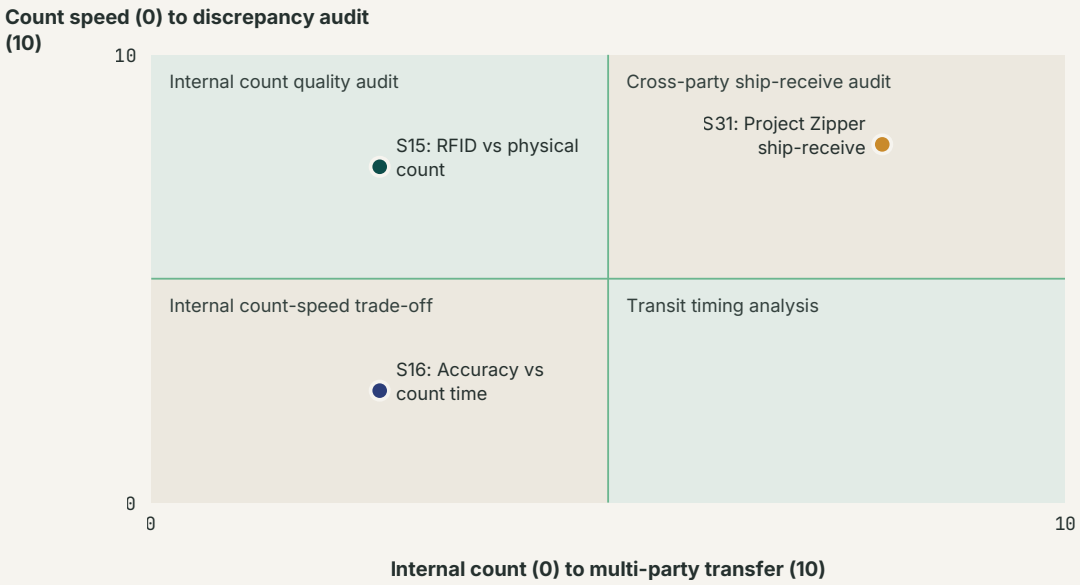
Three published studies in the paper's evidence register illustrate how carefully external RFID findings must be interpreted:

- **Cycle-count reliability against physical truth (`S15`, 2013)**: This fashion-retail study explicitly compares the quality of an RFID cycle count against a physically conducted count as its starting point. [S15] Its methodological lesson is that an RFID scan is a noisy measurement instrument whose read completeness and error rate must be benchmarked against a verified physical baseline, never assumed to be 100 percent accurate. [S15]
- **Accuracy as a function of count duration (`S16`, 2015)**: This apparel-retail field study empirically derives a trend of inventory accuracy as a function of the time required to conduct RFID inventory counts. [S16] It demonstrates an operational trade-off: shortening the time spent sweeping a zone reduces read completeness, so count speed and resulting inventory accuracy must be evaluated jointly rather than treated as independent benefits. [S16]
- **Inter-party shipment and receipt reconciliation (`S31`, 2018)**: The GS1 US and Auburn University RFID Lab Project Zipper study examined data obtained from eight brand owners and five retailers over one year to reconcile brand shipments against retail receipts and purchase orders. [S31] That study addresses multi-party supply-chain data exchange and order fulfilment discrepancies at the shipping/receiving boundary; its findings do not measure internal factory raw-material bin counting or fabric-roll length tracking. [S31]

Exhibit 9, `p. 44` places these three sourced studies on a qualitative two-dimensional scope matrix: the horizontal axis distinguishes single-site internal stock counting from multi-party shipment-to-receipt exchange, while the vertical axis distinguishes count-speed curve analysis from discrepancy-audit and reconciliation design. The coordinates (`0` to `10`) are qualitative layout positions that map each study's analytical boundary, not numerical accuracy percentages. [S15] [S16] [S31]

EXHIBIT 09 / QUALITATIVE STUDY SCOPE INDEX (0 TO 10)

RFID evidence answers different questions in each study



SOURCED [S15] [S16] [S31] Qualitative comparison of analytical scope across three opened sources: S15 (2013) compares RFID cycle-count quality to a physical count in fashion retail; S16 (2015) derives inventory accuracy as a function of count duration in apparel retail; S31 (2018) reconciles ship and receive data across eight brand owners and five retailers. Coordinates are qualitative layout indices mapping operational boundary (x) against primary evaluation focus (y), not measured accuracy scores.

For GFT's Cameroon textile programme, the practical implication of **Exhibit 9**, p. 44 is clear: barcode scanning on handling units (fabric rolls, thread cartons, cut-bundle tickets and finished-goods cartons) is the mandatory baseline capture method for Phase 1, while UHF RFID is a candidate option that requires a bounded local trial before any capital commitment. [S15] [S16] [S31] A local trial must test actual fabric roll density, warehouse racking geometry, operator sweep speed, false-positive stray reads across adjacent bins, tag failure rates in humid conditions and total cost per thousand reads against a controlled physical baseline. [S15] [S16]

2.7

Capture can continue during a network outage without corrupting stock

A stock ledger designed for an emerging-market factory must assume that electrical power or local wireless connectivity will occasionally drop in the middle of a working shift. If the warehouse management interface requires a live, synchronous round-trip to a central server before a storekeeper can receive a truck or issue fabric rolls to the cutting room, a two-hour network outage forces an unacceptable choice: either halt physical material movement and starve the sewing lines, or wave the rolls through on informal paper notes that get lost, misread or keyed in the wrong sequence hours later. [S27] [S31] [S0203] Conversely, poorly designed offline caching can corrupt the ledger even faster than paper notes if reconnected mobile terminals overwrite current bin balances or post duplicate deductions. [S31]

While WP07 owns site power and edge hardware selection and WP18 owns network and software transport architecture across the GFT white-paper library, WP12 specifies the inventory-specific data invariants required to keep the stock ledger consistent during and after an connectivity interruption: [S07] [S27] [S31]

1. **Locally unique, deterministic event identifiers**: Every scan or transaction captured on a mobile terminal while online or offline generates an immutable composite identifier (`event_id`) constructed from the registered device ID, a strictly monotonic local sequence counter, the operator role ID and the local hardware clock timestamp (`t_event`). [S31]
2. **Append-only movement events rather than state overwrites**: Mobile devices must never transmit absolute bin-balance overwrites (such as `SET Bin_A12 = 14 rolls`). Instead, every offline action is recorded in a local non-volatile append-only log as a signed movement delta tied to a specific handling-unit serial ID or lot (`ISSUE Roll_ID R-0492, Lot L-08, Qty 88.5 m from Bin_A12 to WorkOrder WO-204 at t_event`). [S07] [S31] Because signed movements of distinct handling units commute, two operators working in the same zone during an outage can synchronise their logs later without one terminal overwriting the other's deductions.
3. **Dual-timestamp reconciliation (`t_event` versus `t_sync`)**: When connectivity resumes and the terminal uploads its queued log, the ledger server records the ingestion timestamp (`t_sync`) alongside the original physical movement timestamp (`t_event`). [S30] [S31] Ordering the historical ledger by `t_event` reconstructs the true physical sequence of dock receipts, quality releases and line issues, while the gap `t_sync - t_event` measures offline lag.
4. **Idempotent ingestion and duplicate suppression**: Network drops frequently occur halfway through an upload packet, causing the mobile terminal to retry the transmission upon reconnect. The ledger server enforces idempotency by indexing every committed transaction by `event_id`: if a retried batch contains an `event_id` already present in the committed log, the server acknowledges receipt and suppresses the duplicate without deducting stock a second time. [S31]
5. **Adjustment and promise guards while synchronisation is pending**: The most common source of post-outage ledger corruption occurs when a supervisor approves a cycle-count adjustment or manually overrides a bin balance on the central server while a handheld scanner on the dock or cutting floor still holds unsynchronised picks for that same bin. [S27] [S0202] Under GFT's control rule, each warehouse zone tracks the last heartbeat and sequence number of its assigned mobile terminals. If a terminal has unuploaded offline transactions (`sync_pending = true`), any item-bin in that zone is flagged as `reconciliation-pending`: cycle-count variance approvals (`S0202`) are held in queue until the terminal synchronises and the cut-off comparison (`t_count` versus `t_event`) is re-evaluated, and unconfirmed offline picks are deducted conservatively before new available-to-promise commitments are approved. [S07] [S27] [S0202]

GFT NOTE

Never allow an operator or automated script to post an unexplained "miscellaneous stock adjustment" to force a book balance to match a physical bin during or after a connectivity outage. Require every offline terminal in the affected warehouse zone to complete synchronisation, deduplicate by `event_id`, sort movements by physical `t_event` and re-run the cut-off check before the Cycle Count Approver reviews any residual discrepancy. [S27] [S31] [S0202] If an unallocated balance remains uncertain because a terminal cannot be recovered before an urgent production release, verify the specific roll or carton serial IDs physically at the bin and record a provisional reservation against those verified serials rather than overriding the book ledger. [S07] [S15]

BOUNDARY CONDITIONS AND UNRESOLVED LOCAL INPUTS

The ledger controls, unit-conversion equations and offline synchronisation rules in this chapter are engineering specifications, not claims of deployed software or measured error reduction in a GFT facility. Before an emerging-market factory can rely on automated available-to-promise or safety-stock calculations, it must establish baseline exact-match and gross absolute discrepancy rates across its actual fabric, trim and finished-goods bins, validate roll-length and width tolerances against its real suppliers, and test barcode or candidate RFID capture under local warehouse conditions. [S15] [S16] [S27] Where IAS 2 inventory valuation, customs bond records or contractual consignment terms apply, the ledger's reason codes and write-down approvals require review by qualified accounting and legal professionals in the relevant jurisdiction. [S09] [S29]

This stock-truth contract forms the foundation for the rest of WP12 and its interfaces across the GFT research library and tool suite. For WP05 (traceability and provenance), the handling-unit serial ID, dye-bath lot and append-only `event_id` log provide the internal chain of custody linking incoming fibre or fabric lots to finished garment packs, feeding the `trace-a-garment` tool without asserting that end-to-end field traceability is already qualified. [S07] [S31] For WP14 (data and model governance) and WP18 (digital architecture), the separation of `t_event` and `t_sync`, idempotent event ingestion and role-segregated discrepancy approval define the data-hygiene gates that any replenishment algorithm must check before generating an order. [S27] [S31] Most directly, the `safety-stock` tool and the segmentation and replenishment rules in Chapter 3 through Chapter 8 consume **unreserved usable stock** (`U_free`) and **accepted usable stock** (`U_acc`) rather than raw book balances. [S01] [S05] [S07] Once

the factory can prove what stock it actually holds, in what unit, in what quality state and under whose reservation, it is ready to segment commercial promises and choose differentiated stock policies in Chapter 3. [S01] [S05]

A promise rests on what the line can use today, not on what the ledger forgot to deduct yesterday.

CHAPTER SOURCE NOTES

S01 MIT OpenCourseWare. 2006.

S05 MIT OpenCourseWare. 2006.

S07 MIT OpenCourseWare. 2006.

S27 North Carolina Office of the State Controller. 2023-10-24.

S31 GS1 US and Auburn University RFID Lab. 2018.

S0203 MIT OpenCourseWare. 2006.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S09 IFRS Foundation. n.d..

S20 OECD. 2018-03-07.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S15 Association for Information Systems Electronic Library. 2013.

R01 Business in Cameroon. Publication date unrecorded; inherited source-text check.

R02 Fibre2Fashion; Business in Cameroon. Publication date unrecorded; inherited source-text check.

S26 European Environment Agency. 2024.

S13 World Trade Organization. n.d..

S30 World Customs Organization. 2025.

S29 International Chamber of Commerce. 2020.

S22 arXiv, Chen, Jiang, Zhang and Zhou. 2023-10-31.

S0201 North Carolina Office of the State Controller. 2023-07-26.

S0202 North Carolina Office of the State Controller. 2023-07-27.

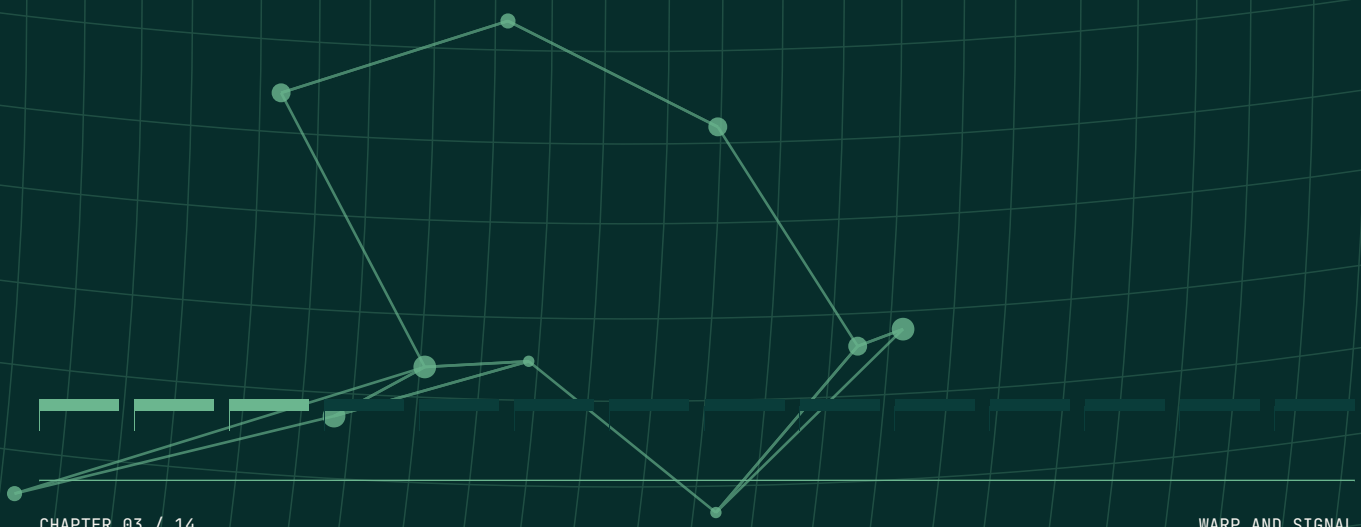
S16 University of Modena and Reggio Emilia institutional repository. 2015.

03

PART I: DEFINE THE STOCK DECISION BEFORE BUYING THE STOCK / SEGMENTATION

Stock policy must follow commercial purpose and supply criticality

A high-value fashion line, a low-cost line-stopping trim and a repeat uniform fabric need different decision rules.



CHAPTER 03
DECISION TEST

A high-value fashion line, a low-cost line-stopping trim and a repeat uniform fabric need different decision rules.

An apparel factory should never apply a single inventory rule across its item register, nor segment materials solely by annual purchase spend. [S01] In a garment plant, a low-cost polyester thread or approved brass zip can halt a sewing line just as completely as a missing roll of shirting fabric, while a high-value seasonal print can lose its recovery value within weeks of a missed selling window. [S06] [S07] Green Future Tech Limited (GFT), Hong Kong SAR, argues that stock segmentation must combine four dimensions before any reorder formula is chosen: annual usage value, demand pattern and intermittency, the decoupling point of the customer promise, and supply criticality. [S01] [S05] Within GFT's Cameroon textile programme, prepared in January 2024 with source-refresh notes retained in 2026, this discipline links physical stock truth to replenishment mathematics so working capital protects confirmed delivery commitments. [R06] [R07]

When an emerging-market manufacturer moves from raw-material export context toward domestic value addition, inventory mistakes compound quickly. Inherited sector context indicates that about 95 percent of Cameroon's cotton fibre is exported raw while about 5 percent is processed locally, alongside a national target of 50 percent local processing by 2030. [R06] [R07] Separate reporting on CICAM records market share declining from about 85 percent three decades ago to about 5 percent, with July 2023 capacity utilisation reported at 33.7 percent at Garoua, 48.9 percent at CICAM 1 and 10.83 percent at CICAM 2, and underutilisation cost near CFAF 2 billion in 2020. [R11] [R12] Those sector figures are background context rather than GFT operating results, yet they show why a garment programme cannot afford material starvation on active lines or capital locked in unsaleable cloth. [R11] [R12] Segmentation prevents a planner from buying a fashion print like a repeat institutional twill, or ignoring a two-franc care label until three hundred shirts sit quarantined at packing. [S01] [S07]

3.1

Value classification is useful but insufficient

Classical inventory segmentation begins with annual usage value, termed ABC analysis, which ranks stock-keeping units by multiplying physical consumption over a defined horizon by a consistent unit cost. [S01] If item i has annual consumption D_i and unit cost c_i in CFA francs (XAF), its annual usage value is $V_i = D_i \times c_i$. [S01] Sorting the register by descending V_i and computing cumulative spend share lets planners focus executive review, supplier negotiation and audit frequency on the small fraction of items driving most cash outflow. [S01] [S27]

3.2 USAGE VALUE REQUIRES A CONSISTENT BOUNDARY AND COST BASIS

Even this basic calculation fails when the observation window, unit of measure or valuation basis is inconsistent. [S01] [S09] First, the planner must state whether D_i reflects historical issues, net requirements exploded from firm orders, or an unconstrained forecast. [S07] [S22] Where historical issues were truncated by stock-outs, observed usage D_i understates true demand; an item out of stock for four months of the prior year can look like a low-value C item precisely because purchasing failed to supply it. [S22] [S0301] Second, unit cost c_i must follow a common boundary across imported and domestic materials. [S09] Valuing imported dyed twill at landed factory-gate cost while recording local grey cloth at ex-mill price

distorts capital intensity. [S09] Whilst International Accounting Standard IAS 2 measures financial-statement inventories at the lower of cost and net realisable value, operational segmentation requires replacement cost in XAF, with foreign-currency quotes converted on a declared reference basis. [R32] [S09]

Third, the cumulative breakpoints separating Class A, Class B and Class C items are local policy choices, not universal constants. [S01] An illustrative convention might place the top 70 to 80 percent of annual usage value into Class A, the next 15 to 20 percent into Class B, and the remaining 5 to 10 percent across most item codes into Class C. [S01] In procedural guidance such as the North Carolina Office of the State Controller cycle-counting framework, ABC classes govern how frequently bins are counted and investigated, focusing verification on high-spend items without exhausting store labour on minor consumables. [S27]

3.3 LINE-STOP CRITICALITY CORRECTS THE BLIND SPOT OF SPEND RANKING

In dependent-demand manufacturing, commercial spend and operational criticality diverge sharply. [S07] [S0302] A garment is an assembly of mutually dependent components governed by a versioned bill of materials. [S07] In a cut-make-trim (CMT) or full-package factory, one hundred percent of fabric value cannot become an accepted invoice if a single mandatory trim, label, button or thread is missing at assembly. [S07] Empirical research in complex operations shows why single-criterion ABC ranking is inadequate alone: in a study of 247,000 spare-part records over nine years at a seaport, researchers combined ABC spend analysis with Analytic Hierarchy Process (AHP) weighting to isolate 1,416 critical items requiring strict inventory control. [S0302] Although a seaport store differs from an apparel trim room, the flaw of single-axis spend ranking is identical: low annual spend does not mean low stockout consequence. [S01] [S0302]

To prevent inexpensive items from starving sewing lines, GFT overlays a three-level operational criticality classification onto the spend ranking, adapted from Vital-Essential-Desirable (VED) and bill-of-material completeness principles: [S01] [S07]

- **Criticality 1 (Line-stop or mandatory specification):** Any component without a pre-approved, available substitute whose absence prevents cutting, sewing, finishing or legal shipment of a confirmed order, or any safety item needed to run the line safely. [S07] [S20]
- **Criticality 2 (Constrained substitute or rework-inducing):** Materials where a shortage forces a line changeover, sub-assembly buffering, or customer concession request, causing setup loss or extra handling without halting the plant. [S07]
- **Criticality 3 (Non-blocking or general consumable):** Standard packaging ancillaries or multi-supplier consumables that have approved local substitutes and do not block primary garment assembly. [S01]

As shown in **Exhibit 10**, p. 51, combining annual usage value (Classes A, B and C) with demand pattern and line-stop criticality produces an actionable priority score rather than a flat spend list. [S01] [S27]

EXHIBIT 10 / ILLUSTRATIVE GOVERNANCE PRIORITY SCORE (1 TO 5)

A policy matrix separates value, variability and criticality

	Stable repeat	Seasonal spike	Lumpy tender	Intermittent
A: High spend (crit. 1)	4	5	5	5
A: High spend (crit. 3)	3	4	4	4
B: Medium spend (crit. 1)	3	4	4	5
C: Low spend (crit. 1 trim)	4	4	5	4
C: Low spend (crit. 3)	1	2	2	2

Cell scale: 1 to 5

ILLUSTRATIVE [S01] [S27] [S0302] Illustrative GFT policy scoring grid (1 = routine automated/periodic rule, 5 = executive cash and line-stop gate). Breakpoints between A, B and C value classes and demand variability tiers are configurable local policy choices based on S01, S27 and S0302, not universal constants.

Exhibit 10, p. 51 shows why a low-spend Class C item carrying Criticality 1 status (such as an imported tunic zip or a customer-nominated woven badge) receives a governance priority score of 4 or 5, whereas a Class C general consumable with Criticality 3 receives a score of 1 or 2. [S01] [S07] What **Exhibit 10**, p. 51 cannot establish on its own is the exact CFA franc threshold between Class A and Class C in a specific Cameroon factory; those cut-offs must be calibrated from the plant's verified material ledger and working-capital ceiling. [S01] [S09]

3.4

Demand variability and intermittency require separate descriptions

Once annual spend and criticality are recorded, the second segmentation axis examines the temporal structure of demand. [S01] [S04] Two garment inputs can consume the same annual quantity of 1,200 units over twelve months while imposing completely different risk profiles on purchasing and stores. [S04] [S05]

3.5 FOUR TEMPORAL PROFILES GOVERN APPAREL AND MATERIAL CONSUMPTION

In an emerging-market apparel operation serving retail basics, seasonal collections and institutional tenders, GFT separates demand into four structural patterns: [S01] [S04] [S05]

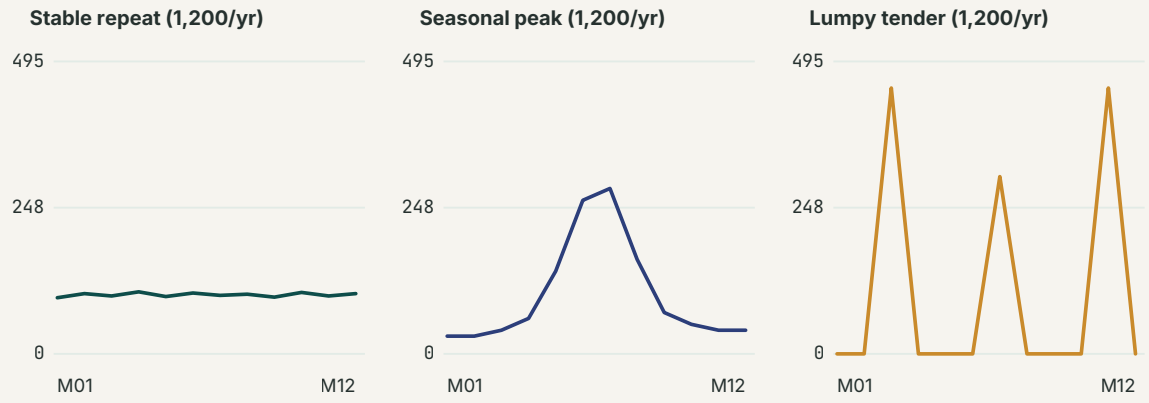
1. **Stable repeat consumption:** Continuous, stationary or gently trending demand across almost every review period, typical of core white or navy school-shirt poplin, standard sewing thread and basic polybags. [S05]

2. **Seasonal or fixed-window spikes:** Demand concentrated within a predictable calendar window, such as back-to-school uniforms, festive printed collections or campaign workwear, where consumption rises sharply and then drops to a low baseline. [S04] [S06]
3. **Lumpy institutional tenders:** Demand arriving in large, discrete batches separated by multiple zero-demand intervals, characteristic of ministry orders, corporate workwear contracts or export batch commitments. [S04] [S07]
4. **Genuinely intermittent demand:** Sporadic, low-volume requests separated by long and irregular zero-demand intervals, typical of specialty spare parts for sewing and cutting machines, uncommon size extremes (such as 4XL or 5XL garments), or bespoke repair trims. [S0301] [S0302]

Exhibit 11, p. 52 compares three illustrative twelve-month demand series that each sum to 1,200 units per year (mean 100 units per month) but exhibit stable, seasonal and lumpy-intermittent behaviour. [S04] [S05]

EXHIBIT 11 / ILLUSTRATIVE UNITS DEMANDED PER MONTH

Equal annual demand can produce different ordering risks



ILLUSTRATIVE [S04] [S05] [S0301] Illustrative constructed sequences over a 12-month horizon (M01 to M12). Each panel totals exactly 1,200 units per year (mean 100 units per month). Methodological contrast follows time-varying and probabilistic demand principles in S04, S05 and zero-demand classification in S0301.

As **Exhibit 11**, p. 52 shows, an annual mean of 100 units per month is an operational fiction for the second and third panels. [S04] [S05] In the first panel, monthly demand stays between 95 and 105 units; a continuous-review reorder point calibrated to stationary variance protects service reliably. [S05] In the second panel, a static buffer calibrated to a 100-unit average causes stockouts during months M05 to M08 (when demand reaches 260 to 280 units) and leaves idle capital during months M10 to M03. [S04] [S06] In the third panel, eight months record zero demand while three tender releases require 300 to 450 units each; a stationary normal safety-stock formula holds dead stock during zero intervals and still fails when a 450-unit order arrives. [S04] [S07] Because **Exhibit 11**, p. 52 uses constructed sequences, it illustrates the failure of single-mean aggregation rather than reporting observed Cameroon factory orders. [S04] [S05]

3.6 ZERO-DEMAND INTERVALS AND UNSTABLE VARIATION RATIOS REQUIRE METHODOLOGICAL GUARDS

Planners often try to automate demand classification using the coefficient of variation ($\text{CV} = \sigma / \mu$), where σ is the sample standard deviation and μ is the sample mean of period demand. [S05] While CV suits positive, stationary series, it becomes unstable when a series contains frequent zeroes, seasonal shifts or censored observations. [S22] [S0301]

METHOD

To classify demand patterns without distorting intermittent or seasonal items, GFT specifies a three-step diagnostic boundary in coordination with WP04: [S04] [S05] [S0301]

1. **Separate natural zeroes from artificial zeroes first:** Cross-check zero-consumption periods against the usable-stock ledger from Chapter 2. As shown by Svetunkov and Sroginis, recorded sales or issue zeroes occur either naturally or artificially because of stockouts and recording errors. [S22] [S0301] Treating a four-week stockout as four weeks of zero demand corrupts both the estimated mean and the inter-demand interval. [S22] [S0301]
2. **Decompose occurrence frequency from non-zero transaction size:** For series with genuine zero-demand periods, compute the average inter-demand interval (ADI, total periods divided by positive-demand periods) separately from the squared coefficient of variation of positive demand sizes ($\text{CV}_+^2 = (\sigma_+ / \mu_+)^2$). [S0301] Never divide overall standard deviation by a near-zero mean $\mu \approx 0$; when μ approaches zero, small rounding changes send the naive ratio σ / μ toward infinity without providing stable classification information. [S05] [S0301]
3. **Test for calendar seasonality and dependent tender schedules before smoothing:** Where demand spikes correspond to school terms, crop-buying seasons or scheduled contract call-offs, classify the item as time-varying or dependent-demand rather than random noise. [S04] [S07]

This separation routes items with high inter-demand intervals and large batch sizes to dependent material requirements planning (MRP) or order-triggered purchasing rather than a normal safety-stock calculator. [S05] [S07]

3.7

The customer promise determines the stock object

Inventory segmentation cannot stop at consumption statistics, because the physical form in which stock is held depends on what the factory promised the customer and how much time the customer allows between order confirmation and delivery. [S01] [S07] In supply-chain engineering, this boundary is the customer order decoupling point: the stage where forecast-driven push inventory ends and confirmed order-driven pull execution begins. [S01] [S17]

3.8 FOUR COMMERCIAL ARCHETYPES DEFINE THE DECOUPLING POINT

For an apparel manufacturer in Cameroon or comparable emerging markets, GFT distinguishes four commercial archetypes, each requiring a different stock object at the decoupling point: [S01] [S06] [S07] [S08]

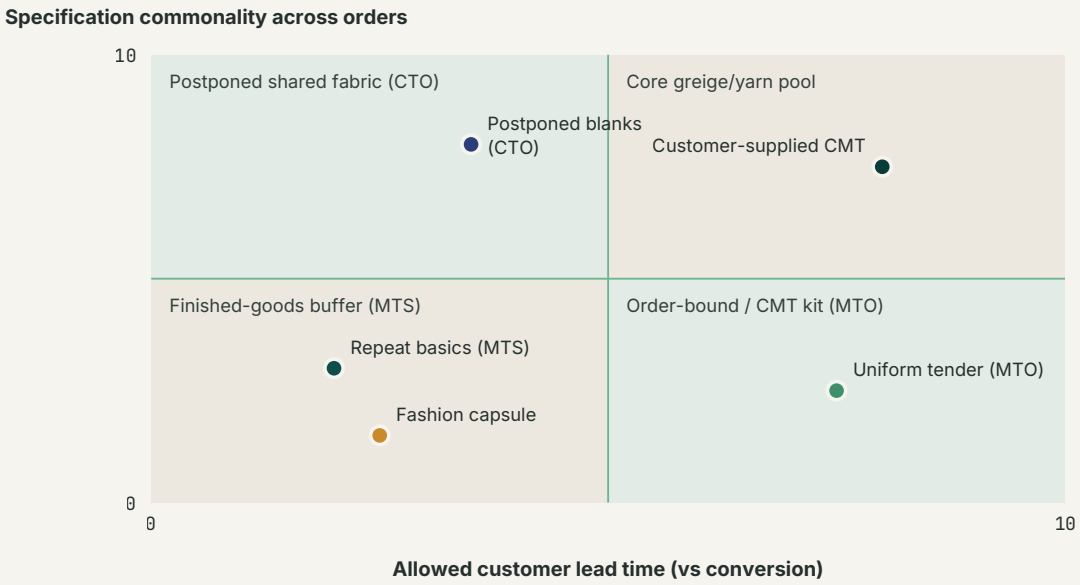
- **Make-to-stock (MTS) repeat basics:** Standardised local retail items such as plain T-shirts, work trousers or core school shirts where buyers expect immediate shelf dispatch. [S01] [S05] The decoupling point sits at finished goods by style, colour and size, backed by core fabrics. [S01] [S05]
- **Postponed or configure-to-order (CTO) apparel:** Garments sharing a common undyed fabric, prepared-for-dyeing (PFD) blank or neutral base garment, where final dyeing, printing, embroidery or badging is delayed until demand is observed. [S07] [S17] The decoupling point sits at shared intermediate fabric or blanks plus generic trims. [S17]
- **Make-to-order (MTO) institutional and uniform tenders:** Confirmed contracts for ministries, utilities, schools or corporate buyers with an agreed delivery window and custom specification. [S07] [S08] No speculative finished-goods buffer is held; commitment is triggered by contract confirmation and deposit, or held as reserved raw fabric only where contract lead time is shorter than mill lead time. [S07] [S08]
- **Customer-supplied-material (CSM / pure CMT) assembly:** Contracts where the buyer retains ownership of primary fabric and supplies it to the factory for cutting, sewing and packing. [S07] [S09] The factory holds custodial quarantine and traceability records for customer-owned cloth while purchasing only its own thread, needles and ancillaries. [S07] [S09]

Empirical evidence from the apparel sector confirms that moving the decoupling point works only when governance matches the commercial channel. [S17] In an empirical study of Japanese apparel firms in the *Quarterly Journal of Marketing*, researchers distinguished postponed manufacturing from postponed product planning and found that while manufacturing postponement was associated with lower inventory, planning postponement could increase inventory when retail control remained externalised. [S17] Although a Japanese wholesaler survey is not a Cameroon factory forecast, its operational mechanism applies directly: holding undyed or uncut fabric reduces stock only if the factory controls the conversion schedule and finishing capacity needed to meet the customer's delivery window. [S07] [S17]

Exhibit 12, p. 55 maps these four archetypes across allowed customer fulfilment lead time relative to conversion time and specification commonality across customers. [S01] [S06] [S07] [S08]

EXHIBIT 12 / ILLUSTRATIVE POSITIONING SCALE (0 TO 10)

Commercial archetypes map to different stock positions



ILLUSTRATIVE [S01] [S06] [S07] [S08] Illustrative GFT positioning of four commercial apparel archetypes on customer lead-time tolerance (x-axis, 0 = immediate shelf dispatch, 10 = full procurement and production horizon) and material commonality across orders (y-axis, 0 = bespoke single-buyer spec, 10 = fully interchangeable core spec). Conceptual basis: S01, S06, S07 and S08.

In **Exhibit 12**, p. 55, points in the lower-left quadrant (repeat basics and fashion capsules) face short customer lead times and low specification substitutability at sale, forcing stock commitment into finished goods or pre-dyed cloth. [S01] [S06] Points in the upper-left quadrant (postponed blanks and shared fabric) use material commonality to hold stock upstream of final printing or trimming, provided conversion fits inside the promise window. [S07] [S17] Points on the right side of **Exhibit 12**, p. 55 (uniform tenders and customer-supplied CMT) have customer lead times long enough to execute procurement and production against firm orders, making speculative finished stock unnecessary. [S07] [S08] Because the coordinates in **Exhibit 12**, p. 55 are illustrative scores on a 0 to 10 scale, they define policy selection rather than Cameroon market shares. [S01] [S08]

3.9 EVERY CUSTOMER PROMISE REQUIRES FIVE EXPLICIT PARAMETERS

To make the decoupling point operational, GFT requires every commercial segment charter to specify five contractual parameters before any buffer is calculated: [S05] [S07] [S08]

- **Stock object and specification boundary:** Whether the promise applies to a finished garment SKU (style-colour-size), an approved semi-finished blank, or a reserved fabric roll within an accepted shade-lot tolerance. [S07]
- **Delivery location and Incoterms boundary:** Whether fulfilment is measured at the factory packing gate, a domestic hub in Douala or Yaounde, or a customer warehouse. [S01] [S07]
- **Confirmed response window:** Maximum elapsed working days between order acceptance (including deposit clearance) and usable dispatch. [S07] [S08]

- **Minimum and maximum order quantity:** The order band within which the lead-time promise holds without triggering a custom production schedule. [S05] [S08]
- **Substitution rule:** Whether the contract forbids any material or size substitution (standard in institutional tenders) or permits pre-agreed colour/size substitution in local wholesale packs. [S01] [S08]

3.10

Service measures must name the event they count

A recurring error in factory planning is invoking a "95 percent service level" without stating the numerator, denominator or counting unit. [S05] In probabilistic inventory theory, different service metrics count distinct physical events and lead to different safety-stock investments and customer outcomes. [S05]

3.11 CYCLE SERVICE LEVEL, UNIT FILL RATE AND ORDER OTIF MEASURE DISTINCT OUTCOMES

Foundational inventory theory distinguishes event-based cycle protection from volume-based demand fulfilment, while operational governance adds order-level delivery: [S05]

1. **Cycle service level (CSL, or P_1):** The probability, or observed fraction of replenishment cycles, in which usable stock on hand meets all lead-time demand without a stockout event. [S05] In a continuous-review model with order quantity Q, CSL counts whether a cycle experiences zero shortage (1 if zero, 0 if any shortage occurs), whether the shortage was 1 garment or 500 garments. [S05]
2. **Unit fill rate (FR, or P_2):** The fraction of requested unit demand satisfied immediately from usable stock without backordering or lost sales. [S05] Under standard continuous-review assumptions where expected shortage per cycle is $E[\text{shortage per cycle}]$ and order quantity is Q, unit fill rate is $\text{FR} = 1 - (E[\text{shortage per cycle}] / Q)$. [S05] Because Q sits in the denominator, changing lot size Q alters the unit fill rate even when the reorder point and cycle service level remain unchanged. [S05]
3. **Order on-time in full (OTIF):** Under GFT's proposed metric contract (aligned with WP02), the fraction of confirmed customer orders delivered with one hundred percent of the agreed complete quantity across all lines and sizes on or before the confirmed delivery date, passing quality acceptance. [S05] [S07]

Exhibit 13, p. 57 contrasts these service definitions, showing their mathematical structure, counting unit and primary failure mode when misapplied. [S05]

EXHIBIT 13 / METRIC DEFINITION AND BOUNDARY COMPARISON

Service definitions cannot be substituted for one another

Metric name	Numerator and denominator	Counting unit	Operational blind spot if used alone
Cycle service level (CSL / P1) [S05]	Stockout-free replenishment cycles / total replenishment cycles	Replenishment cycle	Ignores shortage depth and lot-size frequency; a massive shortage counts as one failed cycle
Unit fill rate (FR / P2) [S05]	Units filled immediately from usable stock / total units requested	Individual physical unit	Can look high in aggregate while missing one critical trim or size that blocks an entire order
Kit readiness rate (GFT / S07)	Production orders with 100% of BOM lines usable by release date / total scheduled orders	Complete production kit	Measures internal line-release readiness rather than final customer dispatch punctuality
Order OTIF (GFT proposed contract)	Confirmed customer orders delivered 100% complete on or before agreed date / total confirmed orders	Complete customer order	Penalises partial shipments strictly; requires line-level diagnostics to locate root cause

DERIVED [S05] [S07] Derived comparison combining standard probabilistic inventory definitions of cycle service level and unit fill rate from S05 with dependent bill-of-material completeness from S07 and GFT's proposed order-level OTIF metric contract. No universal numeric target is prescribed.

The operational lesson of **Exhibit 13**, p. 57 is that high item-level fill rates can coexist with poor order-level OTIF or frequent line stoppages. [S05] [S07] Consider an illustrative uniform order requiring $n = 10$ independent bill-of-material components, each with a 95 percent probability of being available on the required date. [S05] [S07] Under independence, the probability that all 10 components are simultaneously ready for complete-kit release is $0.95^{10} \approx 0.5987$, or just under 60 percent. [S05] [S07] With 20 independent components at 95 percent availability each, the illustrative complete-kit probability falls to $0.95^{20} \approx 0.3585$, or about 36 percent. [S05] [S07] While component stockouts in a real factory are often correlated, this arithmetic explains why an apparel plant cannot assign a uniform 90 or 95 percent availability target to cheap Criticality 1 trims and expect to ship complete garments on time. [S05] [S07]

GFT NOTE

Never accept a software default such as "set 95 percent service across all SKUs." Require planners to state whether the target is cycle service (CSL), unit fill rate (FR), complete-kit readiness or customer order OTIF. For low-cost Criticality 1 trims (thread, care labels, approved zips, buttons) feeding dependent assembly, set high availability protection or hold order-matched kits so cheap ancillaries never govern the completion probability of expensive fabric and sewing labour. [S01] [S05] [S07]

3.12

Seasonal sell-through needs a different risk appetite

Whereas repeat basics and core trims follow multi-cycle replenishment, fashion capsules, festive prints and event merchandise operate inside a finite selling window where unsold stock cannot wait for next month's demand. [S04] [S06] Applying a continuous-review reorder point or a high fill-rate target to a short-season style creates terminal overstock that erodes earlier sales margin. [S06] [S09]

3.13 THE BALANCE BETWEEN UNDERAGE AND OVERAGE GOVERNS SEASONAL COMMITMENT

In single-period inventory theory (the newsvendor model), the planner chooses an order quantity before random seasonal demand D is realised, balancing two asymmetric unit costs: [S06]

- **Unit underage cost (C_u):** Contribution lost when a unit of demand cannot be met ($C_u = p - c$, selling price less incremental cost, plus any contractual penalty). [S06]
- **Unit overage cost (C_o):** Loss on each unit ordered beyond realised demand that remains unsold at season end ($C_o = c - s$, unit cost less net salvage value s). [S06]

Under a continuous demand distribution $F(x)$ and linear costs, the profit-maximising cycle service probability (the critical fractile) is $\text{CF} = C_u / (C_u + C_o)$. [S06] Although Chapter 7 develops single-period ordering and salvage exits in detail, the critical fractile belongs in Chapter 3 as a segmentation rule because it proves that optimal service depends on salvage economics. [S06] For an illustrative seasonal garment selling at $p = \text{XAF } 12,000$, costing $c = \text{XAF } 7,000$, and clearing at net salvage $s = \text{XAF } 2,000$, both C_u and C_o equal $\text{XAF } 5,000$, giving $\text{CF} = 0.50$ (50 percent). [S06] Placing that print into a repeat-stock segment with a 95 percent cycle service target means every unit buffered above median demand has over a 50 percent chance of losing XAF 5,000 on clearance. [S06]

3.14 RESIDUAL SELLING TIME SEPARATES REORDER CANDIDATES FROM CLEARANCE CANDIDATES

An item's policy class also shifts as the calendar advances through a selling season. [S04] [S06] At the start of a twelve-week season, if domestic conversion takes two weeks, strong early sell-through allows a second replenishment run. [S04] [S06] Once the remaining selling window is shorter than usable-receipt and conversion lead time (Chapter 4), replenishment is impossible; the item shifts from replenishment to markdown recovery. [S04] [S06] Under IAS 2 principles, once seasonal or obsolete stock clears below cost, its carrying value must be written down to net realisable value, turning a segmentation error into a financial charge. [S09]

3.15

Cash and physical risk can override an ABC label

Even after accounting for spend, demand pattern, decoupling point and salvage, a statistical classification can recommend a commitment that violates storage, supply continuity or cash constraints. [S07] [S20] [S25] Four physical and commercial risks regularly override standard ABC-XYZ rules: [S07] [S20] [S25]

1. **Single-source imported bottlenecks and subcontractor opacity:** A modest-spend material may come from a single foreign mill or dye house with maritime and customs exposure, or through subcontracting where upstream capacity is shared across buyers. [S20] [S25] OECD due diligence guidance notes that subcontracting is common across the garment sector and requires mapping upstream risks rather than assuming supplier independence. [S20] [S25]
2. **Non-interchangeable shade lots and roll widths:** Dyed fabrics exhibit roll-to-roll shade and shrinkage variation. [S07] Fifteen remnant rolls from six dye lots totaling 300 metres cannot be cut together into a 250-metre uniform order requiring shade continuity. [S07] Segment rules must define usable lot minimums rather than aggregate metre totals. [S07]
3. **Physical shelf-life, humidity and storage degradation:** Elastics, adhesive interlinings, fusible stays, inks and chemical finishes degrade under heat and humidity when stored for long periods, capping the maximum coverage horizon. [S06] [S09]
4. **Fragile buyer commitments and concentration risk:** High-value fabric bought for a single buyer with unverified credit or call-off history carries counterparty default risk that a generic fabric pool does not. [S08] [S09]

3.16 GOVERNANCE REQUIRES ESCALATION RATHER THAN SILENT TARGET RELAXATION

When a segment's target service level requires more working capital or space than the budget permits, an unguided team will often make informal cuts to order quantities or delay supplier payments. [S08] [S20] Both responses are harmful: arbitrary cuts destroy bill-of-material completeness on Criticality 1 items, while late payment or sudden order changes shift stress onto suppliers and workers. [S07] [S20]

> **WARNING**
>
> If the calculated buffer or order quantity for a segment breaches the approved CFA franc cash ceiling, storage capacity or shelf-life limit, the planner must not silently lower the service parameter or issue an incomplete partial kit. GFT governance requires a formal exception escalation presenting leadership with three choices: (a) approve additional working-capital allocation, (b) renegotiate the customer promise (longer lead time, deposit or narrower SKU range), or (c) shift the decoupling point upstream to shared fabric or make-to-order execution. [S01] [S05] [S08]

3.17

A policy matrix must prescribe an action

A segmentation scheme has no operational value if it ends as a spreadsheet leaving planners using the same reorder rule for every row. [S01] In GFT's inventory architecture, every segment code binds six parameters: [S01] [S05] [S27]

- **Review cadence:** Continuous transaction-triggered review, weekly review, monthly review, or order-milestone gate review. [S05] [S27]
- **Demand input source:** Stationary residual distribution from WP04, seasonal critical-fractile distribution, or net dependent demand from confirmed bills of materials (MRP). [S05] [S06] [S07]
- **Replenishment policy family:** Continuous reorder-point (s, Q) or (s, S), periodic review (R, S) or (R, s, S), single-period newsvendor buy with stop-date gate, or lot-for-lot order-matched kit purchasing. [S04] [S05] [S06] [S07]
- **Service metric and protection bound:** Explicit choice of CSL, unit FR or kit readiness, with minimum and maximum coverage bounds. [S05]
- **Cycle-count and audit cadence:** Physical verification frequency linked to spend class and criticality (monthly blind counts for Class A and Criticality 1 trims, quarterly for Class B, semi-annual for Class C non-critical items). [S27]
- **Exception owner and substitution authority:** The named role (planner, procurement lead, quality lead, finance manager or general manager) authorised to approve a lot override, substitution or express freight. [S01] [S27]

WORKED EXAMPLE: SEGMENTING FOUR ILLUSTRATIVE CAMEROON APPAREL ITEMS

Consider an illustrative Cameroon garment factory evaluating four items over a 12-month horizon. All quantities, XAF costs and demand parameters below are illustrative teaching inputs constructed to show the arithmetic, not observations from an operating GFT facility. [S01] [S05] [S06] [S07]

1. **Item 1 (Core navy poly-cotton twill fabric, 150 cm width):** Used across repeat workwear and uniform lines. Illustrative annual usage $D_1 = 30,000\text{ metres}$ at landed cost $c_1 = \text{XAF } 2,200\text{ per metre}$ gives $V_1 = \text{XAF } 66,000,000$. Demand occurs in all 12 months ($\text{ADI} = 1.0$, monthly mean $\mu_1 = 2,500\text{ metres}$, $\sigma_1 = 300\text{ metres}$, $\text{CV}_1 = 0.12$). Absence halts core cutting (Criticality 1).
2. **Item 2 (Heavy-duty 18 cm brass trouser zip, navy):** Mandatory trim for work trousers made from Item 1. Illustrative usage $D_2 = 24,000\text{ pieces}$ at $c_2 = \text{XAF } 150\text{ per piece}$ gives $V_2 = \text{XAF } 3,600,000$ ($\text{ADI} = 1.0$, $\text{CV}_2 = 0.12$, no local substitute, Criticality 1).
3. **Item 3 (Seasonal festive printed dress fabric):** Ordered for an eight-week year-end capsule. Illustrative usage $D_3 = 8,000\text{ metres}$ at $c_3 = \text{XAF } 3,500\text{ per metre}$ gives $V_3 = \text{XAF } 28,000,000$. Salvage value is $s_3 = \text{XAF } 1,400\text{ per metre}$, with underage contribution $C_{u,3} = \text{XAF } 2,100\text{ per metre}$ and overage loss $C_{o,3} = \text{XAF } 2,100\text{ per metre}$.
4. **Item 4 (Bespoke embroidered ministry tunic badge):** Used in two tender releases with illustrative usage $D_4 = 6,000\text{ pieces}$ at $c_4 = \text{XAF } 400\text{ per piece}$, giving $V_4 = \text{XAF } 2,400,000$ ($\text{ADI} = 6.0$, 60-day contract window versus 25-day procurement lead time, Criticality 1 for that tender).

Step 1: Compare naive spend ranking with multi-criteria segmentation. Total illustrative spend is $\text{XAF } 100,000,000$: Item 1 is 66.0 percent (Class A), Item 3 is 28.0 percent (Class A/B), Item 2 is 3.6 percent (Class C) and Item 4 is 2.4 percent (Class C). [S01] Spend-only rules would over-buffer Item 3 and under-protect Item 2. [S01]

Step 2: Apply GFT's four-dimensional policy assignment.

- **Item 1 (Class A, Stable repeat, Shared fabric, Criticality 1):** Continuous-review reorder point (s, Q) calibrated to usable-receipt lead time (Chapter 5), monthly cycle counts and high unit fill-rate target. [S01] [S05] [S27]
- **Item 2 (Class C spend, Stable repeat, Criticality 1 override):** Every missing 150-franc zip blocks trousers using XAF 3,300 of twill ($1.5\text{ metres} \times \text{XAF } 2,200$). Criticality 1 overrides Class C spend: Item 2 gets a high-protection continuous-review or two-bin buffer with monthly counts. [S01] [S05] [S07]
- **Item 3 (Class A/B spend, Seasonal window, Low salvage):** Barred from continuous replenishment; its illustrative critical fractile is $\text{CF}_3 = 2,100 / 4,200 = 0.50$ (50 percent), governed by a single-period buy, stop-order date and clearance rule (Chapter 7). [S06] [S09]

- **Item 4 (Class C spend, Lumpy tender $\text{ADI} = 6.0$, MTO, Criticality 1):** Order-triggered lot-for-lot MRP buy upon tender confirmation and deposit (Chapter 8), with zero speculative buffer between tenders. [S04] [S07] [S08]

3.18 SEGMENT ASSIGNMENTS REQUIRE SCHEDULED REVIEW AND VERSION CONTROL

Markets, lead times and contracts evolve over a garment programme. [S01] [S04] A bespoke uniform fabric for one buyer may later be adopted across four regional clients, shifting it from order-bound MTO purchasing to a pooled fabric buffer. [S07] [S17] Conversely, a formerly stable basic may see rising zero-demand intervals as preferences shift, requiring reclassification from continuous replenishment to make-to-order or clearance before excess stock builds. [S06] [S0301]

GFT recommends reviewing segment classifications quarterly and whenever a major contract, bill-of-material revision or route change occurs. [S01] [S27] Every reclassification is logged in the item-location policy ledger from Chapter 1, preserving the prior code, effective date, trigger and approving role. [S01] [S27] With usable stock reconciled (Chapter 2) and every item assigned to a commercial promise and criticality segment (Chapter 3), the next step is measuring the full replenishment clock and its tail delays (Chapter 4). [S05]

A cheap trim that stops a sewing line is never a minor item, and an expensive seasonal print that cannot be salvaged is never a candidate for routine safety stock.

CHAPTER SOURCE NOTES

S01 MIT OpenCourseWare. 2006.

S06 MIT OpenCourseWare. 2006.

S07 MIT OpenCourseWare. 2006.

S05 MIT OpenCourseWare. 2006.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R11 Business in Cameroon; MMI News. Publication date unrecorded; inherited source-text check.

R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

S27 North Carolina Office of the State Controller. 2023-10-24.

S09 IFRS Foundation. n.d..

S22 arXiv, Chen, Jiang, Zhang and Zhou. 2023-10-31.

S0301 arXiv, Svetunkov and Sroginis. 2025-11-13.

R32 Wise; XE. Publication date unrecorded; inherited source-text check.

S0302 arXiv, Arani, Abdolmaleki, Maleki, Momenitabar and Liu. 2021-05-28.

S20 OECD. 2018-03-07.

S04 MIT OpenCourseWare. 2006.

S17 Japan Marketing Academy, Quarterly Journal of Marketing. 2020-03-31.

S08 MIT OpenCourseWare. 2006.

S25 OECD. 2026-02-26.

04

PART II: MATCH REPLENISHMENT RULES TO MATERIAL AND DEMAND / THE SUPPLY CLOCK

Measure the full replenishment clock and its delays before choosing a buffer

The relevant lead time ends when approved material is usable, not when a supplier announces that it has shipped.

CHAPTER 04
DECISION TEST

The relevant lead time ends when approved material is usable, not when a supplier announces that it has shipped.

A garment factory cannot sew a shipping notice, a bill of lading or a customs pre-declaration. [S07] The replenishment clock that governs inventory policy starts when an internal requirement is raised for approval and stops only when an inspected, shade-verified and unencumbered lot is released to usable stock at the cutting table or trim store. [S30] Treating a vendor's quoted manufacturing window or ocean transit time as the complete lead time omits the administrative, port, border, inland and quality-release intervals where delays concentrate. [S12] Worse, reducing those intervals to a single arithmetic average conceals the upper tail of the distribution, the unfinished open orders still waiting in transit, and the tendency for delays to lengthen when seasonal demand or supplier utilisation peaks. [S22] Green Future Tech Limited (GFT), Hong Kong SAR, therefore requires any apparel operation under its Cameroon textile programme design to map and timestamp the ten-stage order-to-release clock, retain right-censored and split shipments in the empirical record, and evaluate whether compressing procedural variance is cheaper than financing a larger safety buffer. [S30]

This chapter establishes the measurement contract linking physical logistics and customs release to the safety-stock and lot-sizing rules developed later in this paper. [S05] While national industrialisation ambitions frame the context, where source-annex records indicate that about 95 percent of Cameroon's cotton fibre is exported raw against about 5 percent processed locally and policy targets reach 50 percent local processing by 2030, a planner cannot derive a trim or fabric buffer from macro targets. [R06] [R07] Each material route, whether importing dyed shirting and zippers through a coastal port or moving yarn and greige cloth along a domestic corridor from northern ginning centres, carries its own staged time distribution, contractual risk allocation and data-capture requirements. [S14] [S29] The discussion below is an operational and statistical engineering method, not investment, legal, customs-brokerage or tax advice. [S29]

4.1

The usable-receipt boundary includes approval and inspection

A replenishment policy protects the interval between recognising a material need and placing usable stock into production. [S05] If the measured lead time is truncated at the supplier's factory gate or at vessel arrival, the unmeasured downstream days do not vanish; they consume the safety buffer before the material ever reaches an operator. [S12] In apparel manufacturing, where a missing care label, sewing thread cone or approved metal zip can halt an entire cutting lay, the replenishment boundary must coincide with the physical and documentary conditions required for line release. [S07] [S31]

4.2 TEN SEQUENTIAL STAGES SEPARATE AN INTERNAL REQUISITION FROM USABLE STOCK

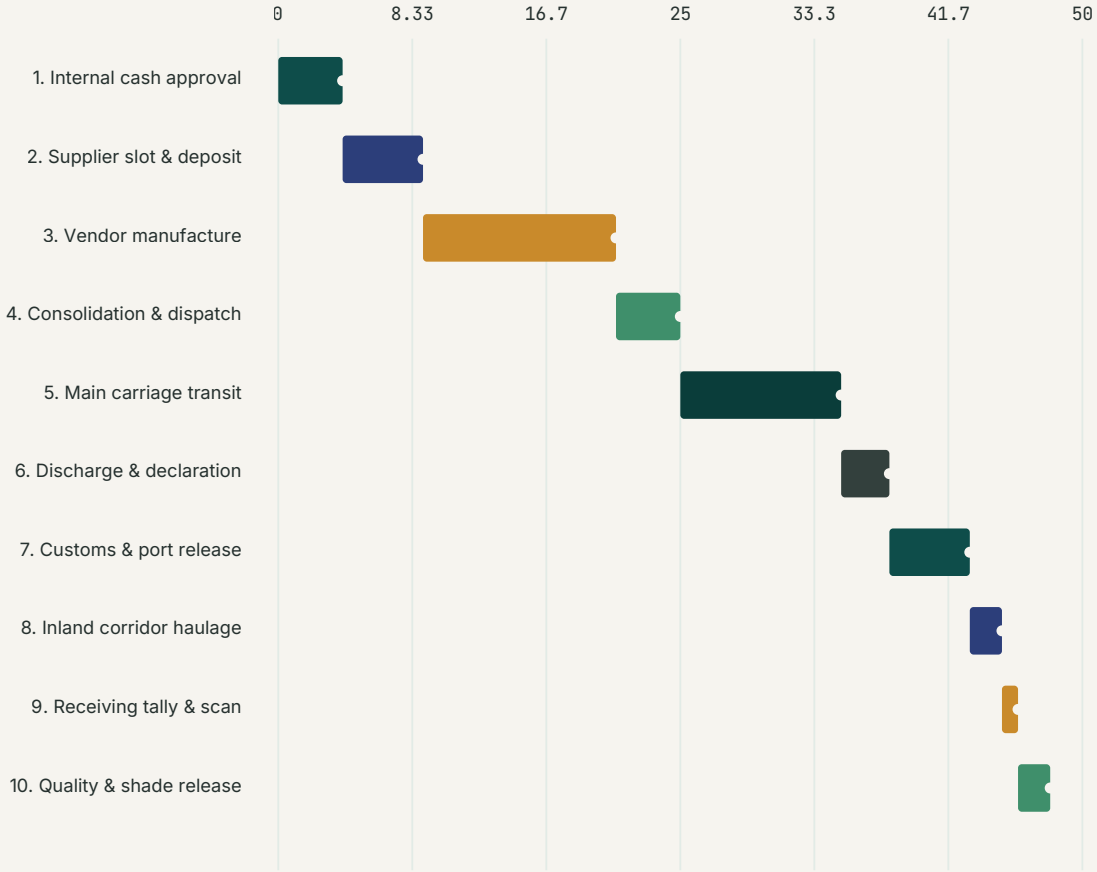
To prevent planners and buyers from using incompatible definitions, GFT decomposes the replenishment clock into ten non-overlapping intervals bounded by eleven auditable timestamps: [S07] [S30]

1. **Internal requisition to purchase approval (T0 to T1):** The interval from the planner's validated net-requirement signal to authorised commercial and cash commitment. [S07] When working capital is tight, requisitions can sit in an internal approval or deposit queue for several illustrative days before the supplier receives a binding order. [S18]
2. **Purchase approval to supplier queue entry (T1 to T2):** Time required to transmit the purchase order, satisfy advance-payment conditions, confirm lab-dip shade references and secure a slot on the supplier's schedule. [S07] [S08]
3. **Supplier queue to manufacturing completion (T2 to T3):** Physical conversion, dyeing, finishing or trim assembly at the vendor's facility, including any subcontracted stages mapped for due-diligence visibility. [S25]
4. **Manufacturing completion to carrier dispatch (T3 to T4):** Packing, export documentation, container consolidation and handover to the origin carrier or forwarder. [S12] [S29]
5. **Origin dispatch to port or border arrival (T4 to T5):** Main-carriage maritime, air or overland transit to the port of entry or land frontier, tracked by carrier event logs. [S12]
6. **Port or border arrival to customs and agency lodgement (T5 to T6):** Vessel berthing, container discharge, manifest registration and formal submission of import declarations and conformity certificates. [S13] [S30]
7. **Lodgement to physical border and terminal release (T6 to T7):** Selectivity routing, documentary or physical examination, duty settlement and issuance of the customs release note and terminal gate-out pass. [S30] [S0401]
8. **Terminal release to factory gate arrival (T7 to T8):** Inland drayage or corridor haulage from the seaport or northern depot, including weighbridges and urban delivery windows. [S12] [S13]
9. **Factory gate arrival to physical receiving count (T8 to T9):** Unloading, carton and roll tally, barcode or candidate radio-frequency identification (RFID) capture, and reconciliation against the shipping notice and purchase order. [S15] [S31]
10. **Physical receiving to quality release into usable stock (T9 to T10):** Quarantine hold while inspectors verify fabric width, weight, shrinkage, defect scores and shade-lot continuity, concluding with the system posting that moves accepted material to unreserved usable stock. [S07] [S27]

The structure of this ten-stage clock is shown in **Exhibit 14**, p. 66, illustrating how a shipment that looks rapid on a vendor's production tracker can take nearly twice as long to reach the cutting table. [S30]

EXHIBIT 14 / ILLUSTRATIVE CALENDAR DAYS

The usable-material clock extends beyond shipment



ILLUSTRATIVE [S12] [S13] [S30] [S31] Illustrative ten-stage order-to-release timeline totaling 48 calendar days on an assumed import replenishment order. No local Cameroon duration is claimed; interval structure adapts WCO Time Release Study boundary mapping (S30) and supply-chain event reconciliation (S12, S13, S31).

As **Exhibit 14**, p. 66 illustrates on a constructed 48-day teaching timeline, vendor manufacturing and consolidation (T2 to T4, days 9 to 25) plus main carriage (T4 to T5, days 25 to 35) account for 26 calendar days out of 48 illustrative days. [S12] The remaining 22 illustrative days occur before vendor production begins (9 days of internal approval and deposit clearance) and after port arrival (13 days across discharge, customs release, inland haulage, receiving and quality inspection). [S30] While **Exhibit 14**, p. 66 cannot establish actual transit or dwell durations for any specific Cameroon shipment, it shows why calibrating a reorder point from a 16-day production quote plus a 10-day maritime quote leaves 22 illustrative days of demand unprotected. [S05]

4.3 CALENDAR DAYS AND WORKING DAYS MUST MATCH THE DEMAND DENOMINATOR

A frequent unit error in inventory spreadsheets is multiplying a lead time measured in calendar days by a daily consumption rate calculated over factory working days. [S05] Vessels steam seven days a week and port storage runs across weekends, whereas a garment factory in Cameroon operates a standard 40-hour week across five or six working days (up to the 48-hour maximum in the shared evidence register). [R16]

[S12] If a line consumes an illustrative 600 metres of pocketing fabric per working day across a six-day week (3,600 metres per seven-day calendar week, or 514.3 metres per calendar day) and the replenishment clock spans an illustrative 35 calendar days (30 working days), expected demand over the lead time is 18,000 metres. [S05] Multiplying 600 metres per working day by 35 calendar days overstates expected demand at 21,000 metres, inflating the trigger by an illustrative 3,000 metres. [S05]

4.4 REQUESTED, PROMISED AND REALISED DATES SERVE DISTINCT AUDIT FUNCTIONS

For every stage from `T0` to `T10`, the ledger must store three date columns alongside the system-capture timestamp: [S30] [S31]

- **Requested date:** The target date requested by the buyer on the purchase order. [S07]
- **Promised date:** The initial contractual date committed by the supplier or carrier at order acceptance, with later revisions stored as versioned records rather than overwriting the baseline. [S08]
- **Realised physical date (`T_phys`):** The actual time when the physical or regulatory transition occurred, recorded separately from the **system-posted date (`T_post`)** when a clerk or offline scanner uploaded the entry. [S27] [S31]

Separating `T_phys` from `T_post` prevents administrative batching from corrupting lead-time variance. [S31] If interlining is released from quarantine on Friday afternoon (`T10_phys`) but posted on Monday morning (`T10_post`) after a connectivity outage, using `T10_post` adds three spurious days to the sample. [S31] Conversely, posting a goods receipt from a packing list while the truck is still queuing outside the gate understates the clock and creates phantom usable stock. [S15] [S27]

> **GFT NOTE**
>
> Never overwrite the original supplier promise date when a vendor sends a delay notice. Overwriting the baseline promise makes a late supplier appear 100 percent punctual against its final revised schedule, destroying the lead-time-variance history needed to size safety buffers. [S05] [S30] Store the original promise, every revision timestamp and the realised usable-receipt date as immutable records. [S31]

4.5

Historical corridor evidence is a diagnostic prompt

When designing a supply chain in Cameroon, analysts are sometimes tempted to lift port dwell times from older macro-economic reports and plug them into an inventory formula. [S14] GFT's evidence rules forbid this shortcut: historical port diagnostics and international indices explain *why* border and corridor stages must be measured separately, but they cannot supply current numeric parameters for a live buffer. [S12] [S14]

4.6 DATED DOUALA DIAGNOSTICS JUSTIFY STAGE GRANULARITY RATHER THAN CURRENT PARAMETERS

In February 2015, the World Bank examined the efficiency of the Port of Douala and noted that vessel turnaround and dwell constraints were tied to channel and terminal conditions: "This was mainly due to a lack of investment and the absence of routine dredging." [S14] For a supply-chain architect, that 2015 diagnosis teaches a structural lesson: maritime arrival (T5), berth discharge and declaration (T6), and customs and terminal release (T7) are separate operational bottlenecks that can dominate the replenishment clock if left unmeasured. [S14] [S30]

Citing a 2015 report as if it described container release times in 2026 would, however, be invalid. [S14] Over the intervening decade, Cameroon's maritime and corridor infrastructure has changed through the commissioning of the deep-water Port of Kribi, updated terminal concessions and revised customs procedures. [S12] [S14] Neither the 2015 Douala study nor macro cotton-sector reports on Camtext SA or SODECOTON provide a verified 2026 distribution of container release times for apparel inputs at Douala or Kribi. [R03] [R09] [S14]

4.7 THE REVISED LOGISTICS PERFORMANCE INDEX AND WCO GUIDE DEFINE THE MEASUREMENT PROTOCOL

A parallel caution applies to global logistics benchmarks. [S12] In its *Logistics Performance Indicators: LPI Report 2025*, the World Bank highlights stage interdependence: "Supply chains are only as good as their weakest link." [S12] The revised LPI methodology combines survey data with shipment-level tracking datasets covering containers, air freight and parcels, measuring corridor turnaround from operational timestamps rather than perceptions alone. [S12]

Meanwhile, the World Trade Organization's Trade Facilitation Agreement (TFA), which entered into force on 22 February 2017, states that "The TFA contains provisions for expediting the movement, release and clearance of goods, including goods in transit." [S13] [S0401] Article 7.6 of the TFA encourages members to measure release times using the World Customs Organization (WCO) Time Release Study (TRS) method. [S30] [S0401] Version 4 of the WCO *Time Release Study Guide* (2025) refines process mapping, interval segmentation and outlier analysis across border agencies, noting that "Improvements have been made to questionnaire design and sampling techniques to ensure data accuracy." [S30]

The distinction between what these institutional sources establish and what a factory planner must measure locally is summarised in **Exhibit 15**, p. 69. [S12] [S14] [S30]

EXHIBIT 15 / SOURCE SCOPE

Historical port evidence cannot calibrate a current buffer

Source ID	Publisher and date	Core scope and method	Permitted use in GFT design	Explicit limitation
S14	World Bank (2015-02-12)	Port of Douala efficiency diagnosis	Justifies separating arrival, discharge and release stages	Historical 2015 study; never use as a current parameter
S13 / S0401	WTO (2017-02-22 / n.d.)	Trade Facilitation Agreement Article 7	Defines pre-arrival processing and guarantee release levers	Legal treaty text; does not prove local clearance speed
S12	World Bank (2025)	LPI shipment-level operational tracking	Establishes weakest-link corridor timestamp tracking	Macro indicators; no SKU-level Cameroon distribution
S30	WCO (2025)	Time Release Study Guide Version 4	Provides protocol for interval sampling and outlier audit	Methodology only; requires local execution on shipments

SOURCED [S12] [S13] [S14] [S30] [S0401] Comparison of opened institutional sources on trade facilitation, port efficiency and release-time measurement. None of these sources supplies a current factory-specific lead-time distribution for Cameroon apparel inputs.

As **Exhibit 15**, p. 69 shows, institutional reports provide diagnostic structure and audit methods while leaving local durations to be measured. [S12] [S30] Before a Cameroon garment operation compares sourcing through Douala, Kribi or northern corridors from Garoua, GFT's evidence gate requires dated carrier and clearing quotes, mapped regulatory approvals for the relevant tariff headings, and timestamped shipment samples. [R03] [R10] [S14] [S29] [S30] Because shared annex references R19 to R22 are missing from the register, this chapter makes no claim regarding regional tariff bands or trade-preference dates. [S13]

4.8

Incomplete shipments must remain visible in the sample

Even when a factory records accurate timestamps from T0 to T10, computing mean lead time across "completed orders in the last 90 days" introduces a systematic downward bias by ignoring right censoring and partial fulfilment. [S22] [S0402]

4.9 RIGHT-CENSORED OPEN ORDERS CONTAIN THE LONGEST DELAYS IN THE SYSTEM

At any review cutoff date, recent orders that moved quickly have already passed T10 and appear in the completed receipts table. [S30] Orders placed weeks earlier that are stuck at a transshipment port, held in a customs query at T6, or waiting in quarantine at T9 for shade resolution have not yet reached T10. [S22] [S30] As Chen, Jiang, Zhang

and Zhou (2023) note, "A key challenge in this learning problem is that both demand and supply data can be censored; hence only truncated values are observable." [S22] Filtering for completed orders drops every in-flight delay from the calculation. [S22]

To eliminate survivorship bias without inventing arrival dates for open cargoes, GFT uses the non-parametric Kaplan-Meier (product-limit) estimator documented by the National Institute of Standards and Technology (NIST): "No underlying model (such as Weibull or lognormal) is assumed; K-M estimation is an empirical (non-parametric) procedure." [S0402] An order still open at the cutoff date with elapsed duration `c_j` is recorded as right-censored at `c_j`: its eventual usable lead time exceeds `c_j`, so it remains in the risk set `n_i` for all observed completion times `t_i` up to `c_j`. [S0402]

METHOD

Non-parametric Kaplan-Meier quantile estimation for right-censored lead times

1. **Define the cohort:** Select all purchase-order lines for a homogeneous route-material group issued within a defined window (such as an illustrative 180-day window), excluding only orders cancelled before supplier confirmation (`T2`) for documented customer changes. [S30]
2. **Assign cutoff status:** For each order `j` in the cohort of size `N`, record completed elapsed days `t_j = T10_j - T0_j` (`delta_j = 1`) or right-censored elapsed days `c_j = T_cutoff - T0_j` if open or only partially released (`delta_j = 0`). [S22] [S0402]
3. **Compute product-limit survival:** Across sorted completion durations `t_(1)` through `t_(k)`, let `d_i` be the count of orders completing at `t_(i)` and `n_i` be the count of all orders (completed or censored) with elapsed time at least `t_(i)`. [S0402] Estimate unfulfilled survival as `S_hat(t) = prod((1 - d_i / n_i))` for `t_(i)` up to `t`, and cumulative completion as `F_hat(t) = 1 - S_hat(t)`. [S0402]
4. **Extract planning quantiles:** For target probability `p` (such as `0.50`, `0.80` or `0.90`), select the smallest `t_(i)` where `F_hat(t_(i))` is at least `p`; if censoring keeps `F_hat(t)` below `p`, report the censored lower bound and apply a stress scenario. [S05] [S0402]

4.10 PARTIAL DELIVERIES AND SPLIT LOTS REQUIRE KIT-COMPLETE TRACKING

Supply censoring also occurs when an order arrives in split shipments or when quality control releases 70 percent of a fabric lot at `T10` while holding 30 percent in quarantine. [S22] [S31] Stamping the order with the *first* partial receipt date misleads the planner when an illustrative 10,000-shirt uniform order needs all 15,000 metres of fabric before the full contract can be fulfilled. [S07] GFT tracks three distinct metrics on split orders: **first-usable-lot lead time** (for progressive daily cutting), **quantity-weighted lead time** (for supplier flow monitoring), and **order-complete (kit-complete) lead time** (elapsed days until 100 percent of the ordered usable quantity and dependent trims pass `T10`). [S07] [S30] [S31]

WORKED EXAMPLE

How omitting open orders and partial lots distorts the 80th-percentile lead time

Consider an illustrative review of `N = 10` purchase orders for imported pocketing fabric, using constructed teaching figures to demonstrate the Kaplan-Meier method. [S05] [S0402] Seven orders completed full usable release (`T10`) in `22, 24, 25, 26, 28, 30` and `35` calendar days. Three orders remain incomplete at the cutoff date: [S22]

- Order 8 delivered 50 percent on day 27, with the remaining 50 percent delayed at a transshipment port and still open at 40 illustrative days (`delta_8 = 0, c_8 = 40`). [S22]
- Order 9 is held in customs query at `T6` at 42 illustrative days (`delta_9 = 0, c_9 = 42`). [S30]
- Order 10 failed initial shade checks at `T9` and awaits replacement rolls at 45 illustrative days (`delta_10 = 0, c_10 = 45`). [S27]

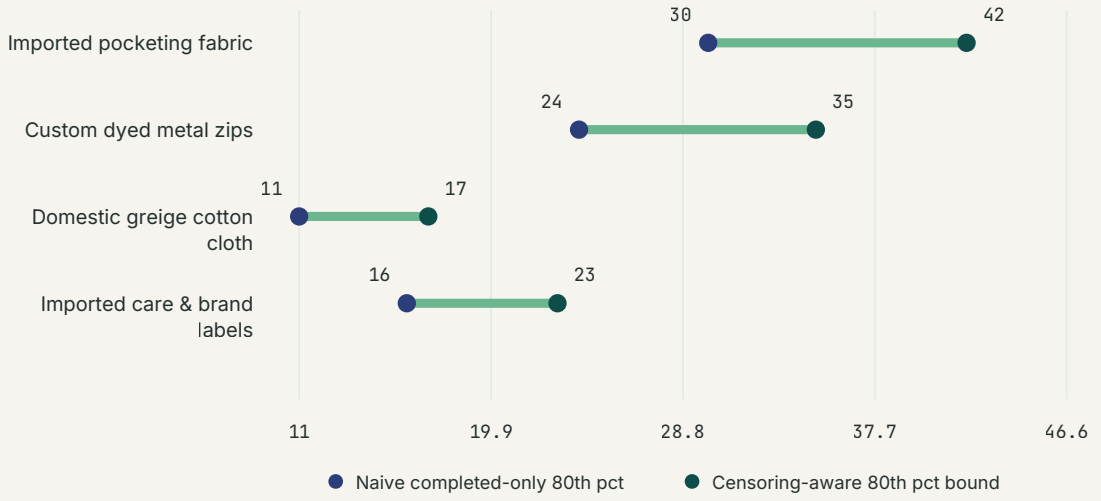
Method A (Naive completed-only plus first-receipt stamp): Counting the 7 completed orders plus Order 8 at its first partial receipt (day 27) yields 8 durations (`22, 24, 25, 26, 27, 28, 30, 35`) with a naive mean of `27.1` days, median of `26.5` days, and empirical 80th percentile (`7 / 8 = 0.875`) of **30.0 calendar days**. [S05]

Method B (Kaplan-Meier on order-complete usable receipts): Retaining all `N = 10` orders with Orders 8, 9 and 10 right-censored at 40, 42 and 45 days, all 7 completed durations (22 to 35 days) occur while all 3 censored orders remain in the risk set. [S0402] Each completion reduces `S_hat(t)` by `0.10`: the Kaplan-Meier median (`F_hat = 0.50`) is **28.0 days**, and at day 35 `F_hat(35)` reaches only `0.70` (`S_hat(35) = 0.30`). [S0402] Because two orders remain unfinished at days 42 and 45 (20 percent of the cohort), the true 80th percentile is **at least 42.0 calendar days**, proving that the naive query understated the 80th-percentile clock by at least **12.0 calendar days** (a 40 percent downward distortion relative to 30.0 days). [S22] [S0402]

The gap between naive completed-only reporting and censoring-aware bounds across four illustrative material categories is shown in **Exhibit 16**, p. 72. [S22] [S0402]

EXHIBIT 16 / ILLUSTRATIVE CALENDAR DAYS

Omitting open orders and partial lots understates the upper tail



DERIVED [S22] [S30] [S0402] Derived from constructed 10-order illustrative cohorts using the Kaplan-Meier product-limit method (S0402) and supply-censoring principles (S22, S30). Row 1 reflects the worked example in the text (30 days naive vs 42 days censored lower bound).

As **Exhibit 16**, p. 72 demonstrates across the four constructed cohorts, ignoring open orders and partial deliveries compresses the apparent 80th-percentile lead time by 6 to 12 illustrative calendar days. [S22] [S0402] While the day counts are illustrative, the mechanism holds universally whenever open orders are older than the completed median. [S0402]

4.11

Tail measures and sample size constrain confidence

Lead-time distributions in emerging-market corridors are rarely symmetric or Gaussian: physical transit has a hard lower bound set by distance and speed, whereas transshipment rollovers, border inspections and quality re-tests create a skewed right tail. [S12] [S24] [S30]

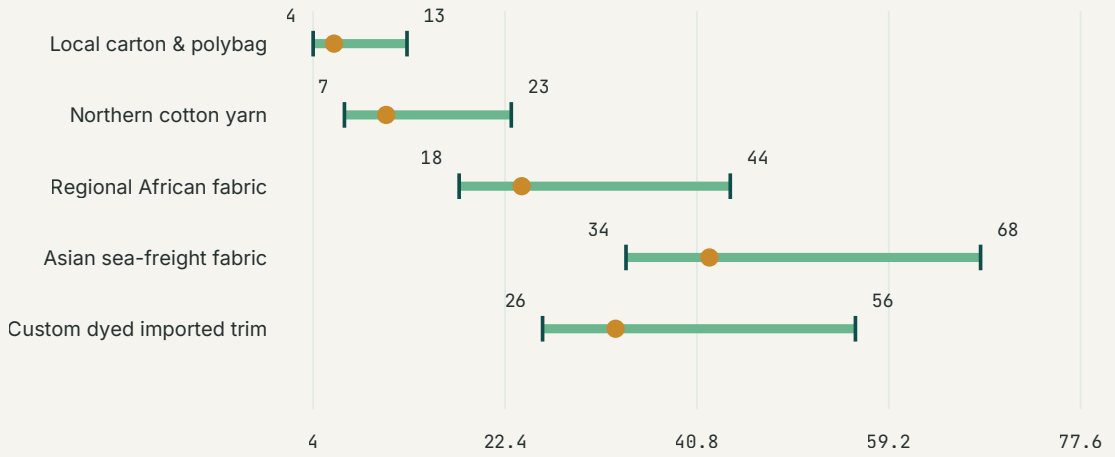
4.12 MEANS AND STANDARD DEVIATIONS FAIL WHEN DISTRIBUTIONS ARE SKEWED OR MULTIMODAL

When 80 percent of shipments clear through a documentary channel in 2 to 4 illustrative days while 20 percent undergo physical scanning or valuation review taking 14 to 22 illustrative days, the distribution is bimodal. [S30] A single Gaussian mean and standard deviation over-buffers routine shipments while under-protecting against inspection holds. [S05] [S30] Consistent with the WCO *Time Release Study Guide Version 4*, GFT requires lead-time summaries to report six mandatory metrics: total and completed sample counts (`N_total`, `N_complete`), observation window dates, median (`P50`), interquartile range (`P75 - P25`) alongside standard deviation (`sigma_L`), upper quantiles (`P80`, `P90`, `P95`) or censored bounds, and the maximum observed and oldest open duration. [S05] [S30] [S0402]

The spread between routine medians and 90th-percentile tails across five illustrative lanes is shown in **Exhibit 17**, p. 73. [S05] [S24] [S30]

EXHIBIT 17 / ILLUSTRATIVE CALENDAR DAYS

A mean hides the tail that exhausts stock



ILLUSTRATIVE [S05] [S24] [S30] Illustrative lead-time intervals across five constructed sourcing lanes. Low = 20th percentile (fast routine receipt); Mid = 50th percentile median; High = 90th percentile delayed tail where physical inspection, transshipment rollovers or shade re-tests occur. Basis: S05, S24, S30.

In **Exhibit 17**, p. 73, every illustrative lane shows strong right skew: the distance from the median (`mid`) to the 90th percentile (`high`) is roughly three times the distance from the 20th percentile (`low`) to the median. [S24] [S30] For the illustrative Asian sea-freight lane, the median is 42 days (20th percentile 34 days), but the 90th percentile reaches 68 days, a 26-day tail extension. [S12] [S30] While **Exhibit 17**, p. 73 is an illustrative model rather than measured Cameroon factory data, it shows why buffering only for a 42-day median leaves the line vulnerable whenever a shipment falls into the upper decile. [S05]

4.13 SMALL SAMPLES AND SPECIAL-CAUSE OUTLIERS REQUIRE EXPLICIT GOVERNANCE

Claiming an empirical 95th percentile from a small sample is mathematically unsound: with `N = 12` monthly orders per year, the sample maximum has an expected cumulative probability of `12 / 13` (about 92.3 percent), and the chance of observing *zero* draws from the worst 5 percent tail is `(0.95)^12` (about 54.0 percent). [S05] GFT therefore enforces three sample-size tiers: below 10 observations, planners must use structured three-point scenario bounds and manual review; from 10 to 29 observations, planners combine the empirical Kaplan-Meier 80th percentile with a scenario tail buffer; and at 30 or more observations, empirical quantiles up to the 90th percentile may be used if stationarity holds. [S05] [S30] [S0402]

Similarly, planners must not delete inconvenient delays (such as an illustrative 75-day port breakdown or storm hold) from the database. [S30] Following WCO outlier guidance, every valid timestamp remains in the master ledger: an extreme point may be excluded from the routine reorder-point layer *only* through an auditable exception

record signed by planning and finance (never for internal payment delays), and every excluded point is automatically transferred into the disruption stress register for Chapter 5 and Chapter 13 scenario testing. [S23] [S27] [S30]

4.14

Demand and lead time may worsen together

Standard safety-stock formulas assume that lead time `L` is independent of daily demand `D_t` and independent across consecutive orders. [S05] [S24] As Michna, Nielsen and Nielsen (2015) note, "The assumption that the lead time is constant is rather unrealistic." [S24] Assuming that random lead times are independent of demand is equally risky when seasonal peaks hit a shared corridor. [S24]

4.15 FOUR MECHANISMS GENERATE POSITIVE COVARIANCE AND AMPLIFY VARIANCE

Positive dependence between demand and replenishment delay (`Cov(D, L)` above zero) arises through four operational channels: [S12] [S24]

- **Synchronised seasonal peaks:** School-uniform, workwear and festive apparel demand clusters around calendar windows while northern cotton harvesting and ginning follow agricultural cycles, causing mill queues and port volumes to peak when factory consumption rises. [R01] [R03] [S12] [S24]
- **Rainy-season corridor slowdowns:** Heavy seasonal rains along overland corridors linking northern Cameroon to southern hubs can slow road transit while rural distribution is active. [R03] [S12]
- **Supplier capacity saturation:** Larger replenishment orders `Q` placed during demand spikes take longer to dye, finish (`T2` to `T3`) and inspect (`T9` to `T10`). [S07] [S24]
- **Order crossovers:** When consecutive lead times fluctuate, a later order can arrive before an earlier delayed order, distorting effective inter-arrival intervals and upstream order variance. [S24]

Under independent daily demand (mean `mu_D`, variance `sigma_D^2`) and independent random lead time (mean `mu_L`, variance `sigma_L^2`), lead-time demand variance is `Var(D_L)_ind = mu_L * sigma_D^2 + (mu_D^2) * (sigma_L^2)`. [S05] [S24] For an illustrative uniform fabric with `mu_D = 200` metres per day, `sigma_D = 30` metres per day and `mu_L = 30` days, assuming a constant lead time (`sigma_L = 0`) gives a standard deviation of `sqrt(30 * 30^2) = 164.3` metres, whereas a modest lead-time standard deviation of `sigma_L = 6` days raises the standard deviation to `sqrt(27,000 + 1,440,000) = 1,211.2` metres, more than **seven times larger**. [S05] [S24] When high daily demand (`240` illustrative metres) coincides with peak-season lead times (`38` illustrative days, requiring `9,120` metres) rather than off-peak lead times (`160` metres over 22 days, or `3,520` metres), positive covariance pushes the upper tail even higher. [S05] [S24]

Planners must first strip predictable calendar shifts from demand using the companion forecasting framework (WP04) and time-phased scheduling so safety stock covers only unpredicted forecast residuals. [S04] [S05] [S07] GFT then requires checking rank

correlation between L_k, demand residuals, lot size Q_k and vendor backlog, stratifying parameters by dry versus rainy/peak season or passing joint (D_L, L) scenarios to the Chapter 5 quantile calculator. [S05] [S23] [S24]

4.16

Incoterms allocate responsibilities but do not erase elapsed time

Switching a purchase contract from Ex Works (EXW) or Free On Board (FOB) to Cost, Insurance and Freight (CIF) or Delivered Duty Paid (DDP) does not shorten physical transit. [S29] The International Chamber of Commerce (ICC) *Incoterms 2020* rules govern the contractual division of cost, risk and carriage responsibilities, noting that "It recognises that not all commercial trade transactions from the seller to the buyer are conducted by a third-party carrier." [S29] Every contract must specify the rule edition and named place (such as FCA Ningbo Terminal, Incoterms 2020), but all ten physical stages from T0 to T10 still occur: [S29]

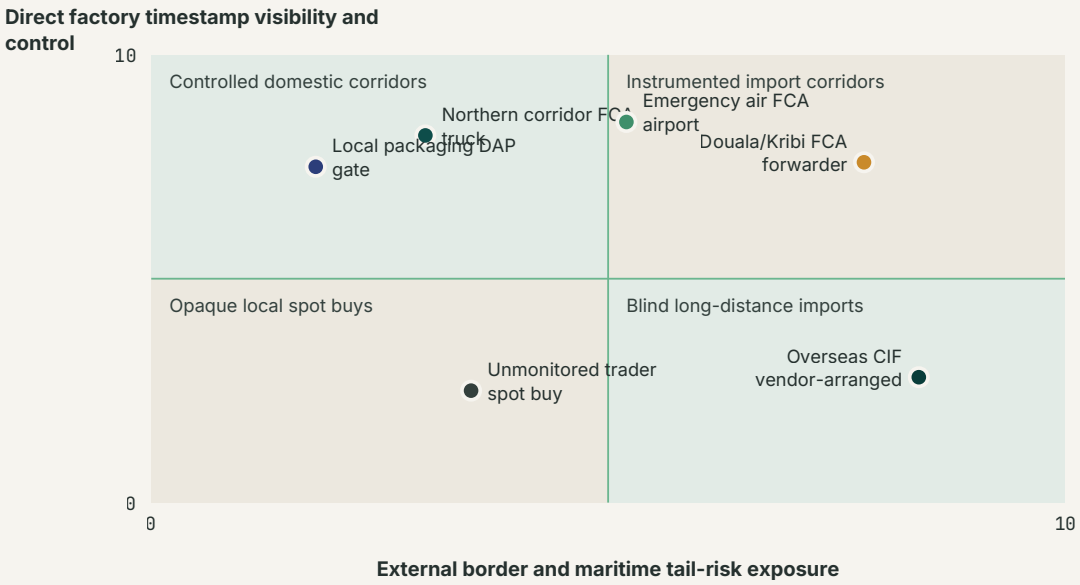
- **Under EXW or FCA (origin):** The buyer's forwarder controls main carriage (T4 to T5), port clearance (T6 to T7) and inland haulage (T8), giving the factory direct timestamp visibility alongside transit risk. [S12] [S29]
- **Under CPT, CIP, CFR or CIF (destination port):** The seller pays main carriage while risk transfers at origin handover, which can leave the buyer blind to transshipment rollovers during T4 to T5 unless milestone reporting is mandated, while the buyer still owns customs clearance (T6 to T7) and haulage (T8). [S12] [S29]
- **Under DAP, DPU or DDP (factory gate):** The seller bears delivery responsibility to the destination, yet an overseas mill lacking licensed customs capability in Cameroon can lengthen T6 and T7 under DDP, and receiving (T9) plus quality release (T10) remain inside the buyer's clock. [S13] [S29] [S30]

Furthermore, *Incoterms 2020* rules do not govern title transfer, payment timing or breach remedies, which require separate contract clauses and legal review (not legal or tax advice). [S08] [S29] Documentary and financing clocks directly affect port release (T6 to T7): delayed original bills of lading hold containers at T6; post-shipment receivables finance such as the International Finance Corporation's Global Trade Supplier Finance programme ("The GTSF Program provides post-shipment finance to suppliers based upon acceptance of receivables by select buyers approved by IFC.") does not fund pre-shipment buyer deposits at T1 or import duties at T7; and under *IAS 2 Inventories* ("Inventories are measured at the lower of cost and net realisable value."), goods in transit may sit on the balance sheet weeks before passing T10. [S09] [S18] [S29]

How route and contract choices reshape risk exposure and timestamp visibility across six sourcing archetypes is mapped in **Exhibit 18**, p. 76. [S12] [S14] [S29]

EXHIBIT 18 / ILLUSTRATIVE INDEX (0 TO 10)

Route choice changes the measurement requirement



ILLUSTRATIVE [S12] [S14] [S29] [S30] Illustrative scoring of six sourcing route archetypes on two planning dimensions: x = External border and maritime tail-risk exposure (0 = minimal border friction, 10 = multi-ocean transshipment plus full customs clearance); y = Direct factory visibility and control over transit timestamps under standard contracts (0 = opaque vendor-arranged chain, 10 = direct factory/forwarder telemetry). Basis: S12, S14, S29, S30.

Exhibit 18, p. 76 contrasts external border and maritime tail exposure (x) with direct factory timestamp visibility (y). [S12] [S29] An overseas `CIF` order with vendor-arranged shipping sits in `Blind long-distance imports` (illustrative score `x = 8.4, y = 2.8`), whereas pairing an `FCA` contract with forwarder milestone telemetry moves the lane into `Instrumented import corridors` (`x = 7.8, y = 7.6`). [S29] [S30] Domestic movements (`Northern corridor FCA truck`, `x = 3.0, y = 8.2`) avoid maritime and customs stages (`T6` to `T7`) but require inland road monitoring (`T8`). [R03] [S12] These illustrative scores do not rank specific Cameroon ports or carriers; they specify the telemetry required for each route. [S14] [S29]

> **WARNING**
>
> Never assume that domestic sourcing automatically guarantees a shorter usable-receipt clock than imports. [S12] Where a domestic mill faces low capacity utilisation or maintenance queues (compare the historical July 2023 CICAM utilisation context in [^R12], cited strictly as sector background), supplier queue and manufacturing stages (`T2` to `T3`) can exceed ocean transit times. [R12] [S30] Every route must be calibrated from measured `T0-to-T10` observations. [S30]

4.17

Reducing uncertainty can be preferable to buying more buffer

When a lead-time audit shows an illustrative 65-day 90th percentile against a 40-day median, buying 25 extra days of safety stock ties up cash and increases obsolescence risk. [S05] [S06] [S09] Because the lead-time term in demand variance scales with `mu_D^2 * sigma_L^2`, halving `sigma_L` removes **75 percent** (`1 - 0.5^2 = 0.75`) of lead-time-driven variance without adding a metre of stock. [S05] [S24]

4.18 FOUR OPERATIONAL LEVERS AND A MULTI-CYCLE EVIDENCE GATE GOVERN LEAD-TIME REDUCTIONS

Engineering management should target four procedural bottlenecks before expanding buffers:

1. **Pre-arrival customs lodgement (`T5` to `T7`):** Article 7.1 of the WTO Trade Facilitation Agreement requires procedures "allowing for the submission of import documentation and other required information, including manifests, in order to begin processing prior to the arrival of goods with a view to expediting the release of goods upon arrival." [S0401] Enforcing a document-complete gate at vessel departure (`T4`) at least 5 illustrative days before berthing eliminates avoidable declaration delays at `T6`. [S30] [S0401]
2. **Customs guarantee separation (`T6` to `T7`):** Using WTO TFA Article 7.3 surety release mechanisms, clean HS-code master tables and pre-funded duty accounts prevents routine valuation queries from causing 15-day container holds. [S30] [S0401]
3. **Pre-shipment shade-band approval (`T3` and `T9` to `T10`):** Inspecting physical or calibrated digital shade cuttings from production rolls at `T3` before container sealing prevents full-order replacement delays at `T9`. [S07] [S27] [S31] (Where high-volume lines later evaluate automated optical inspection, sector literature in `[^R29]` reports manual accuracy of 60 to 70 percent versus computer vision above 90 percent at up to 60 metres a minute, which remains a separate WP06 decision.) [R29]
4. **Scheduled internal approval gates (`T0` to `T2`):** Replacing ad-hoc payment sign-offs (which can add 1 to 11 illustrative days of administrative jitter) with a fixed twice-weekly approval cadence bounds `T0`-to-`T2` within 48 hours. [S07] [S18]

Finally, GFT enforces an asymmetric **lead-time ratcheting rule** (linked to Chapter 12 and Chapter 13): upward delay shifts observed in right-censored or completed orders trigger immediate protection reviews, whereas a promised lead-time reduction may *not* reduce the reorder-point parameter `mu_L` or `P90` until verified across at least **one**

full seasonal cycle and 10 consecutive uncensored shipments, releasing any resulting buffer reduction in two staged steps as fill rate and complete-kit readiness hold steady. [S05] [S22] [S23] [S30]

A shipping notice is a promise in transit; only released, usable material on the factory floor can stop a sewing line from starving.

CHAPTER SOURCE NOTES

S07 MIT OpenCourseWare. 2006.

S30 World Customs Organization. 2025.

S12 World Bank. 2025.

S22 arXiv, Chen, Jiang, Zhang and Zhou. 2023-10-31.

S05 MIT OpenCourseWare. 2006.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S14 World Bank. 2015-02-12.

S29 International Chamber of Commerce. 2020.

S31 GS1 US and Auburn University RFID Lab. 2018.

S18 International Finance Corporation. 2026.

S08 MIT OpenCourseWare. 2006.

S25 OECD. 2026-02-26.

S13 World Trade Organization. n.d..

S0401 World Trade Organization. 2017-02-22.

S15 Association for Information Systems Electronic Library. 2013.

S27 North Carolina Office of the State Controller. 2023-10-24.

R16 WageIndicator. Publication date unrecorded; inherited source-text check.

R03 SODECOTON. Publication date unrecorded; inherited source-text check.

R09 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R10 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S0402 National Institute of Standards and Technology. 2012-04.

S24 arXiv, Michna, Nielsen and Nielsen. 2015-03-20.

S23 arXiv, Sarit Maitra. 2024-06-27.

R01 Business in Cameroon. Publication date unrecorded; inherited source-text check.

S04 MIT OpenCourseWare. 2006.

S09 IFRS Foundation. n.d..

R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

S06 MIT OpenCourseWare. 2006.

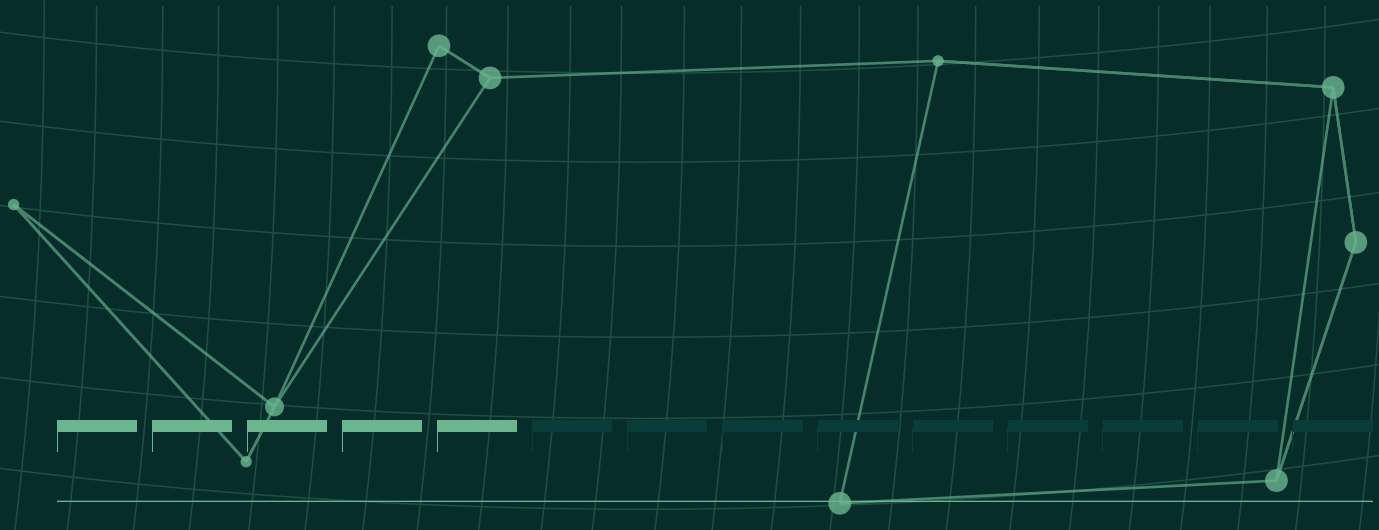
R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

GREEN FUTURE TECH
LIMITED / HONG KONG SAR

05

Safety stock works only when its assumptions match the replenishment process

Calculate protection against uncertainty with consistent time units, explicit service definitions and checks for non-normal or dependent demand.



CHAPTER 05
DECISION TEST

Calculate protection against uncertainty with consistent time units, explicit service definitions and checks for non-normal or dependent demand.

Safety stock is not a universal allowance added to a forecast. It is the quantity required to protect a defined service event over a defined protection horizon, under an explicit model of demand, usable supply and review timing. Green Future Tech Limited (GFT), Hong Kong SAR, therefore starts with the replenishment process rather than the normal-distribution formula. A continuous-review reorder point covers demand until a firm order becomes usable. A periodic-review order-up-to level also covers the wait until the next review. Random lead time adds variance only through a stated dependence model, while cycle service and unit fill rate answer different questions. [S05] [S24] This chapter derives those rules, tests them against empirical and joint scenarios, and specifies calculator guards and a manual fallback. Every numerical case is illustrative. No factory demand history, supplier receipt sample or operating result has been supplied, and nothing here authorises an automatic purchase order.

5.1

Inventory position and available stock answer different questions

Available stock is a physical and quality judgement. Inventory position is a planning judgement. Confusing them can make a reorder rule fire when sufficient supply is already committed, or fail to fire when the recorded supply cannot be used. GFT defines usable on-hand stock as accepted material at the item location, less quantities whose reservation has not already been netted from requirements. Quarantine, rejected lots, unapproved shade lots and stock reserved for another order are not interchangeable with free stock. Material requirements planning likewise depends on consistent item definitions, requirements and supply netting rather than an undifferentiated stock total. [S07]

For this chapter, inventory position is `IP = U + F - B - C`. `U` is usable unreserved on-hand quantity. `F` is firm outstanding supply expected to become usable, expressed in the same item unit. `B` is backorders, meaning unfilled demand retained as an obligation. `C` is other relevant commitments that have not already reduced `U` or the demand forecast. This is a proposed GFT convention built from standard inventory-position logic, not a claim that one formula fits every ledger. [S05] [S07] The implementation must document whether allocations sit inside `U`, inside `C` or inside netted demand. They must appear once, never twice.

Consider an illustrative fabric record with 800 accepted metres in the store, 150 metres quarantined for shade inspection, 250 metres reserved to a confirmed uniform order, and a firm purchase order for 600 metres. Backorders total 100 metres. No other commitment remains outside the demand calculation. Usable unreserved stock is `800 - 150 - 250 = 400` metres if the 800-metre book balance includes both quarantined and reserved material. Inventory position is then `400 + 600 - 100 = 900` metres. Subtracting the reservation again when netting requirements would understate supply by 250 metres; treating quarantine as usable would overstate it by 150 metres. The example demonstrates reconciliation, not a recommended stock level.

The firm-receipt term needs the same discipline. A supplier confirmation is not proof that the full quantity will arrive on the stated date or pass inspection. Research on uncertain supply identifies censoring in both demand and supply observations, so an

accepted or observed quantity can be only a truncated view of what was requested or dispatched. [S22] In the Cameroon sector context, where the inherited annex records about 95 percent of cotton fibre exported raw and about 5 percent processed locally against a national 50 percent local-processing ambition by 2030, a garment factory scaling repeat uniform or basics production cannot equate gross inbound deliveries with usable stock. [R06] [R07] Inspection literature stored in the same annex reports manual fabric-defect detection at 60 to 70 percent accuracy, missing 20 to 30 percent of defects, compared with computer-vision inspection above 90 percent across more than 40 defect types at up to 60 metres a minute. [R29] That inspection boundary is an inherited literature benchmark rather than a measured GFT factory result, but its planning implication is direct: GFT would admit a receipt into `F` only when its quantity, item version, promised usable date and quality-release status meet an approved rule. Disputed quantities, unapproved specifications and unreliable dates should be excluded or represented through explicit receipt scenarios rather than counted at full face value. [S22] [R29]

> **METHOD**
>
> Use a four-way reconciliation before calculating any reorder quantity:
>
> 1. **Physical to usable.** Match the counted balance to accepted, quarantined, rejected and reserved states, with unit conversions preserved. [S07]
> 2. **Commitments once.** Trace each customer allocation and backorder to either the stock deduction or the commitment deduction, not both.
> 3. **Receipts to usable dates.** Link each firm receipt to its expected quality-release date and accepted quantity, retaining partial and censored observations. [S22] [S30]
> 4. **Demand to the same horizon.** Confirm that forecast demand, reservations and backorders describe the same item, location and period before comparing inventory position with the reorder point. [S05]

This boundary also clarifies what the reorder point is not. It is not a physical minimum that must always remain on the shelf. Because the trigger uses inventory position, outstanding firm supply can legitimately keep `IP` above the threshold while physical stock is low. Conversely, physical stock can look ample while unusable lots and outstanding commitments make the effective position inadequate. The planner should be able to inspect every component of the calculation and the event time of every receipt.

5.2

The reorder point protects demand over lead time

Under continuous review, the reorder decision should protect demand from the effective order release until usable receipt. Let `μ` be mean daily demand and `σ_d` the standard deviation of the daily forecast residual, both in pieces per day. Let `L` be a fixed lead time measured in days. With stationary, independent daily residuals and a normal approximation to lead-time demand, expected demand over the horizon is `μL`

and its variance is $L\sigma_d^2$. The resulting rule is $ROP = \mu L + z\sigma_{dVL}$, where safety stock is $z\sigma_{dVL}$. [S05] The square root follows from adding independent daily variances, not daily standard deviations.

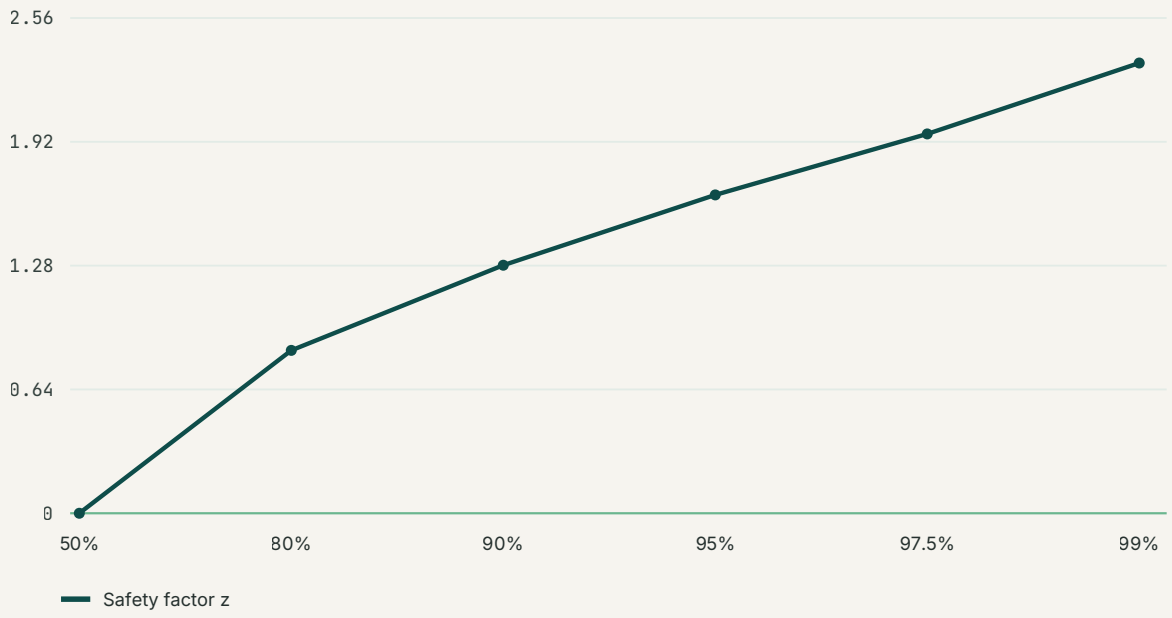
The input must be residual uncertainty when the mean demand term already uses a forecast. Raw sales variation can include explainable seasonality, price effects or order calendars that the forecast has separately represented. Using raw variation alongside that forecast can count predictable movement twice. WP04 owns estimation of the forecast and its residual distribution; this chapter specifies how the replenishment rule consumes those outputs. If no defensible residual estimate exists, the honest response is a simpler reviewed policy with declared uncertainty, not a precise safety factor attached to an arbitrary standard deviation.

The safety factor z targets cycle service under the normal approximation. Cycle service is the probability that demand during the replenishment cycle does not exceed available protection, so no stockout occurs in that cycle. It is not the share of requested units filled immediately. MIT's teaching material presents cycle service and item fill rate as distinct service measures with different calculations. [S05] A manager who sets z from a desired fill rate without calculating expected shortages has changed the meaning of the target without changing the formula.

The illustrative quantiles in **Exhibit 19**, p. 82 show the nonlinear price of reducing stockout probability. They are inverse standard-normal values, not observed factory performance. [S05]

EXHIBIT 19 / STANDARD NORMAL SAFETY FACTOR Z

Higher cycle service buys diminishing reductions in stockout probability



ILLUSTRATIVE [S05] Illustrative inverse standard-normal quantiles computed for selected cycle-service probabilities. The complementary stockout probabilities are 50, 20, 10, 5, 2.5 and 1 percent respectively. Values use unrounded normal quantiles before display rounding; they are not empirical service results or recommended universal targets.

Moving from an illustrative 90 percent to 95 percent cycle service raises z from 1.281552 to 1.644854, an addition of 0.363302. Moving from 97.5 percent to 99 percent adds 0.366384, while halving the nominal stockout probability from 5 percent to 2.5 percent between the intervening targets adds 0.315110. The lesson is not that a higher target is always wasteful. It is that service ambition must be converted into stock and cash with the actual residual scale, then approved explicitly. [S05]

> **WORKED EXAMPLE: A FIXED-HORIZON REORDER POINT**
>
> An illustrative repeat fabric has mean demand of 100 pieces per day, a daily forecast-residual standard deviation of 20 pieces, and a fixed usable-receipt lead time of 10 days. Assume independent daily residuals, stationary demand and normal lead-time demand. For an illustrative 95 percent cycle-service target, `z = 1.644854` before rounding. [S05]
>
> Expected lead-time demand is `100 × 10 = 1,000` pieces. Its standard deviation is `20 × √10 = 63.25` pieces. Safety stock is therefore `1.644854 × 63.245553 = 104.03` pieces, and the reorder point is `1,104.03` pieces before integer rounding. If pieces are indivisible and the policy rounds upward, the trigger is 1,105 pieces of inventory position. These values demonstrate arithmetic only; they are not a calibrated buffer for any factory. [S05]

The worked result is highly sensitive to the time unit. Entering a weekly residual standard deviation beside daily demand would produce a buffer with no coherent unit. Using calendar days for demand but working days for receipt would create a similar error. The protection horizon must be the actual elapsed interval between effective decision and usable availability, with demand transformed to that same calendar basis. [S05] [S30]

5.3

Random lead time adds a different variance term

A fixed-lead-time model removes an uncertainty that many replenishment processes actually contain. The stochastic-lead-time literature explicitly challenges the constant-lead-time assumption and shows why its distribution matters to inventory behaviour. [S24] However, replacing a fixed `L` with an average lead time inside `σ_dVL` is not a sufficient correction. Lead-time variability changes both the number of daily demands accumulated and the variability of that accumulation.

Let `D_L` denote demand during a random lead time `L`. Conditional on a particular realised lead time, independent identically distributed daily demand gives `E[D_L | L] = μL` and `Var(D_L | L) = Lσ_d²`. The law of total variance then gives `Var(D_L) = E[Var(D_L | L)] + Var(E[D_L | L])`. Under stationarity and independence between daily demand and lead time, this becomes `Var(D_L) = E[L]σ_d² + μ²Var(L)`. The first term is demand variability accumulated over the expected horizon. The second is variability caused by an uncertain number of demand days. [S05] [S24]

The units verify the derivation. If μ is pieces per day and σ_d^2 is pieces squared per day, while L is days, both terms reduce to pieces squared. A formula that adds Var(L) directly to daily demand variance would mix days squared with pieces squared and should be rejected. This check is simple enough to perform manually and catches a consequential class of spreadsheet errors.

For an illustrative comparison, retain mean daily demand of 100 pieces and daily residual standard deviation of 20 pieces. Let random lead time have mean 10 days and standard deviation 3 days, independently of demand. Then $E[L]\sigma_d^2 = 10 \times 400 = 4{,}000$ pieces squared, while $\mu^2 Var(L) = 100^2 \times 3^2 = 90{,}000$ pieces squared. Total lead-time-demand variance is 94,000 pieces squared and its standard deviation is 306.59 pieces. At the same illustrative $z = 1.644854$, safety stock becomes 504.30 pieces and expected lead-time demand remains 1,000 pieces, giving a reorder point of 1,504.30 pieces before rounding. [S05] [S24]

The comparison with the fixed-lead-time case is stark: the illustrative safety stock rises from 104.03 to 504.30 pieces because lead-time standard deviation of 3 days introduces the second variance term. This is not a general ratio that should be transferred to another item. Lower mean demand would reduce the $\mu^2 Var(L)$ contribution, while a lower lead-time variance would reduce it directly. A model that combines a shorter average lead time with a wider tail can therefore require more protection even when the average looks improved.

The independence condition is equally important. Demand and lead time can worsen together when seasonal order peaks consume supplier capacity or when a disruption affects both material availability and customer requests. The stochastic-lead-time manuscript identifies dependence between lead times and demand as an area requiring further work, which cautions against treating independent-model outputs as universal. [S24] Similarly, inventory research on correlated merchandise notes that "it is difficult to control the stock management of correlated marketable merchandise when using an inventory control method that was developed under the assumption that the demands are not correlated." [S0501] When multiple garment sizes, colours or component trims share a fabric base or respond to the same institutional tender, serial and cross-item correlations invalidate simple sum-of-variances rules. If orders also change supplier congestion, lead time is endogenous to the buying rule rather than an external random draw. In those circumstances, GFT would use joint historical blocks, probability-inequality bounds or explicit scenarios linking order size, demand state and usable receipt instead of applying the two-term formula as though independence had been proved. [S24] [S0501]

> **GFT NOTE**
>
> Before accepting a larger buffer from a random-lead-time calculation, ask whether the business can reduce the variance itself. Separate supplier queue, carriage, clearance, receipt and quality-release times; retain open orders and exceptional delays; then test whether better confirmation, document preparation or inspection scheduling changes the usable-receipt distribution. [S24] [S30] A buffer can protect service, but it cannot by itself repair an unreliable process. Compare the recurring stock commitment with the cost and feasibility of removing the delay mechanism.

A practical policy record should consequently store more than mean and standard deviation. It should retain the observation window, number of completed and open orders, calendar convention, exclusions and their reasons, the assumed dependence model, and the person or role authorised to approve it. Sparse data may support a broad scenario range but not a confident variance estimate. The output should say so rather than conceal sampling weakness behind additional decimal places.

5.4

Periodic review protects the review gap as well

Continuous review assumes that the planner can act whenever inventory position crosses the trigger. Periodic review instead evaluates the item at scheduled events. Just after one review, the next opportunity to release an order is one review interval away. The policy must therefore protect demand across the review interval `P` plus the lead time `L`, not across `L` alone. This longer protection horizon is the central mathematical difference, not a supplementary allowance chosen by habit. [S05]

With fixed `L`, stationary independent daily residuals and the normal approximation, the order-up-to level is `S = μ(P + L) + zσ_dv(P + L)`. At a review, the unconstrained order quantity is `max(0, S - IP)`, using the same reconciled inventory position defined earlier. The calculation must include all firm open orders because another order placed now cannot cancel supply already committed. It must also respect the timing of those receipts: a distant or unreliable order should not be treated as equivalent to stock arriving inside the protection horizon. [S05] [S22]

Periodic review can reduce monitoring and ordering effort by grouping decisions, but it changes both service exposure and order cadence. Simulation research comparing periodic and continuous policies provides a useful method for examining that trade-off on common demand paths, not a universal ranking that can be transferred to an apparel factory. [S23] The appropriate comparison holds item identity, usable-receipt rules, service definition, lot constraints and demand paths constant. Otherwise a difference attributed to review policy may actually come from a changed forecast or supplier assumption.

The conditional equations and their principal assumptions are consolidated in **Exhibit 20**, p. 86. The table separates the mathematical object being calculated from the service claim it can support. [S05] [S24]

EXHIBIT 20 / SYMBOLS DEFINED IN NOTES

The safety-stock equation is conditional on explicit assumptions

Decision	Conditional expression	Material assumptions	Unit
Continuous review	ROP = μL + zσ_d√L	Fixed L; independent stationary daily residuals; normal lead-time demand	pieces
Random lead-time variance	Var(D_L) = E[L]σ_d² + μ²Var(L)	Independent daily demand and lead time; stationary first and second moments	pieces squared
Periodic review	S = μ(P + L) + zσ_d√(P + L)	Fixed L; review gap P; independent stationary daily residuals	pieces
Cycle service	CSL = Pr(D_H ≤ ROP)	A cycle event under the stated demand and supply model	probability
Unit fill rate	FR = 1 - E[B]/Q	Fixed-lot backorder model with consistent cycle and demand definitions	share of units

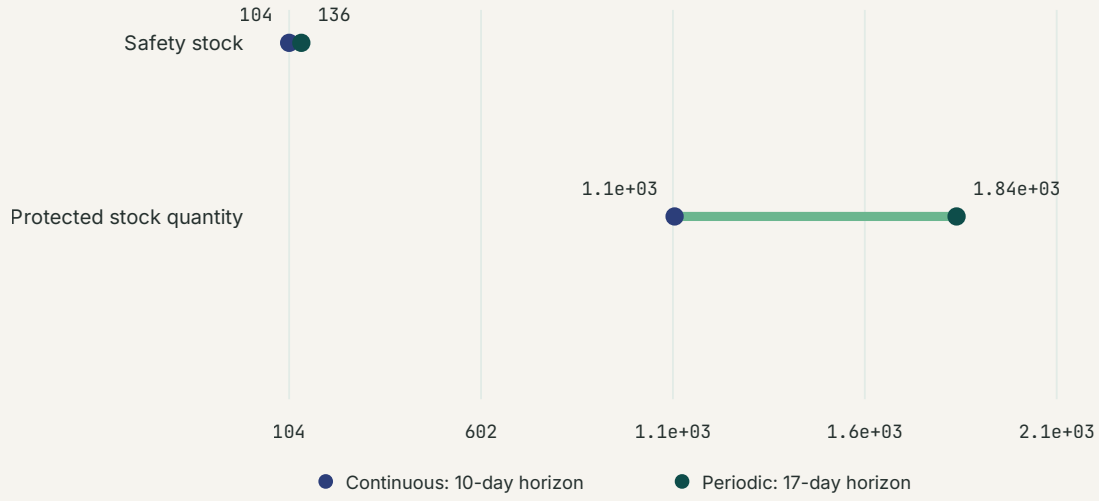
DERIVED [S05] [S24] Derived method summary from S05 and the law of total variance, with stochastic-lead-time limitations from S24. μ is mean daily demand, σ_d is daily forecast-residual standard deviation, L is usable-receipt lead time in days, P is review interval in days, H is protection horizon, Q is fixed lot size and E[B] is expected backorders per replenishment cycle. The fill-rate expression excludes reorder undershoot and assumes a continuous-review fixed-lot backorder model. No factory values are asserted.

To isolate the review-gap effect, use the earlier illustrative inputs of mean daily demand 100 pieces, daily residual standard deviation 20 pieces, fixed lead time 10 days and the same unrounded `z = 1.644854`. Continuous review protects 10 days and requires 104.03 pieces of safety stock. With an illustrative review interval of 7 days, periodic review protects 17 days, so expected demand is 1,700 pieces, safety stock is `1.644854 × 20 × √17 = 135.64` pieces, and the order-up-to level is 1,835.64 pieces before integer rounding. [S05]

Exhibit 21, p. 87 compares those quantities. It does not compare expected physical stock, because cycle stock, order timing, lot size and actual demand also determine what is held. [S05]

EXHIBIT 21 / ILLUSTRATIVE PIECES

Adding a review gap changes required protection



ILLUSTRATIVE [S05] Illustrative arithmetic using μ = 100 pieces/day, σ_d = 20 pieces/day, fixed L = 10 days, P = 7 days for periodic review and unrounded z = 1.6448536269514715. Continuous-review safety stock = zσ_d√L = 104.029678 pieces and ROP = μL + safety stock = 1,104.029678 pieces. Periodic-review safety stock = zσ_d√(P+L) = 135.638105 pieces and order-up-to level = μ(P+L) + safety stock = 1,835.638105 pieces. Values shown before integer rounding; not observed stock or a recommended target.

The additional 31.61 pieces of illustrative safety stock is the price of protecting the extra 7 days under these assumptions, while the 731.61-piece difference between trigger and target also includes 700 pieces of additional expected demand. Confusing those two components would misdescribe the buffer. A planner should report horizon demand and uncertainty protection separately.

Lot constraints come after the unconstrained requirement, not inside the statistical horizon. Supplier minimums, pack multiples and finite replenishment can change feasible order size and timing. [S03] Industrial inventory optimisation models that integrate review periods, minimum order quantities and distinct service indicators (fill rate and cycle service levels) have to represent those constraints explicitly and validate the resulting buffers through simulation. [S0502] If the calculated order is 420 metres but the approved pack is 500 metres, rounding to the pack changes cycle stock and possibly cash exposure; it does not retroactively change lead-time demand. Where several open orders coexist, the calculation should display each receipt and its expected usable date before another release is proposed. A policy that silently treats the order-up-to level as an order quantity will repeatedly overbuy.

5.5

An empirical quantile can be safer than an unsupported bell curve

The normal approximation is useful because it makes the protection calculation transparent. It is not evidence that actual lead-time demand is normal. Apparel material demand can contain zeros, tender-sized jumps and size-colour substitutions, while usable lead time can have a discrete clearance or inspection tail. A symmetric model can understate one side of that process. The relevant alternative is not to

abandon mathematics, but to calculate the quantile from a distribution that better represents the observed mechanism and to disclose how little the tail may contain. [S05] [S23]

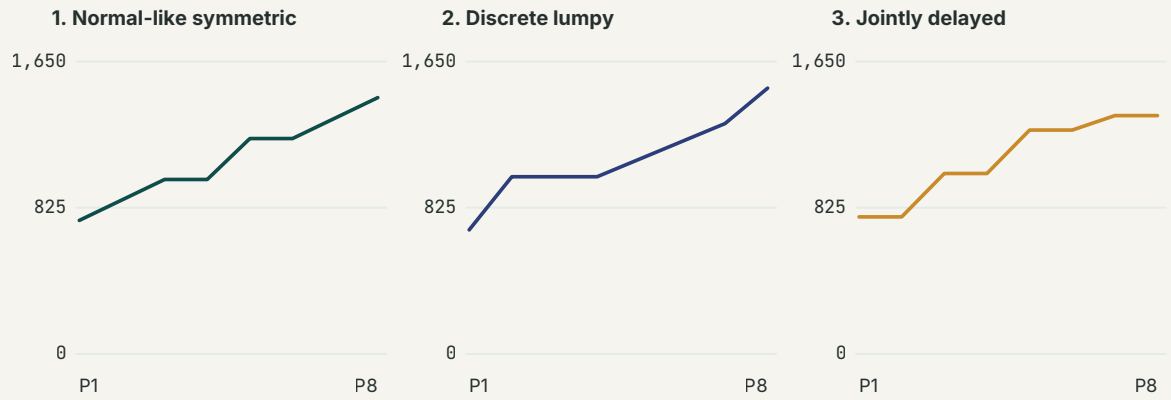
An empirical lead-time-demand quantile begins by constructing complete historical protection intervals. Each interval needs the demand that was knowable only afterwards, the usable-receipt duration and the information available at its start. Sales during a stockout cannot simply be counted as total demand, because lost sales censor the observation. Supply observations can also be truncated. [S22] Where those records are incomplete, a sensitivity range for unobserved demand is more honest than a single calibrated percentile.

Block resampling offers another route when consecutive days or joint disruptions must remain together. Rather than drawing each daily demand and each lead time independently, the method samples contiguous demand blocks and preserves their association with supplier or route states. Simulation provides a way to replay policies over these constructed paths and compare service, stock and order behaviour. [S23] It does not validate the chosen generator. A simulation built from the wrong dependence structure will reproduce that error with great numerical precision.

The constructed paths in **Exhibit 22**, p. 88 make the distribution issue visible. All three have the same unrounded mean of 1,100 pieces and population variance of 50,000 pieces squared, but their shapes and maxima differ. They are teaching scenarios, not fitted distributions or probabilities for a Cameroon factory. [S22] [S23] [S24]

EXHIBIT 22 / ILLUSTRATIVE PIECES OF LEAD-TIME DEMAND

The same mean and variance need not imply the same tail risk



ILLUSTRATIVE [S22] [S23] [S24] Eight constructed paths per panel. Unrounded, all share mean 1,100 pieces, population variance 50,000 pieces squared and standard deviation 223.606798 pieces. Panel 1 scales symmetric deviations [-300,-200,-100,-100,100,100,200,300] by √(4/3). Panel 2 uses lumpy deviations [-400,-100,-100,-100,0,100,200,400]. Panel 3 scales joint delays [-400,-400,-100,-100,200,200,300,300] by √(2/3), producing negative skew. No probabilities are assigned. Censoring and lead-time cautions follow S22 and S24; scenario comparison follows S23.

Under a normal approximation, the shared moments imply an illustrative 95 percent lead-time-demand quantile of `1,100 + 1.644854 × 223.606798 = 1,467.80` pieces. The lumpy panel contains a 1,500-piece path above that value, while the symmetric and jointly delayed panels have maxima of 1,446.41 and 1,344.95 pieces respect-

ively. This does not make the maximum a valid 95 percent estimator. With only eight constructed paths, the upper tail is deliberately sparse. The exhibit proves only that mean and variance do not determine the quantile.

An empirical policy should therefore show the sample count, the exact quantile convention and the sensitivity to individual tail observations. A planner should see whether removing one disruption changes the recommendation materially. If it does, the appropriate output is a range and an evidence request, not a confident single buffer. New demand regimes, supplier changes and unrecorded lost sales can all break the resemblance between the historical sample and the next replenishment cycle. [S22] [S24]

The choice between parametric and empirical methods is consequently conditional. Use the normal formula when its assumptions are defensible and its transparency aids review. Use empirical quantiles, probability-inequality bounds or block-resampled quantiles when the record supports their tail and dependence structure without justifying a Gaussian shape. [S23] [S0501] Use explicit scenarios when neither sample size nor independence is adequate. In every case, evaluate the selected rule on paths not used to choose its parameters before giving it operational authority. [S23]

5.6

Fill rate requires expected shortages and order quantity

Cycle service counts replenishment cycles without a stockout. Unit fill rate counts requested units filled immediately. A cycle can contain a stockout of one unit or many units, and those outcomes receive the same cycle-service failure mark but different fill-rate consequences. For a continuous-review, fixed-lot model with backorders, the standard approximation is `FR = 1 - E[B]/Q`, where `E[B]` is expected shortage per replenishment cycle and `Q` is order quantity. [S05] The order quantity belongs in the denominator because expected units short per cycle are spread across the units replenished in that cycle.

For normally distributed lead-time demand with standard deviation `σ_L` and safety factor `k`, expected shortage is `E[B] = σ_LG(k)`. The standard-normal loss function is `G(k) = φ(k) - k[1 - Φ(k)]`, where `φ` is the density and `Φ` the cumulative distribution. This calculation uses the full tail beyond the reorder point, not just the probability of crossing it. [S05] It therefore contains information that cycle service alone discards.

Return to the illustrative fixed-lead-time case with `σ_L = 63.245553` pieces and `k = 1.644854`. The unrounded loss-function value is 0.020892959, so expected shortage is 1.321387 pieces per cycle. With an illustrative lot size of 500 pieces, the fill-rate approximation is `1 - 1.321387/500 = 99.7357 percent`. With the same demand path, reorder point and cycle service but a lot size of 2,000 pieces, it becomes `1 - 1.321387/2,000 = 99.9339 percent`. The same 95 percent cycle-service assumption therefore does not imply a 95 percent fill rate, and changing the lot changes the unit-service result without changing `z`. [S05]

This arithmetic is an approximation with a specific inventory process. Reorder undershoot can occur when demand crosses the reorder point by more than a small amount, so the order is triggered below the nominal threshold. That increases exposure beyond the simple fixed-lot derivation. Under lost sales, unmet demand disappears rather than becoming a backorder, and recorded sales no longer reveal the full request. [S22] A calculator must identify which process it represents rather than silently applying the backorder formula to every item.

The practical target-setting sequence is therefore reversed from common spreadsheet practice. First define whether the commercial promise concerns cycles, units, complete orders or another event. Then choose a model that calculates that event, including lot size, minimum order quantities and undershoot where material; recent multi-echelon optimisation work notes that setting safety stocks against a fill-rate target requires an explicit mathematical approximation of the fill-rate relationship validated by simulation. [S05] [S0502] Finally compare the resulting stock, expected shortages and cash with the approved promise. A high item fill rate can still coexist with an incomplete customer order when one required component is missing, so the item-level result must not be relabelled as order on-time-in-full performance. [S05] [S07]

5.7

The calculator must reject misleading inputs

A safety-stock calculator should behave as a controlled decision instrument, not as a form that returns a number whenever every cell is populated. GFT's proposed `safety-stock` tool contract begins with validation. It must reject negative demand, stock or lead time; zero or negative lot size where fill rate is requested; missing residual variance; incommensurable time units; and service probabilities outside the open interval between zero and one. It must also reject an inventory-position record whose reservations, backorders or firm receipts cannot be reconciled. [S05] [S07]

The tool should require the user to select a replenishment process before selecting an equation. Continuous review needs a fixed or explicitly random lead-time model. Periodic review needs the review interval and all open orders. A fill-rate request needs lot size, a backorder or lost-sales convention and an undershoot treatment. An empirical quantile needs the sample, interval construction and quantile rule. If the user selects independent demand and lead time, the tool should display that assumption beside the output and provide a joint-scenario route when it is not supported. [S22] [S24]

Every output should carry the input vintage and the calculation boundary. The minimum useful record includes item, location, unit, forecast version, residual definition, lead-time sample window, review calendar, service definition, horizon, formula, unrounded result, rounding rule and resulting order action. Derived values should show their inputs rather than only a final reorder point. This is particularly important when a supplier or quality status changes: the planner must be able to distinguish a demand-driven change from a receipt-driven one. [S05] [S30]

METHOD

GFT would test the calculator in this order:

1. **Arithmetic identities.** Reproduce the fixed-horizon, random-lead-time, review-gap and loss-function calculations independently, including unit checks and rounding.
2. **Invalid-input tests.** Confirm that missing variance, mixed days and weeks, negative quantities, duplicate receipts and out-of-range probabilities stop the recommendation rather than producing a warning beside an order.
3. **Common-path comparison.** Replay continuous, periodic and empirical policies on identical synthetic demand and usable-receipt paths, with the generator and seeds retained. [S23]
4. **Stress and censoring tests.** Add joint demand-delay states, partial receipts, reorder undershoot and lost-sales intervals, and expose the effect of each assumption. [S22] [S24]
5. **Shadow reconciliation.** Compare the proposed recommendation with a planner's calculation using only information available at the decision time, without executing an order.

Validation does not end when the equations match. The tool can still be wrong for the process if it counts quarantined stock, omits a review gap or interprets a cycle target as unit fill. Its output should therefore remain a recommendation with named approval until field evidence supports greater authority. The mathematical contract in this chapter specifies the `safety-stock` tool; it does not claim that the tool has been implemented, calibrated or accepted.

GFT NOTE

Adopt a manual fallback before automating the rule. For each approved repeat item, retain a dated card showing the reconciled inventory-position convention, protection horizon, demand and receipt evidence, service definition, calculation or scenario range, order bounds and next review. Where the sample cannot support calibration, use a planner-approved coverage rule with an explicit review date and an exception trigger. Do not fill the evidential gap with an invented normal distribution. Reassess the policy after a material demand, supplier, route, quality or calendar change. [S05] [S22] [S24]

The decision consequence should remain measurable. GFT would compare recommended and realised cycle service, requested-unit fill, complete-order performance, average and aged stock, order frequency, expediting and cash exposure on the same item-location boundary. It would also record overrides and their reasons. A fallback is appropriate when stock capture is unreliable, when tail evidence is too sparse, or

when correlated disruptions dominate the independent model. A larger buffer is appropriate only when it protects the approved promise at an accepted cash and operating cost, not when the spreadsheet merely makes it calculable. [S05] [S23]

Protect the horizon that the process actually creates, and make every assumption visible before stock becomes a commitment.

CHAPTER SOURCE NOTES

S05 MIT OpenCourseWare. 2006.

S24 arXiv, Michna, Nielsen and Nielsen. 2015-03-20.

S07 MIT OpenCourseWare. 2006.

S22 arXiv, Chen, Jiang, Zhang and Zhou. 2023-10-31.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R29 MDPI; industry inspection literature. Publication date unrecorded; inherited source-text check.

S30 World Customs Organization. 2025.

S0501 arXiv, Takashi Shinzato. 2017-01-09.

S23 arXiv, Sarit Maitra. 2024-06-27.

S03 MIT OpenCourseWare. 2006.

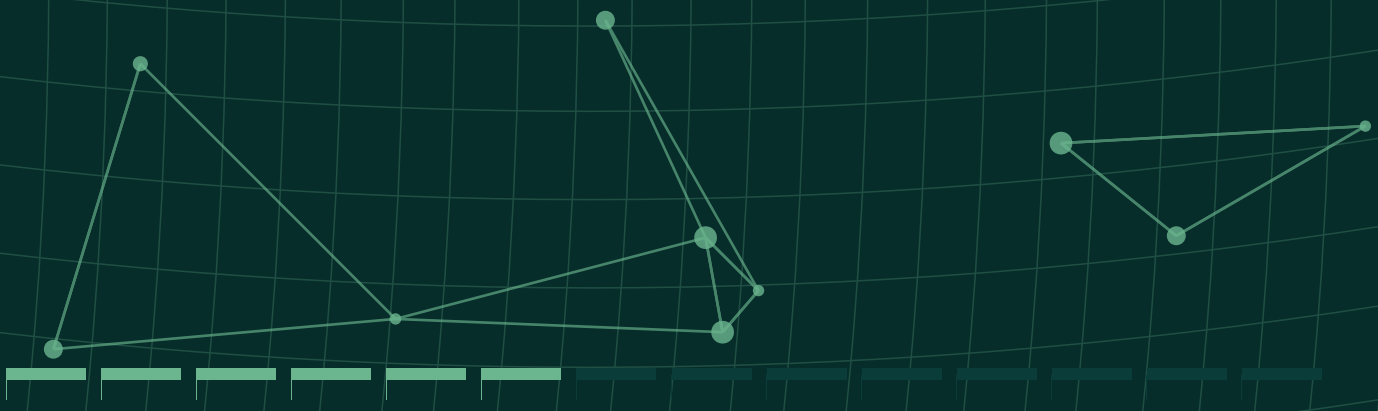
S0502 arXiv, Achkar, Brunaud, Perez, Musa, Mendez and Grossmann. 2023-06-19.

06

PART II: MATCH REPLENISHMENT RULES TO MATERIAL AND DEMAND / LOT SIZING

A cheaper purchase price can produce a more expensive inventory policy

Order quantity must reflect holding costs, pack sizes, supplier minimums, time-phased needs and available cash.



CHAPTER 06
DECISION TEST

Order quantity must reflect holding costs, pack sizes, supplier minimums, time-phased needs and available cash.

An order quantity is a cash commitment before it is a purchasing bargain. Green Future Tech Limited (GFT), Hong Kong SAR, would therefore approve a lot only after comparing total relevant cost, dated cash exposure and terminal material risk across feasible quantities. The economic order quantity (EOQ) supplies a useful benchmark for stable repeat demand, but it does not choose a service policy, accept a supplier minimum or make an unaffordable order feasible. A quantity discount is valuable only when its price and logistics savings exceed the additional ordering, holding, financing and residual exposure. Time-phased requirements, pack sizes, storage, shade continuity and supplier capacity then determine which quantity can actually be released. [S02] [S03] [S07] Every numerical case in this chapter is illustrative. No factory demand ledger, supplier quotation, warehouse tariff, funding offer or realised write-down has been supplied, and nothing here authorises a purchase order.

6.1

EOQ is a benchmark with restrictive conditions

The basic EOQ model answers a narrow question: under constant known demand, instantaneous replenishment and no quantity constraints, which fixed lot balances the annual cost of placing orders against the annual cost of holding cycle stock? It does not answer how much protection to hold against uncertain demand, when to release an order or whether a seasonal residual can be sold. Those decisions require their own replenishment and demand models. [S02] [S05]

Let D be annual demand in units, K the fixed incremental cost of placing and receiving one order in CFA franc (XAF), Q the order quantity in units and h the annual holding cost per unit in XAF. With demand consumed evenly and each order arriving in full just as stock is exhausted, average cycle stock is $Q/2$. Annual relevant cost before any constant purchase-cost term is $TRC = DK/Q + hQ/2$. The first term falls as orders become larger and less frequent. The second rises with average cycle stock. [S02]

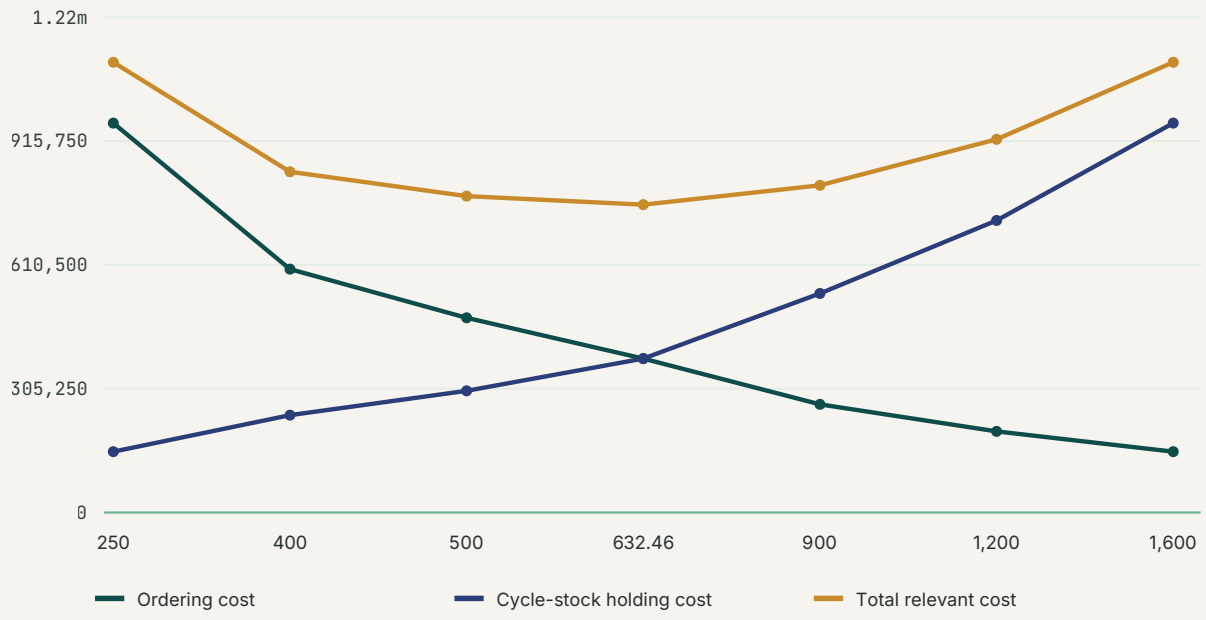
The minimum follows from differentiating this convex expression and setting the result to zero: $-DK/Q^2 + h/2 = 0$. The resulting benchmark is $Q^* = \sqrt{(2DK/h)}$. At that quantity, annual ordering cost equals annual cycle-stock holding cost. The second derivative is positive for positive Q, confirming a minimum rather than a maximum. [S02] This derivation also explains why a spreadsheet should show both cost components. A total without its mechanism cannot reveal whether an implausible order cost or holding charge is driving the answer.

The purchase price c can be added as Dc when it is constant across the quantities compared. It then changes the level of annual cost but not the EOQ. Once price depends on quantity, however, Dc becomes part of the choice and each feasible price bracket must be evaluated. Freight brackets, minimum lots, pack multiples, finite replenishment and cash limits similarly move the problem beyond the basic formula. [S03]

Exhibit 23, p. 95 uses the outline's constructed repeat-item case. It shows the falling ordering component, the rising holding component and their total. The curve is arithmetic, not a supplier recommendation or an observed factory cost function. [S02]

EXHIBIT 23 / ILLUSTRATIVE XAF PER YEAR

The EOQ minimum is a benchmark, not a purchase instruction



ILLUSTRATIVE [S02] Illustrative deterministic case with annual demand D = 12,000 units, fixed order cost K = XAF 20,000 per order and annual holding cost h = XAF 1,200 per unit. Ordering cost = DK/Q, holding cost = hQ/2 and total relevant cost = their sum. EOQ = √(2DK/h) = 632.455532 units and minimum annual relevant cost = XAF 758,946.64 before purchase cost and all excluded constraints. Values are calculated before display rounding.

WORKED EXAMPLE: CALCULATE, THEN CHALLENGE, THE EOQ

For the illustrative case, annual demand is 12,000 units, fixed order cost is XAF 20,000 and annual holding cost is XAF 1,200 per unit. The benchmark is `√(2 × 12,000 × 20,000 / 1,200) = 632.46` units. It implies 18.97 orders per year on average, annual ordering cost of XAF 379,473.32 and the same annual cycle-stock holding cost, giving XAF 758,946.64 in total before purchase cost. [S02]

The fractional order count is an annual average, not an instruction to place part of an order. If a feasible pack contains 100 units, the unconstrained quantity must be rounded and retested. At an illustrative 600 units, annual ordering cost is XAF 400,000 and holding cost is XAF 360,000, so total relevant cost is XAF 760,000. The difference from the unconstrained minimum is only XAF 1,053.36 per year, but feasibility still has to be established separately. [S02] [S03]

The curve also warns against false precision. In the same illustration, quantities of 500 and 800 units both produce annual relevant cost of XAF 780,000, only 2.77 percent above the minimum. This local insensitivity does not make every nearby quantity acceptable. It means that a feasible pack or negotiated delivery pattern may matter more than preserving decimal places, provided the costs and constraints outside the basic model remain unchanged. [S02]

The service boundary remains separate. A fixed EOQ can sit inside a continuous-review policy with a reorder point or a periodic process with scheduled reviews, but the reorder trigger must protect the relevant lead-time and review horizon. [S05] Nor

should the cycle-stock average be confused with total average inventory. Safety stock, goods in transit, quality holds, residual lots and work in progress require their own boundaries. EOQ chooses a lot under restrictive conditions; it does not validate the stock record or the customer promise.

6.2

Holding cost includes opportunity cost without duplicating charges

The annual holding charge h should represent costs that change because one more unit is held for one more year. It is not a standard percentage applied to every item, and it is not the accounting value of stock. A local estimate begins with the opportunity cost of cash committed to inventory, then adds stock-dependent storage, insurance, handling, deterioration, loss and obsolescence where those effects are not already represented elsewhere. [S02] [S09]

Each component needs a driver and an owner. Finance should state whether the capital charge reflects incremental borrowing, cash that could otherwise reduce another obligation, or an approved opportunity-cost convention. Stores should identify whether space and handling genuinely vary with average units held or only with receipts, pallets or a rented area. Quality and commercial roles should distinguish ordinary ageing from a specific revision or selling-window exposure. Fixed warehouse rent does not become incremental merely because it can be allocated per unit.

Order cost K needs the same discipline. Include the incremental work of supplier confirmation, document preparation, receipt, inspection and invoice processing when those activities change with the number of orders. Do not include the same inspection charge in both K and h. A per-unit freight charge belongs in landed unit cost, while a fixed charge per shipment may belong in K or in a separate freight-bracket model. The correct placement follows the cost driver, not the accounting ledger label. [S03] [S29]

For an illustrative decomposition consistent with the EOQ case, assume a unit cost of XAF 6,000 and an annual holding charge of XAF 1,200. The charge could comprise XAF 720 for capital, XAF 180 for stock-dependent space and handling, XAF 60 for insurance and XAF 240 for expected loss and obsolescence. These are teaching assumptions, not Cameroon tariffs or a recommended 20 percent rate. Their purpose is to make the components contestable before the total is used. [S02] [S09]

A terminal-loss scenario requires particular care. If the discount comparison already includes an expected residual loss for buying beyond the selling window, adding the same expected obsolescence to h would count it twice. One option is to keep routine, time-driven carrying exposure in h and model a discrete design revision or season-end residue separately. Another is to estimate the complete age-dependent loss curve and remove the overlapping allowance. The chosen boundary should be visible in the calculation record.

IAS 2 addresses a different question: inventories are measured at the lower of cost and net realisable value. [S09] That measurement does not supply the opportunity cost of cash, prove a future recovery price or convert a write-down into an annual holding

rate. Conversely, a decision-model carrying charge is not an accounting entry. Local reporting treatment and net realisable value require professional review; this chapter is not accounting, investment, legal or tax advice.

> **GFT NOTE**
>
> Approve a holding-cost worksheet before approving the lot formula. For every component, record its unit, driver, observation or estimate basis, owner and review date. Reconcile it to the order-cost and terminal-loss worksheets so no charge appears twice. Where no defensible local estimate exists, show low, central and high carrying scenarios and identify which one changes the selected feasible quantity. Do not conceal an unknown carrying charge inside a universal percentage. [S02] [S09]

6.3

A discount must beat its extra stock and cash cost

A quantity discount changes the purchase-cost term that the basic EOQ often holds constant. The buyer must therefore recompute the feasible choice rather than subtract the advertised saving from an existing order. In an all-units discount, the lower price applies to every unit once the relevant bracket is reached. In an incremental discount, it applies only to units inside that bracket. The two structures produce different average unit costs and different marginal incentives at the breakpoint. [S03]

A sound comparison evaluates each price bracket at its best feasible quantity and at any minimum quantity needed to obtain the bracket. It then applies pack multiples, minimum order quantities and other constraints before comparing totals. The lowest unconstrained cost inside a bracket can be infeasible, while the discounted breakpoint can create more stock than the demand horizon can absorb. The candidate set should include the un-discounted alternative so the analysis does not assume that accepting some discount is compulsory. [S03]

Annual total relevant cost can be written as purchase cost plus ordering, cycle-stock holding and decision-relevant logistics and residual costs. The comparison must use the same demand, service and time boundary for every candidate. If one option creates terminal stock, its expected recovery or disposal exposure belongs in that option, not in a footnote after the preferred answer has been selected. IAS 2 provides accounting context for net realisable value, but the scenario still needs a commercial recovery basis. [S03] [S09]

WORKED EXAMPLE: A LOWER UNIT PRICE LOSES THE ANNUAL COMPARISON

Consider an illustrative all-units offer for the same annual demand of 12,000 units. The baseline candidate orders 600 units at XAF 6,000 each. Its annual purchase cost is XAF 72,000,000, ordering cost XAF 400,000, cycle-stock holding cost XAF 360,000, freight XAF 1,200,000 and residual-loss allowance XAF 100,000. Total annual relevant cost is XAF 74,060,000. Every amount is a constructed teaching assumption, not a quotation or factory result. [S02] [S03]

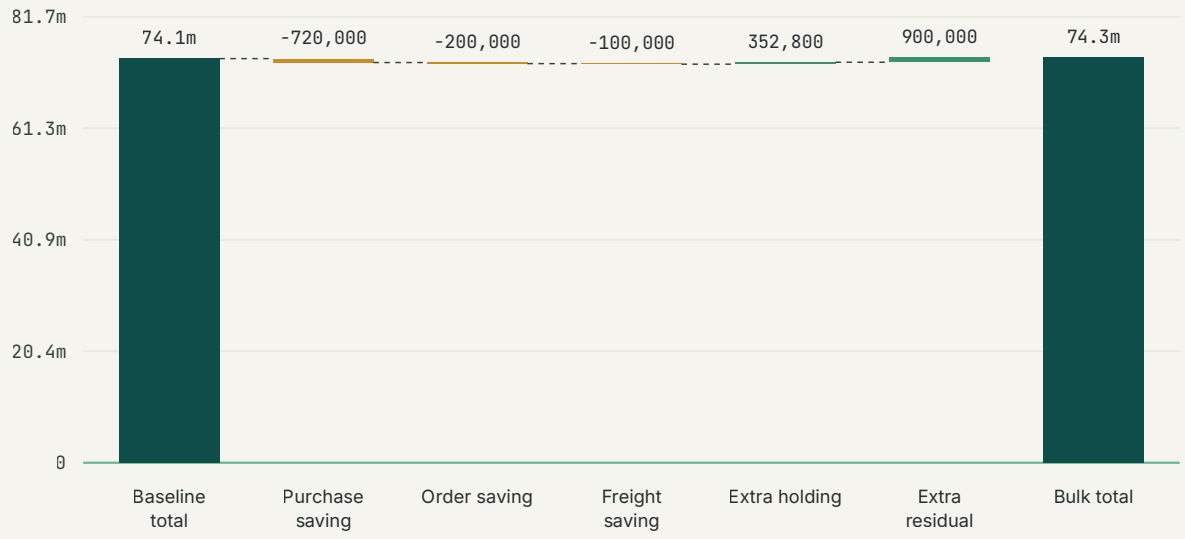
The bulk candidate orders 1,200 units at an illustrative price of XAF 5,940, a 1 per-cent reduction. Annual purchase cost becomes XAF 71,280,000, ordering cost XAF 200,000, holding cost XAF 712,800 using an illustrative charge of XAF 1,188 per unit-year, freight XAF 1,100,000 and residual-loss allowance XAF 1,000,000. Total annual relevant cost is XAF 74,292,800. The bulk candidate is therefore XAF 232,800 more expensive despite the lower price. [S03] [S09]

The annual comparison does not capture the whole approval problem. Each baseline order commits an illustrative XAF 3,600,000 at the purchase invoice, while each bulk order commits XAF 7,128,000, an additional XAF 3,528,000 before freight and other payments. If that cash is needed for another material or wage commitment, an annual cost saving alone would not make the larger lot feasible. Actual payment dates belong in the cash model rather than in this simplified annual arithmetic. [S03]

Exhibit 24, p. 98 reconciles the two illustrative totals. Savings from purchase price, order frequency and freight are not enough to cover the additional holding and residual allowances. The waterfall does not establish that a real discount will behave in the same way; it shows which evidence a real comparison must supply. [S03] [S09]

EXHIBIT 24 / ILLUSTRATIVE XAF PER YEAR

A bulk discount must pay for the stock it creates



ILLUSTRATIVE [S03] [S09] Constructed all-units discount case. Baseline 600-unit policy: purchase XAF 72,000,000 + ordering XAF 400,000 + holding XAF 360,000 + freight XAF 1,200,000 + residual allowance XAF 100,000 = XAF 74,060,000. Bulk 1,200-unit policy: purchase XAF 71,280,000 + ordering XAF 200,000 + holding XAF 712,800 + freight XAF 1,100,000 + residual allowance XAF 1,000,000 = XAF 74,292,800. Waterfall: 74,060,000 - 720,000 - 200,000 - 100,000 + 352,800 + 900,000 = 74,292,800. Residual allowances are decision-model teaching assumptions, not IAS 2 write-downs or realised losses.

The residual line deserves as much scrutiny as the price. It should be built from the quantity expected to remain, its usable condition, approved alternative demand, net recovery, handling and timing. A large allowance should not be accepted merely because it produces the desired rejection of bulk buying. Equally, a zero residual allowance is not neutral when the larger lot extends beyond a confirmed requirement. The honest analysis displays the quantity and recovery assumptions behind either value. [S06] [S09]

Price and cash must also retain their quotation boundaries. A usable comparison records supplier, item specification, quantity bracket, currency, quotation date, payment schedule, expected usable-receipt date and any expiry. Incoterms should state the rule edition and named place so carriage and related responsibilities are allocated consistently, while title and payment terms are checked separately. [S29] No inherited exchange rate or undated freight assumption should be presented as a current Cameroon purchasing input.

Finally, currency and customs assumptions belong in scenarios unless dated operational inputs support them. A lower foreign-currency price can coincide with a higher XAF commitment at the actual payment date, while a consolidated load can change clearance, damage or receipt costs. This chapter does not insert a tariff or conversion value that has not been opened and dated. The decision record should instead identify the missing quote and show which assumption would reverse the preferred quantity.

6.4

Time-phased demand needs time-phased lots

Constant annual demand is a poor description of a garment maker preparing discrete orders, style launches or institutional deliveries. A fabric requirement can be large in one production week, absent in the next and revised when a confirmed order changes. Material requirements planning explodes and nets those dependent requirements by date; fixed-horizon lot-sizing methods then decide which requirements to combine into receipts. [S04] [S07] An annual EOQ applied to the total can hide the very timing that determines whether stock is useful.

Lot-for-lot ordering places a receipt for each positive net requirement. It minimises cycle stock under a convention of immediate availability, but can create many ordering and receiving events. Fixed-coverage methods combine a chosen number of future periods. Finite-horizon optimisation instead evaluates alternative combinations of orders and holding across the stated schedule. The useful comparison holds demand, initial stock, order cost, holding convention and service constraints constant. [S04]

The following constructed schedule demonstrates that comparison. Demand over six periods is 100, 0, 180, 60, 240 and 80 units, totalling 660. Initial stock is zero, receipts occur at the start of a period, demand is consumed during it, and holding is charged on end-period stock at an illustrative XAF 100 per unit-period. Each order costs an illustrative XAF 20,000. Shortages are prohibited and no capacity, pack or cash constraint is applied in this first comparison. [S04]

Exhibit 25, p. 100 reports the resulting receipts, costs and terminal stock. The finite-horizon optimum was found by enumerating every possible next-order point in a dynamic programme. It happens to match the two-period coverage rule in this constructed case; that coincidence is not a general recommendation for two-period coverage. [S04]

EXHIBIT 25 / UNITS AND ILLUSTRATIVE XAF; SEE COLUMN HEADINGS

Time-phased methods produce different terminal stock

Method	Receipts by period	Terminal units	Order XAF	Holding XAF	Total XAF
Lot for lot	P1 100; P3 180; P4 60; P5 240; P6 80	0	100,000	0	100,000
Two-period cover	P1 100; P3 240; P5 320	0	60,000	14,000	74,000
Three-period cover	P1 280; P4 380	0	40,000	76,000	116,000
Fixed lot 300	P1 300; P4 300; P6 300	240	60,000	94,000	154,000
Finite-horizon optimum	P1 100; P3 240; P5 320	0	60,000	14,000	74,000

ILLUSTRATIVE [S04] [S06] [S07] Illustrative demand [100, 0, 180, 60, 240, 80], zero opening stock, K = XAF 20,000 and end-period holding XAF 100 per unit. Receipts occur at period start. Total = ordering + holding before purchase cost or terminal credit. Fixed lot 300 ends [200, 200, 20, 260, 20, 240]. Dynamic programming confirms the XAF 74,000 optimum. Terminal recovery follows S06 and S09.

The illustrative fixed-lot policy purchases 900 units to cover 660 units of demand and ends with 240 units. Its ordering plus holding cost cannot be compared fairly with the other rows without valuing that residue. A reusable standard input may retain useful value, while a revised logo fabric or incompatible shade lot may have little approved recovery. The single-period framework makes overage, salvage and disposal explicit, so terminal stock must enter the objective rather than disappear at the horizon. [S06] [S09]

The table also shows why fewer orders are not automatically better. Moving from lot-for-lot to three-period coverage reduces illustrative ordering cost from XAF 100,000 to XAF 40,000 but adds XAF 76,000 of holding. The best combination depends on when stock arrives and remains, not on order count alone. A method selected because it produces a tidy calendar can be more expensive than a less regular schedule. [S04]

Real planning adds net requirements and receipt status before this lot decision. Usable stock, valid allocations and firm dated receipts must be reconciled once, with quarantined or incompatible material excluded. The versioned bill of materials determines which requirement belongs to which garment order. [S07] Lot sizing should then operate on the remaining dated need, not on gross annual consumption or a style total that conceals different material specifications.

6.5

Capacity and pack sizes make feasible quantities discrete

An optimiser can propose a quantity that no supplier, store or cash account can accept. Feasibility must therefore be imposed explicitly, not left to a planner's silent adjustment after the calculation. The minimum useful constraint set includes supplier minimum order quantity, pack multiple, available purchase cash, storage capacity, maximum permitted material age and any customer or production date that limits useful consumption. [S03] [S07]

Rounding is a new candidate, not a cosmetic edit. If the unconstrained EOQ lies between packs, each adjacent pack should be checked against cash, space and total relevant cost. Rounding upward may buy unneeded exposure; rounding downward may increase order frequency or miss a discount bracket. A minimum lot can force the analysis to include a call-off, another supplier, a changed promise or no purchase, rather than pretending that the unconstrained answer survives. [S03] [S08]

Exhibit 26, p. 101 applies illustrative constraints to the earlier EOQ. The pack is 100 units, minimum order quantity 400 units, purchase-only cash capacity XAF 4,000,000 and storage capacity 650 units. Unit price remains an illustrative XAF 6,000. The unconstrained 632.46-unit result fails the pack rule, while a rounded 700-unit candidate fails both cash and storage. A value of one means that the stated gate is met; it does not mean that the order is approved. [S02] [S03]

EXHIBIT 26 / 1 = GATE MET; 0 = GATE FAILED

Cash and storage limits can rule out the unconstrained lot

	Pack	Cash	Storage	Combined
400	1	1	1	1
500	1	1	1	1
600	1	1	1	1
632 (EOQ)	0	1	1	0
700	1	0	0	0
800	1	0	0	0

Cell scale: 0 to 1

ILLUSTRATIVE [S02] [S03] Illustrative constraints: pack multiple 100 units, minimum order 400 units, unit price XAF 6,000, purchase-only cash ceiling XAF 4,000,000 and storage ceiling 650 units. The EOQ row uses 632 units for gate display; the unrounded EOQ is 632.455532. Purchase cash at rows 400 to 800 is XAF 2,400,000; 3,000,000; 3,600,000; 3,792,000; 4,200,000; and 4,800,000 respectively. Every row meets the minimum-order gate. Combined feasibility is the product of pack, cash and storage gates. No other cash use, freight or quality condition is included.

In this illustration, 600 units is the largest displayed candidate satisfying all three gates and costs XAF 3,600,000 per purchase invoice before other payments. Its annual ordering plus holding cost of XAF 760,000 was already calculated from the

EOQ case. That makes it an admissible candidate for further review, not an optimum: another quantity below 400 units is barred by the illustrative minimum, quantities between the displayed rows require enumeration, and terminal or service effects remain outside the grid. [S02] [S03]

The cash constraint must use payment dates rather than only invoice value. Deposits, balances, freight, duties where applicable and conversion costs can fall on different dates, while several material orders can compete for the same cash pool. WP13 owns financing structures and the paper's working-capital chapter maps the complete cycle. For lot sizing, the immediate requirement is simpler: reject or escalate a quantity whose dated payments breach an approved limit, and show the unmet material or service consequence rather than silently reducing another order. [S18] [S29]

Supplier capacity changes the receipt profile as well as the quantity. In a finite-replenishment model with constant production or supply rate `p` exceeding demand rate `d`, maximum cycle stock is `Q(1 - d/p)` and average cycle stock is `Q(1 - d/p)/2`. The corresponding benchmark is `Q* = √(2DK / [h(1 - d/p)])`. This is the economic-production-quantity extension, valid only under its own stable-rate and no-shortage assumptions. [S03] If `p` does not exceed `d`, that cycle-stock construction cannot be used.

A partial delivery is not merely a smaller instant receipt. Its accepted quantity and usable date determine when production can consume it, and observed supply can be censored when only a truncated delivery is visible. [S22] The lot record should therefore retain ordered, dispatched, received, accepted and still-open quantities with separate dates. Finite capacity, inspection and transport scenarios should be tested before the entire lot is credited to the period in which the first carton arrives.

Age and space constraints can also be directional. A store may physically accept a lot while leaving no workable location for inspection or another line-stopping component. A fabric may remain physically sound while exceeding the commercial window in which its shade or logo can be used. Feasibility therefore includes usable capacity and approved consumption time, not just cubic volume. Where the limits conflict with the requested service, the planner should expose the conflict for a sourcing or commercial decision rather than hide it in a larger safety allowance.

6.6

Fabric choices create hidden obsolete stock

An additional metre is not always an additional usable metre. Fabric can be tied to a construction, width, weight, colour, shade lot, print, logo and approved supplier. A revised specification can make a physically perfect roll unavailable for the next garment order. Material requirements planning depends on those item and version distinctions; aggregate stock value does not remove them. [S07]

Shade continuity creates a particularly important lot boundary. A factory may need enough material from one approved dye lot to complete a garment order without an unapproved variation. Buying a larger lot can support that continuity, but splitting

requirements across incompatible lots can leave remnants that cannot be combined merely because their total metres look sufficient. The useful measure is accepted quantity by approved lot and requirement, not total fabric alone. [S07]

Roll remnants require the same treatment. The remaining length on each roll, its inspection status and its suitability for the marker or later approved style determine whether it can be consumed. Measured process loss and yield belong to the dedicated process-yield analysis. This chapter does not apply a universal fabric yield or assume that every remnant has a recovery route. It requires the lot calculation to preserve those physical and commercial exclusions.

Consider a separate illustrative minimum-lot decision. A supplier will sell only 1,000 metres at XAF 5,940 per metre, while the confirmed requirement is 760 metres. The full invoice is XAF 5,940,000 and 240 metres have no confirmed use, representing XAF 1,425,600 of unassigned purchase commitment. Dividing the full outlay by confirmed use gives XAF 7,815.79 per confirmed metre before any residual recovery. This is a decision-exposure ratio, not an accounting unit cost or a claim that all remaining fabric is worthless. [S03] [S09]

That ratio can reverse the apparent discount even before annual holding is considered. It also shows why the recovery plan must name the actual later requirement. A generic assertion that the residue can be used someday supplies neither a date nor an approved specification. If an alternative style is genuinely approved, its demand, material compatibility, consumption date and recovery value should enter the scenario. If not, the residual remains an exposure rather than a fictitious asset with full value. [S06] [S09]

Design revisions and obsolete logos make terminal risk lumpy rather than smoothly proportional to time. Routine carrying cost may describe storage and capital, but a revision event can change the usable quantity immediately. The model should keep that discrete event separate from the annual holding rate, consistent with the no-duplication rule. IAS 2 net-realisable-value assessment may later recognise a write-down, but that accounting outcome should not be used to predetermine the commercial recovery assumption. [S09]

> **GFT NOTE**
>
> Require three approvals before buying beyond confirmed requirements. Material approval establishes which later products can lawfully and technically consume the residue. Commercial approval identifies the demand or recovery route and its date. Finance approval accepts the cash exposure and downside scenario. Record the unassigned quantity separately and do not count it as usable supply until those conditions are met. Supplier selection and any recovery arrangement must also respect adverse-impact due diligence rather than transferring the cost pressure invisibly. [S07] [S09] [S20]

6.7

Negotiating a smaller lot can change more than the price

The cheapest quoted unit is only one negotiable term. A buyer can instead discuss a smaller initial commitment, scheduled call-offs, split deliveries, reserved supplier capacity, a framework price or an agreed return route. These arrangements change order frequency, receipt timing, residual ownership and supplier risk as well as price. Supply-chain contract analysis therefore belongs inside the lot decision, because the terms allocate overage and underage exposure rather than making physical stock disappear. [S08]

A call-off agreement can separate a commercial volume commitment from the date on which each quantity must be received. Its value depends on what is binding: a non-binding forecast, a reserved capacity option and a firm call-off are different events. The contract should identify revision windows, minimum total quantity, price validity, cancellation liability, receipt acceptance and the party carrying unsold material. Without those definitions, an apparently flexible arrangement can defer the same bulk commitment rather than reduce it. [S08]

Split deliveries can reduce the peak receipt and improve material flow while adding transport, documentation, inspection or order-processing cost. The comparison should use the same total quantity and usable-receipt dates, then test whether each partial receipt arrives early enough for the consuming operation. A supplier's dispatch of a first instalment does not establish that the remainder is firm or accepted. Partial and censored supply observations must stay visible in the planning record. [S03] [S22]

Reserved capacity offers another boundary. The buyer may pay for the option to call material later, while the supplier holds capacity or an agreed input rather than manufacturing the full order immediately. That option has value only if the reserved resource can actually produce the approved item within the remaining requirement window. Its fee, expiry, conversion lead time and downside if demand fails should be compared with the holding and terminal cost of buying everything now. [S06] [S08]

Supplier-owned or consignment stock also requires a whole-chain view. Moving stock to another owner's balance sheet can release buyer cash without reducing physical quantity, ageing or disruption exposure. Returns rights and buybacks can redistribute residual loss, but their value depends on contractual scope, timing and counterparty performance. The analysis should state who owns each lot, who can consume it, who bears a revision and what happens when the agreement ends. [S08] [S09]

A smaller-lot request must also respect the supplier's cash and operating position. International Labour Organization research examines how purchasing practices affect wages and working conditions, based on particular surveys and garment-industry cases rather than a current Cameroon prevalence estimate. [S10] [S11] The practical implication for GFT is a negotiation safeguard: do not finance buyer flexibility through unagreed supplier credit, harmful working practices or an unrealistic delivery promise. Payment terms, change control and adverse-impact review belong beside the unit-price discussion. [S20]

Financing claims need equal precision. The opened IFC programme description concerns post-shipment finance based on acceptance of receivables by approved buyers. [S18] It does not establish funding for the initial deposit or pre-production material purchase, and it implies no eligibility for GFT or a proposed venture. WP13 owns the financing analysis. The lot-size model should show the unfunded interval and required cash approval rather than assuming that a later receivable will pay for today's fabric.

METHOD

GFT would attach one reproducible decision record to every material lot review:

- **Requirement:** item and specification version, location, dated net demand, confirmed and forecast components, usable opening stock and firm receipts. [S07]
- **Economics:** quotation date, currency, price brackets, landed-cost boundary, order-cost drivers, holding-cost components and terminal recovery scenarios. [S03] [S09] [S29]
- **Feasibility:** minimum lot, pack multiple, supplier and receipt capacity, storage, material-age limit, payment calendar and available cash. [S03] [S18]
- **Decision:** candidates evaluated, arithmetic and input vintage, selected quantity, rejected alternatives, service consequence, residual owner and approval role.
- **Review:** actual usable receipts, consumption, stock age, realised recovery, cash payments and the event that would trigger recalculation or stop another release. [S06] [S22]

The final choice should be presented as a small set of feasible policies, not a single decontextualised optimum. One candidate may minimise annual relevant cost but exceed cash; another may cost more per unit while reducing terminal exposure; a third may preserve an option through a call-off fee. Decision-makers need to see those trade-offs and the unmet promise if no candidate fits. GFT would release the lot only after planning, procurement and finance accept the same quantities, dates and exceptions.

This is also the limit of the chapter's evidence. The equations and cases establish a method for comparing lots, not a saving that any factory will realise. Fresh supplier quotes, local cost drivers, dated material requirements, usable-receipt observations

and cash limits must replace the teaching assumptions. Until then, the defensible output is a transparent decision record with explicit uncertainty, not an automatic purchase instruction.

Buy the quantity whose full commitment the business can use, finance and recover, not merely the quantity with the lowest unit price.

CHAPTER SOURCE NOTES

S02 MIT OpenCourseWare. 2006.

S03 MIT OpenCourseWare. 2006.

S07 MIT OpenCourseWare. 2006.

S05 MIT OpenCourseWare. 2006.

S09 IFRS Foundation. n.d..

S29 International Chamber of Commerce. 2020.

S06 MIT OpenCourseWare. 2006.

S04 MIT OpenCourseWare. 2006.

S08 MIT OpenCourseWare. 2006.

S18 International Finance Corporation. 2026.

S22 arXiv, Chen, Jiang, Zhang and Zhou. 2023-10-31.

S20 OECD. 2018-03-07.

S10 International Labour Organization. 2017-06-09.

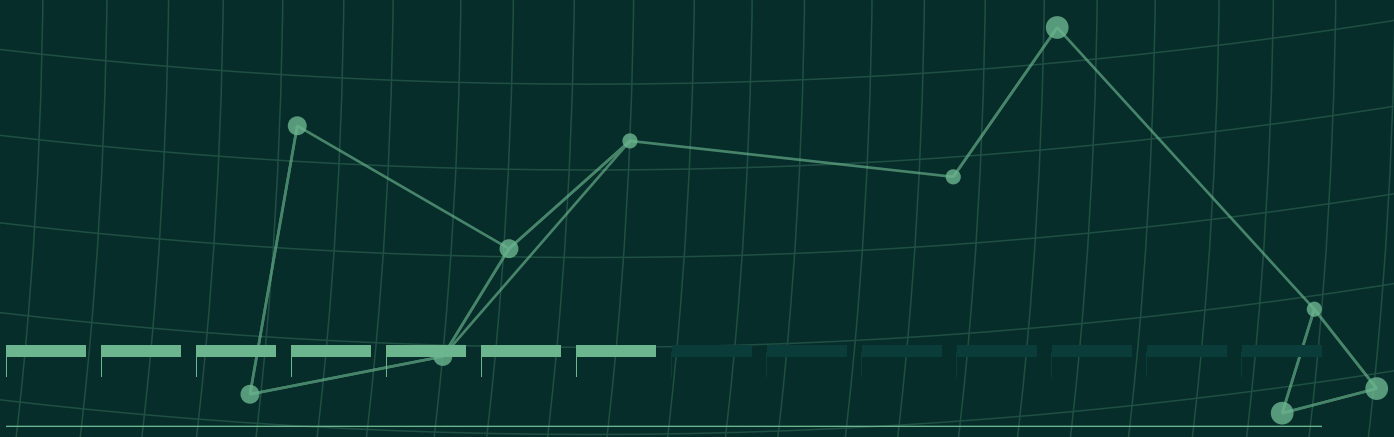
S11 International Labour Organization. 2017-05-15.

07

PART II: MATCH REPLENISHMENT RULES TO MATERIAL AND DEMAND / SEASONAL RISK

A seasonal buy needs a recovery plan before the first unit is ordered

Limited selling windows make overage, underage and replenishment opportunities central to the order decision.



CHAPTER 07
DECISION TEST

Limited selling windows make overage, underage and replenishment opportunities central to the order decision.

A seasonal buy should be approved only when its demand evidence, final useful replenishment date, variant allocation and residual recovery route are explicit. The central decision is not how to avoid every stockout or every leftover unit. It is how much economic exposure to accept before demand is known, and which later decisions can still reduce that exposure. Green Future Tech Limited (GFT), Hong Kong SAR, therefore treats the initial quantity, any contingent second buy and the terminal exit as one policy. The single-period model supplies the marginal logic, but it does not supply the demand distribution, prove a salvage price or make a late receipt useful. [S06] [S24] Every numerical case in this chapter is illustrative. No factory demand ledger, supplier quotation, returns history or realised recovery proceeds have been supplied, and no result below is a GFT operating claim.

7.1

The single-period model prices both mistakes

A finite selling window changes the meaning of an ordering error. In a repeat-stock setting, an unmet request may remain a backorder until the next receipt. In a seasonal setting, demand can disappear when the window closes, while a unit left behind cannot be rescued by the next ordinary replenishment cycle. The decision must therefore compare the incremental loss from ordering too little with the incremental loss from ordering too much, over the same commercial horizon. This is the purpose of the single-period, or newsvendor, model. [S06]

Define `p` as net selling revenue per unit after variable selling costs but before acquisition cost, `c` as incremental acquisition cost at the required location and condition, `s` as residual proceeds per unit before residual handling, `h` as handling or disposal cost per residual unit, and `g` as any incremental contractual cost of failing to supply a unit. Underage cost is `Cu = p - c + g`. Overage cost is `Co = c - s + h`. The residual term can include salvage revenue or disposal cost, as the opened MIT teaching material allows, but each component needs a dated commercial basis rather than a standard percentage. [S06]

This construction deliberately excludes costs that do not change with the order decision. A fixed design expense already incurred does not alter the marginal quantity. A capacity-reservation fee belongs only if the decision actually changes that fee. Selling revenue must also be net of any expedite, commission or packaging expense that occurs only when a sale is made. Otherwise the model can count the same expense in both `p` and `c`, or call a full selling price incremental contribution when it is not.

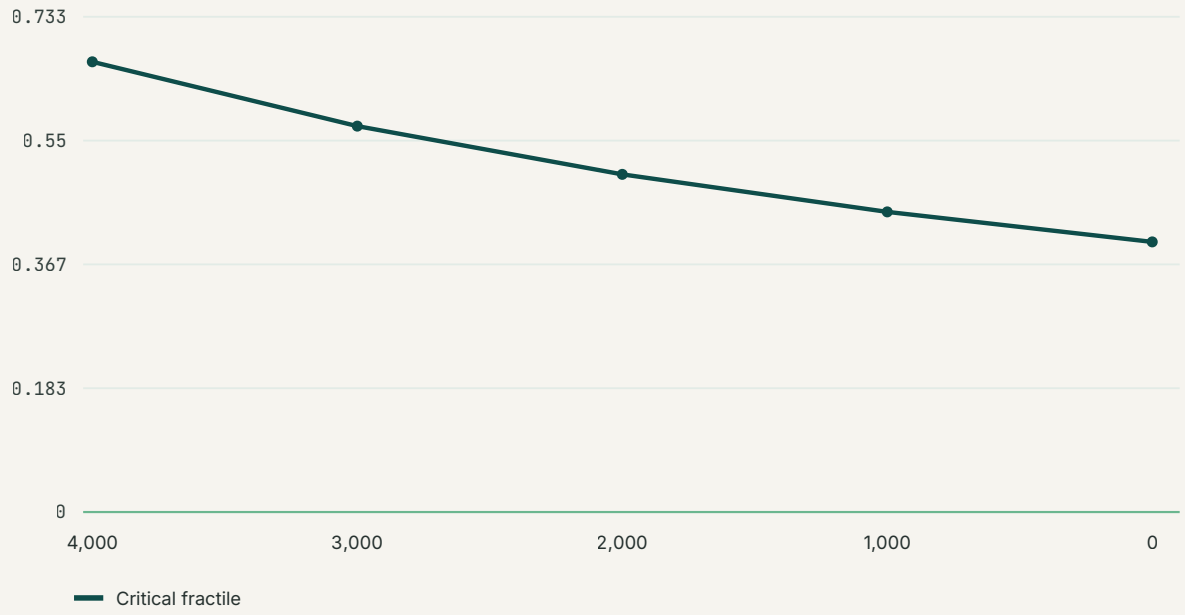
For a continuous demand distribution with cumulative distribution `F`, the critical fractile is `Cu / (Cu + Co)`. The order quantity satisfies `F(Q*) = Cu / (Cu + Co)` under the basic model. For discrete demand, choose the smallest feasible quantity whose cumulative probability reaches that ratio, then apply pack sizes and other constraints. The marginal interpretation is more useful than memorising the formula: an additional unit earns the underage margin when demand reaches it, but incurs the overage loss when demand stops before it. [S06]

The ratio is not a service promise chosen by a manager. It is an economic quantity implied by the cost model. A high critical fractile can justify a high probability of meeting demand when the lost contribution is large relative to residual loss. A low ratio indicates the opposite economic balance. Neither outcome is inherently virtuous. If the inputs imply a service level that conflicts with a contractual commitment, the business must change the commitment, the economics or the replenishment design, not conceal the conflict inside a forecast.

Exhibit 27, p. 109 keeps price and acquisition cost fixed in an illustrative case and lowers net residual proceeds. As recovery weakens, overage becomes more expensive and the critical fractile falls. The line is arithmetic from the stated inputs, not an observed relationship in an apparel market. [S06]

EXHIBIT 27 / CRITICAL FRACTILE AS A PROBABILITY

The critical fractile changes when residual recovery changes



ILLUSTRATIVE [S06] Illustrative arithmetic. Net selling revenue p is XAF 10,000, acquisition cost c is XAF 6,000, contractual shortage cost g is zero and residual handling h is zero. Thus Cu = 4,000. Labels are residual proceeds s in XAF, Co = c - s, and the plotted value is Cu/(Cu + Co), calculated before display rounding. These are teaching assumptions, not market prices or recommended service levels.

WORKED EXAMPLE: A DISCRETE SEASONAL ORDER

An illustrative garment has net selling revenue of XAF 10,000, acquisition cost of XAF 6,000, residual proceeds of XAF 4,000 and no additional shortage, handling or disposal cost. Therefore `Cu = 4,000`, `Co = 2,000` and the critical fractile is `4,000 / 6,000 = 2/3`. These are constructed teaching values in CFA franc (XAF), not a quotation or an estimate for a factory. [S06]

Assume demand has five possible outcomes. Demand of 60, 80, 100, 120 and 140 units carries illustrative probabilities of 0.10, 0.20, 0.35, 0.25 and 0.10 respectively. Cumulative probabilities are 0.10, 0.30, 0.65, 0.90 and 1.00. The first cumulative value reaching 2/3 occurs at 120 units, so the unconstrained single-period choice is `Q = 120`. The narrow gap between 0.65 and 0.666667 also shows why false precision in the input probabilities would be dangerous.

At `Q = 120`, illustrative expected demand is 101 units, expected sales are 99 units, expected residual stock is 21 units and expected unmet demand is 2 units. Expected revenue is XAF 990,000, acquisition cost is XAF 720,000 and residual proceeds are XAF 84,000, giving expected contribution of XAF 354,000 before any excluded fixed costs. The corresponding contributions are XAF 240,000 at 60 units, XAF 308,000 at 80 units, XAF 352,000 at 100 units and XAF 326,000 at 140 units. Arithmetic across the stated grid confirms the fractile choice; it does not validate the invented distribution or prices.

The basic result carries restrictive assumptions. It uses one order decision, a known demand distribution, linear incremental costs and no later replenishment. It also assumes that the act of ordering does not change demand, that residual recovery is available at the stated value and that every unit has the same economics. A size-specific buy, a supplier option, a cash ceiling or a markdown that responds to leftover stock requires an extended model. The critical fractile remains a useful benchmark, but it is not a universal optimiser.

Contractual shortage costs deserve particular restraint. A documented cancellation charge or service credit can enter `g` if the incremental exposure is genuinely triggered by the shortage. Speculative damage to goodwill should not be converted into a precise cash penalty merely to justify more stock. If management wishes to consider that effect, it should appear as a separate scenario with an explicit and challengeable basis. The same caution applies to terminal stock: IAS 2 addresses measurement at the lower of cost and net realisable value, which is an accounting question distinct from the incremental overage cost used here. [S09]

7.2

Demand evidence must include the days when stock was unavailable

Recorded sales are not necessarily recorded demand. When a requested variant is unavailable, the resulting sales history can show zero units even though customers requested the item, accepted a substitute or abandoned the order. Research on lost-sales inventory systems identifies this as censoring: only truncated demand and

supply values may be observable. [S22] A seasonal model fitted directly to those sales can therefore conclude that the missing variant has weak demand and repeat the shortage in the next buy.

The required record begins before the sale. For each request event, retain the item and variant requested, channel, event time, requested quantity, accepted order quantity, immediately fulfilled quantity, later fulfilled quantity, cancellation and substitution. Link the request to usable-stock exposure at that location and time. A customer who accepts another colour creates two related events: unmet demand for the original variant and fulfilled demand for the substitute. Combining them into one sale erases the allocation problem that the seasonal policy needs to solve.

Availability exposure needs an interval, not only a daily closing balance. A variant can begin the day in stock, sell out early and finish with a zero balance. Closing stock alone cannot distinguish a day with no requests from an afternoon of missed requests. The capture rule should record when usable availability became zero, when it resumed and whether the sales channel actually stopped accepting requests. Unknown exposure should remain unknown rather than being silently classified as no demand.

Cancellations also require separation. A customer may cancel because of delay, price, fit, a changed requirement or an offered substitute. Only some of those events represent demand that a larger seasonal buy would have retained. Treating every cancellation as a lost sale overstates latent demand just as treating every stockout day as zero demand understates it. The analysis should present bounds or scenarios when the reason cannot be established, rather than selecting whichever convention produces the more attractive purchase.

A practical demand table therefore distinguishes four quantities: requested demand, accepted demand, fulfilled demand and demand whose observation is censored. These are not interchangeable denominators. Unit fill rate concerns requested units filled immediately under a stated convention, while cycle service concerns whether a replenishment cycle experiences a stockout. [S05] For a seasonal policy, the commercial scorecard may also include accepted-order fulfilment by the promised date. That measure should not be relabelled as immediate fill.

> **GFT NOTE**
>
> Instrument the request-to-fulfilment path before increasing model complexity. At each selling location, GFT would retain requested variant, usable-availability interval, accepted quantity, substitution, cancellation reason and fulfilment event. Where the channel cannot capture a rejected request, mark the affected interval as censored and carry low and high demand scenarios into the order review. [S22] A modest order log with honest missing-data flags is more useful than an elegant forecast trained on a sales series whose stockout periods have been erased.

WP04 owns demand estimation and forecast uncertainty. This chapter consumes a predictive distribution or a set of demand scenarios available at the buying date. [S22] It should not refit a precise distribution from a few visible sales and pretend that censoring, promotions and price changes are absent. Each order record needs the fore-

cast vintage, observation window, stockout treatment and known commercial interventions. A later actual outcome must not be allowed to rewrite the distribution that was supposedly available when the initial order was approved.

Sparse evidence does not prevent a decision, but it changes the form of the answer. The planner can use explicitly illustrative low, central and high demand cases, calculate the exposure under each and show which assumption changes the chosen quantity. A single critical fractile applied to three invented distributions does not create statistical confidence. The value of the exercise is to reveal sensitivity and to identify which observations would most improve the next review.

Demand should also be retained at the variant level when sizes and colours are not interchangeable. Aggregating requests across a style can make total demand look well measured while concealing that the mix is unstable or observed only for variants that remained available. Where early-season sales will inform a second buy, the update must use the same request and exposure definitions as the initial forecast. Otherwise the apparent learning may simply reflect a change in data capture.

7.3

A second buy is valuable only if usable stock arrives in time

A contingent replenishment order is an option on remaining demand, not an automatic remedy for a cautious initial buy. Its value depends on what will be known at the second decision, how much unmet demand will remain, and whether the additional units can become usable before the commercial window closes. A supplier dispatch date is not the relevant endpoint. Production, carriage, receipt, inspection and quality release must finish early enough for the stock to meet demand. Stochastic-lead-time research cautions against replacing that chain with a fixed average. [S24]

Let t be the second-order decision date, L the realised usable-receipt lead time and T the last date on which an additional unit can usefully enter the selling or fulfilment process. The receipt contributes seasonal sales only in scenarios where $t + L$ is early enough relative to T, and where residual demand exceeds sales that existing stock and earlier receipts can satisfy. A receipt after T may still create cash exposure and terminal stock, but it should not be credited with full-season contribution.

The comparison must include the cost of preserving the option. A supplier may require a capacity-reservation payment, a minimum call-off, a price premium or an earlier commitment than the planner would prefer. Supply-chain contract analysis is relevant because these terms allocate overage and underage risk between the parties rather than eliminating it. [S08] The planner should compare the whole policy with and without the option, including the reservation fee and the likely residual outcomes, rather than treating the second order as free flexibility.

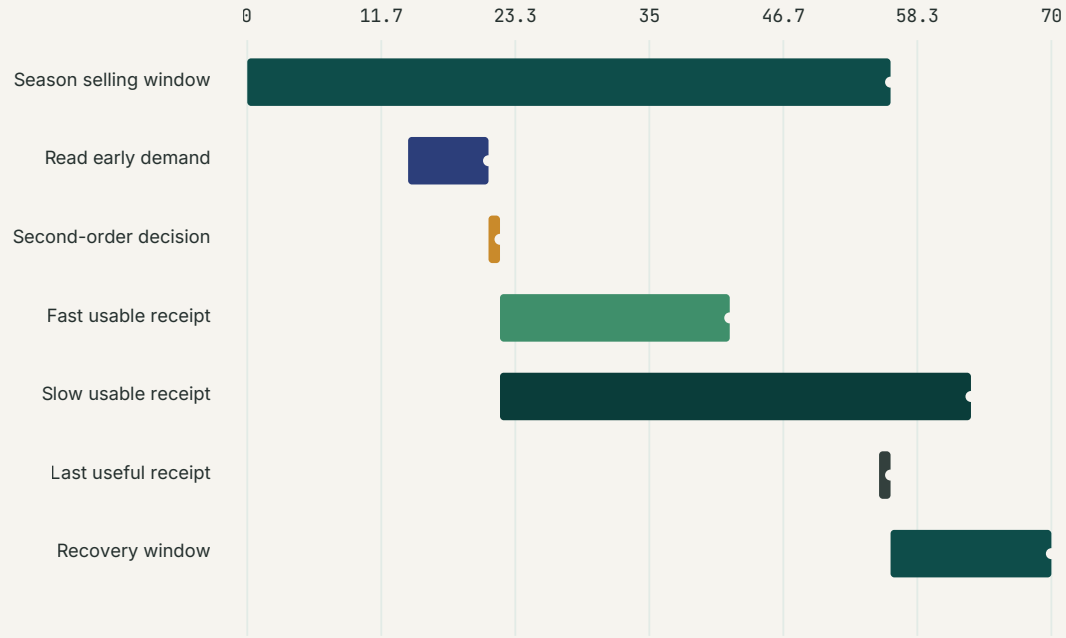
The initial and second quantities are connected. A larger initial buy reduces early shortages but also reduces the residual demand available to a later order and increases the stock that may need recovery. A smaller initial buy preserves learning but can sacrifice contribution before the second decision and can make the supplier option more valuable. The appropriate model therefore evaluates both decisions on

common demand and lead-time scenarios. Time-phased inventory methods provide the fixed-horizon discipline, while the single-period model supplies the overage and underage economics. [S04] [S06]

Exhibit 28, p. 113 shows an illustrative season with an early demand read, a second-order decision and two usable-receipt paths. The fast path leaves selling time. The slow path arrives after the last useful receipt and enters the recovery window instead. The dates are constructed to make the timing test visible, not to describe a supplier.

EXHIBIT 28 / DAYS FROM SEASON LAUNCH, ILLUSTRATIVE

A second order creates value only before the last useful receipt



ILLUSTRATIVE [S04] [S06] [S24] Illustrative planning geometry, not a supplier forecast. The season runs from day 0 to day 56. Early evidence is reviewed before a second-order decision at day 21. Both receipt scenarios begin after that decision. Fast usable receipt occurs at day 42, leaving selling time. Slow usable receipt occurs at day 63, after the last useful receipt at day 56, and therefore cannot be credited with ordinary seasonal sales. Durations are teaching assumptions.

In this illustration, the fast receipt at day 42 has until day 56 to contribute to seasonal sales. The slow receipt at day 63 has none of that window left. Counting both as equivalent supply would overstate the second option. The correct comparison uses scenario-specific sales, residual stock and cash. It also preserves the probability or scenario weight assigned to each lead-time path without presenting a sparse sample as a reliable distribution. [S24]

A usable option can lose value even when its average lead time appears short. Supplier congestion may increase precisely when strong demand triggers the second buy, making demand and lead time adverse together. The second-order contract should consequently state the reserved capacity, call-off deadline, expected usable-receipt stages, acceptance rules and treatment of partial delivery. If those terms are not known, the scenario model should include failure to supply rather than assuming the reserved quantity is guaranteed.

The decision record at the second review should reconcile remaining requested demand, usable stock by variant, firm inbound supply, outstanding customer commitments and the recovery exposure of another order. A positive demand signal alone is insufficient. Strong early sales may have exhausted the very sizes that the proposed replenishment pack does not contain, while apparently weak sales may reflect stockout censoring. [S22] The option should be exercised only when its incremental expected outcome clears the approved commercial and cash conditions after those reconciliations.

7.4

Sizes and colours create non-interchangeable residuals

A style total is not a sufficient inventory unit when the customer requests a particular size and colour. One variant can remain unsold while another misses demand, even when total stock exceeds total demand for the style. This is not a forecasting detail to be repaired after purchase. It changes the feasible order, the immediate fill measure and the residual recovery plan. Segmentation by product and customer requirement supports treating these commercial promises distinctly rather than imposing one aggregate rule. [S01]

The allocation model begins with a total style commitment and divides it across variants. Let q_v be the quantity assigned to variant v, and let D_v be its uncertain demand. The allocation maximises expected contribution subject to the total quantity, order packs, minimum lots, cash and any capacity limits. If substitutions are genuinely permitted, a separate matrix must identify which variant can satisfy which request, at what additional cost and with whose approval. Without that evidence, a transfer from one size or colour to another is not feasible supply.

The objective cannot be evaluated from style-level demand alone. Two allocations with the same total can produce different expected sales, residuals and customer outcomes because their mismatch occurs in different cells. The model should retain variant demand scenarios and, where relevant, their dependence. A common fashion shock can move several colours together, while a size-curve error can affect every colour in the same size. Independent variant distributions would miss those joint outcomes.

METHOD

For a seasonal variant allocation, GFT would test four records before approval:

- **Demand:** requested quantity and censored exposure by size, colour, channel and selling interval. [S22]
- **Feasibility:** pack quantities, production or purchase minimums, approved substitutions and any shared-material constraints. [S07]
- **Economics:** variant-specific net revenue, incremental acquisition cost, residual route and recovery expense, using the single-period boundary. [S06]
- **Service:** immediate variant fill, complete-order fulfilment and a separately defined size-curve availability measure. [S05]

Exhibit 29, p. 115 constructs a four-size, three-colour example. The heatmap shows unmet demand after matching stock to requests in each cell. Full demand and stock matrices appear in the notes so the arithmetic can be inspected. No substitution is permitted between cells. [S01] [S06]

EXHIBIT 29 / ILLUSTRATIVE UNITS OF UNMET DEMAND

A style can be overstocked while demanded sizes are missing



ILLUSTRATIVE [S01] [S06] Constructed allocation case, not observed sales. Rows are S, M, L and XL; columns are Indigo, Forest and Natural. Demand by row is S [18,14,10], M [24,22,16], L [20,18,12], XL [10,8,6], totalling 178 units. Stock by row is S [16,16,8], M [20,18,20], L [22,16,12], XL [12,6,14], totalling 180 units. Each plotted cell is max(demand - stock, 0), totalling 16 units of unmet demand. No substitution is allowed.

The illustrative style has 180 units of stock against 178 units of demand, yet 16 requested units cannot be filled in their requested variant. Expected sales under this realised matching are 162 units and residual stock is 18 units. Variant unit fill is therefore `162 / 178 = 91.0 percent`, despite aggregate stock exceeding aggregate demand by 2 units. These are teaching arithmetic values, not a service benchmark or a forecast. [S05]

The example also shows why an aggregate markdown can be a weak recovery instrument. The residual units sit in cells different from the unmet requests. Lowering the price of every variant sacrifices revenue on stock that could have sold without correcting the original mismatch, and it may still fail to move unwanted sizes or colours. A targeted commercial action needs the same cell-level record as the buying decision.

Complete size-curve availability should be defined rather than invoked. One possible GFT convention is the share of scheduled selling intervals in which every designated core variant has positive usable stock. That is an availability measure, not a fill rate: one remaining unit can make a cell available without satisfying subsequent demand. It should therefore sit beside variant unit fill and complete-order fulfilment, with its core range and review times stated in advance. [S05]

Early demand evidence can improve a second allocation only if it is read at the same level. A supplier prepack may force a different size curve from the one suggested by requests, and a production minimum may prevent small corrections in a particular colour. These are real feasibility constraints, not reasons to revert to a style average. If the feasible pack cannot address the observed shortage, the policy should compare a targeted alternative, a permitted substitution or no second buy rather than ordering more of a convenient but mismatched pack.

7.5

Postponed manufacturing and postponed planning are different

Delaying a decision is not the same as delaying physical differentiation. Postponed manufacturing keeps material in a more generic form until a better demand signal arrives, then converts it into a specified product. Postponed planning delays the commercial or production decision itself. The first can preserve feasible options. The second can simply move commitment closer to the deadline, increase urgency or transfer stock risk to another party.

The Japanese apparel study makes this distinction empirically important. In its survey of Japanese apparel wholesalers, postponed manufacturing was associated with lower inventory, while postponed planning increased inventory when firms externalised the retail function. The authors also identify retail controllability and the accuracy of additional-production decisions as conditions to consider. [S17] These are conditional associations from that sample, not a causal forecast for a Cameroon factory or a promise that every postponement programme reduces stock.

For an apparel maker, candidate manufacturing postponements include holding an approved common fabric, delaying a print or keeping a label generic until the order specification is confirmed. Each candidate requires a technical test. The base material must genuinely serve the later variants; conversion equipment and skills must be available; setup and changeover must fit the remaining window; and final inspection must still leave time for the promised delivery. Material planning and item-version controls provide the specification discipline, but they do not establish that a particular late-stage process is feasible. [S07]

The economic comparison should follow the decoupling point. Generic stock may serve several later demands, but conversion adds cost and can fail. Finished stock avoids that later conversion risk while carrying more variant-specific residual exposure. Postponement is worthwhile only when the expected reduction in mismatch and terminal loss exceeds the additional conversion, reservation, coordination and service costs under common demand scenarios. No universal inventory reduction should be inserted into that comparison.

A delayed planning decision needs an equally explicit owner and information rule. If a retailer, distributor or customer controls the signal used for additional production, the factory must know whether that signal is a forecast, a reservation or a binding order. Otherwise planning postponement can produce repeated speculative commitments while each party believes another owns the residual. Contract definitions and call-off liabilities should be settled before the season, not during an urgent second buy. [S08] [S17]

GFT would therefore approve postponement as a timed operating design rather than a general principle. The record would name the generic stock object, permissible later variants, latest conversion start, usable-completion deadline, responsible approval and fallback route. WP20 owns the broader process and product-design choices. This chapter retains the narrower inventory test: the option must remain technically convertible and commercially useful after the demand information arrives. [S07] [S17]

7.6

Markdowns and returns change the residual-stock calculation

Residual recovery is a process with time, cost and uncertain eligibility, not a single salvage rate. A markdown can accelerate sales but changes revenue and can alter demand for other variants. A customer return does not immediately become usable stock. It must be received, inspected, graded and, where necessary, repaired, cleaned or repacked before it can re-enter the selling process. European Environment Agency evidence describes complex return handling and potentially significant effects on resale timing and value. [S26] Its European return rates and destruction estimates must not be transferred to Cameroon.

The recovery calculation should follow each route separately. Expected net resale recovery is the expected selling proceeds less incremental discount support, handling, repair, repacking, channel and disposal costs, with the timing and probability of each outcome stated. An item that misses the remaining selling window may have a different route from an item returned early. A damaged return may have no resale route even though an undamaged return can be restocked. Combining these conditions into one average salvage price hides the decisions that create the value.

Returns also require a stock-state control. Until inspection establishes condition and ownership, the returned unit should not inflate promiseable stock. If it becomes eligible for resale, its release time belongs in the remaining-window calculation. If it is repaired, the repair capacity and additional cost belong in the same comparison. A return may reverse revenue commercially while still failing to restore a unit that can be sold again in the season.

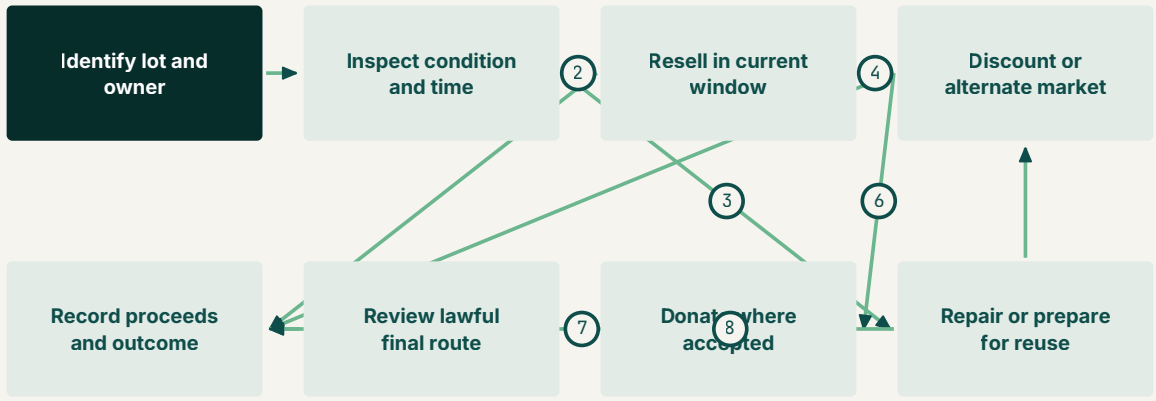
Accounting and decision economics remain distinct. IAS 2 describes inventories as measured at the lower of cost and net realisable value, with net realisable value reflecting estimated ordinary selling price less completion and necessary selling costs. [S09] That measurement does not prove that a buyer exists or that the estimated proceeds will be realised. The seasonal order model needs expected incremental outcomes, while the financial record needs the applicable accounting assessment. Local professional review is required; this chapter is not accounting or legal advice.

An EU-facing scenario adds a scope test. The European Commission states that, from 19 July 2026, large companies across the EU are prohibited from destroying unsold clothes, clothing accessories and footwear, subject to specified circumstances and exemptions. [S21] This is not a blanket obligation on a Cameroon manufacturer. The entity, company size, market, product and contractual role must be established by legal review, with WP17 providing the detailed interpretation.

The Commission's explanation prioritises selling through discounts or alternative markets, donation, and preparation for reuse through repair, refurbishment or remanufacturing, while restricting destruction to specified cases. [S21] **Exhibit 30**, p. 118 converts that principle into an illustrative factory decision flow. The flow is a proposed control design, not a claim that every route is locally available or legally sufficient.

EXHIBIT 30 / PROPOSED DECISION SEQUENCE

Residual goods require an approved recovery route



2. eligible and timely
3. recoverable
4. unsold
6. unsold and suitable
7. not accepted
8. not feasible

ILLUSTRATIVE [S09] [S21] [S26] Illustrative GFT control flow informed by IAS 2 measurement concepts, European returns evidence and the Commission's EU-facing recovery hierarchy. Actual route availability, net realisable value, ownership, donation acceptance, waste requirements and legal scope need local confirmation. No recovery rate or realised value is asserted.

7.7

A seasonal stop rule protects cash from wishful demand

The stop rule should be written before the initial order, while the cost of another unit can still be assessed dispassionately. It should name the last order date, last useful receipt, demand-evidence refresh, residual exposure limit and role authorised to approve an exception. Without those commitments, every new sales narrative can justify another buy after the feasible selling window has narrowed. The single-period model supplies the marginal economics, while inventory measurement and recovery records supply the terminal exposure. [S06] [S09]

At each review, the planner should project terminal stock under declared demand and usable-receipt scenarios. The exposure includes current usable stock, firm inbound supply and any proposed order, less sales expected before the window closes. Variant-level residuals must remain visible. The review then compares expected net recovery, handling cost and cash committed with the approved limit. A statistically attractive reorder that breaches the limit requires an explicit business decision rather than an automatic exception.

The action ladder begins with preventing additional mismatch. Stop or reduce uncommitted orders where the marginal case fails, convert only where specifications and capacity permit, and do not count an unreliable cancellation right as certain relief. Commercial actions can then target price, channel, repair, resale, donation or reuse according to actual eligibility. Final disposal requires documented ownership and lawful route review. The appropriate sequence depends on time and net recovery, not on a universal preference for markdowns. [S21] [S26]

A stop rule must also avoid the opposite error. Ending all replenishment on a fixed calendar date can sacrifice profitable demand when a fast, qualified option still clears its economic and cash tests. The rule should therefore stop unsupported exposure, not every later order. A second buy may proceed when fresh requested-demand evidence, usable-receipt timing, variant fit and residual exposure all remain within the approved conditions. Governance should make both continuation and cessation explainable.

> **GFT NOTE**
>
> Attach a one-page exit record to every seasonal purchase approval. Include the initial quantity and variant allocation, demand vintage, option terms, last useful receipt, recovery routes, maximum approved residual exposure and named owner roles. At each review, record actual requests, sales, returns, usable receipts, markdowns, residual units and net recovery. Close the record only after proceeds and final outcomes reconcile to the physical stock. This is a proposed GFT control, not an implemented system or a universal threshold.

The final evaluation should compare the original assumptions with realised outcomes without rewriting history. Which demand scenarios were available at approval, which receipts were expected to be usable, and which recovery values were authorised? Actual contribution, terminal stock, cash recovery and service should be reported

against that record. A high sell-through achieved by early stockouts is not unqualified success, and a low residual achieved through indiscriminate discounting may conceal lost margin. The policy earns renewal when its complete seasonal result is understood.

A seasonal order is not complete until its exit is priced, timed and owned.

CHAPTER SOURCE NOTES

S06 MIT OpenCourseWare. 2006.

S24 arXiv, Michna, Nielsen and Nielsen. 2015-03-20.

S09 IFRS Foundation. n.d..

S22 arXiv, Chen, Jiang, Zhang and Zhou. 2023-10-31.

S05 MIT OpenCourseWare. 2006.

S08 MIT OpenCourseWare. 2006.

S04 MIT OpenCourseWare. 2006.

S01 MIT OpenCourseWare. 2006.

S07 MIT OpenCourseWare. 2006.

S17 Japan Marketing Academy, Quarterly Journal of Marketing. 2020-03-31.

S26 European Environment Agency. 2024.

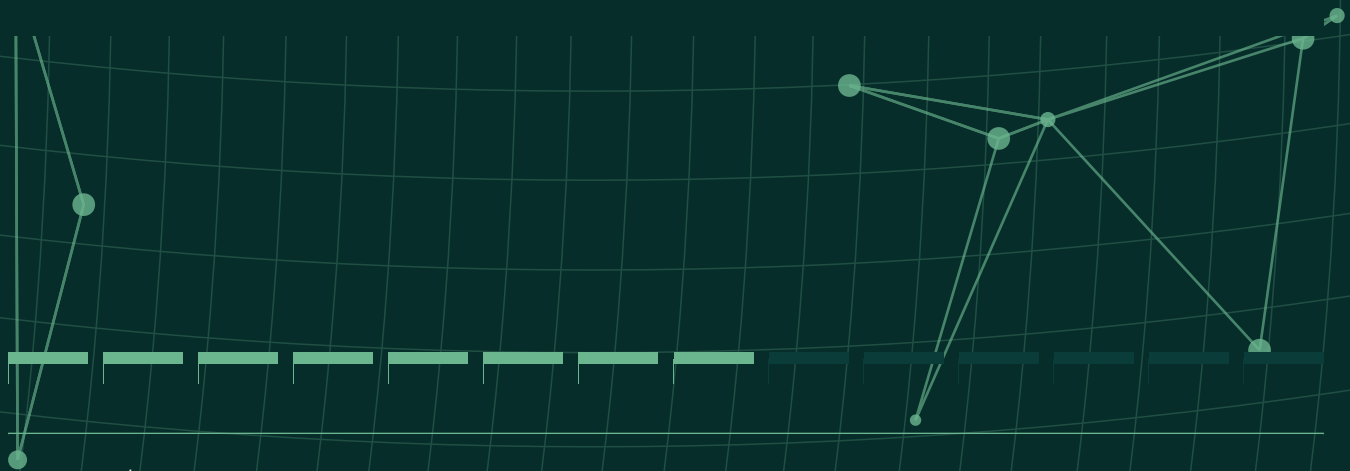
S21 European Commission, Directorate-General for Environment. 2026-07-17.

08

PART III: RESOLVE CONSTRAINTS ACROSS THE SUPPLY CHAIN / MATERIALS AND FLOW

A material plan is feasible only when every component and process can meet the date

Inventory optimisation must respect dependent material requirements, quality release and finite conversion capacity.



CHAPTER 08
DECISION TEST

Inventory optimisation must respect dependent material requirements, quality release and finite conversion capacity.

A material plan is feasible only when the complete approved kit becomes usable before each consuming operation and the required conversion work fits the available process time. A calculated requirement is not yet a delivery promise. Green Future Tech Limited (GFT), Hong Kong SAR, therefore separates four tests: explode the correct order through the correct bill of materials, net only supply that can actually be consumed, back-schedule from the operation date, and reconcile the result with finite capacity. Material requirements planning supplies the dependent-demand logic, while finite-buffer and flow analysis expose what an unconstrained calculation can miss. [S07] [S19] [S28]

This chapter applies that sequence to garment production. It treats quality release, specification version, shade compatibility, packaging and process capability as release conditions rather than administrative detail. Every numerical case is illustrative and constructed to show arithmetic. No factory ledger, routing, supplier quotation or measured line performance has been supplied. GFT's recommendation is a complete-kit gate with recorded exceptions, not a claim that such a gate has already produced operating results.

8.1

The bill of materials converts orders into dependent demand

A confirmed garment order creates independent demand for the parent item and dependent demand for its components. The planner should not forecast every zip, label and metre of thread as though each were unrelated to the order. Material requirements planning instead derives component requirements from the parent schedule, product structure and available supply. [S07] For an apparel order, that product structure must connect the ordered style, colour and size set to the exact fabric, trims, thread, labels and packaging permitted for production.

The bill of materials is therefore a technical control, not merely a purchasing list. Fabric identity may require construction, width, weight, colour and approved shade-lot rules. A zip may require length, teeth, slider and approved supplier. A label may carry a customer, language or care-instruction revision. GFT would retain these attributes in the versioned specification and treat substitution as an engineering approval. Two items with similar names cannot be netted against one another merely because the warehouse has both.

The explosion begins with accepted parent demand. For parent order `i` requiring quantity `Q_i` in period `t`, and component `c` used at `q_ic` per parent unit, the base dependent requirement is `G_ict = Q_i × q_ic`. This expression produces a gross requirement before losses, stock and receipts. If the parent is a pack or a size assortment, the calculation must first preserve that assortment and then apply the correct component usage to each variant. A single average garment can conceal different zip lengths, fabric consumption or labels across the order. [S07]

Multi-level structures require the same discipline at each stage. A finished garment may consume a cut bundle, which consumes fabric and an approved marker; a finished pack may consume garments, a polybag and a carton. The timing of the child requirement follows the operation that consumes it, not automatically the final ship-

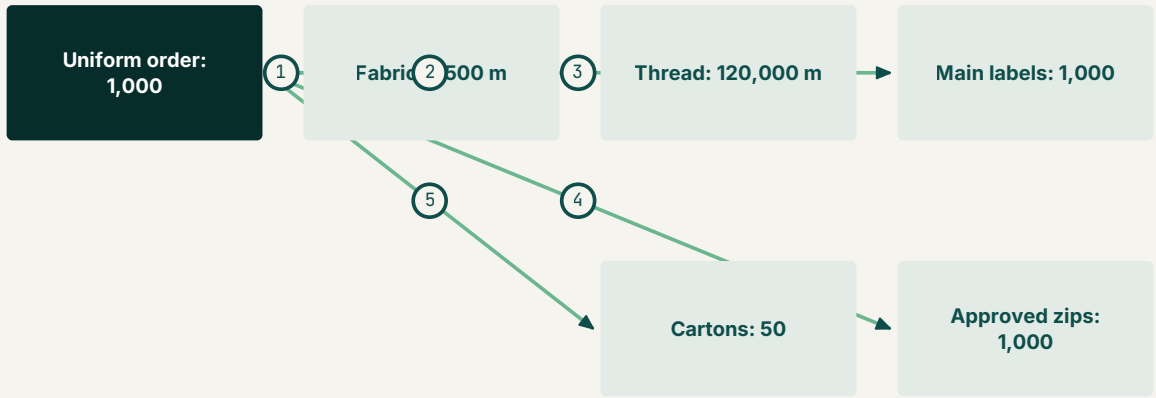
ment date. If an intermediate item is made to stock, its inventory record needs its own status and lead time. Flattening every level into one undated shopping list would destroy precisely the information needed to release orders.

Demand streams also need explicit reconciliation. A customer order, a forecast for possible future orders and a replenishment order already released to a supplier are different objects. If a confirmed order consumes part of the forecast, the remaining forecast must be reduced according to a stated consumption rule. Otherwise the same expected demand enters the plan twice. Likewise, a component purchase order is supply against an existing requirement, not a reason to create the requirement again. GFT would link each parent requirement, allocation and replenishment order through stable identifiers so the planner can trace why each quantity exists. [S07]

The illustrative explosion in **Exhibit 31**, p. 123 uses one uniform order to show these links. Quantities are teaching assumptions, not a product specification or consumption benchmark. The method follows bill-of-material explosion; only the arithmetic is GFT's. [S07]

EXHIBIT 31 / ILLUSTRATIVE BASE QUANTITIES FOR 1,000 GARMENTS

Dependent demand links one garment order to every required component



1. 1.5 m each
2. 120 m each
3. 1 each
4. 1 each
5. 20 per carton

ILLUSTRATIVE [S07] Illustrative bill-of-material explosion. Assumes 1,000 accepted garments, 1.5 metres of fabric and 120 metres of thread per garment, one main label and one zip per garment, and 20 garments per carton. Base quantities are order quantity multiplied by usage per garment, before scrap, stock, receipts and capacity checks. These are not observed consumption rates or an approved garment specification.

The exhibit deliberately stops at base requirements. It does not say how much to buy. Fabric may be waiting for shade approval, zips may arrive after cutting, and a supplier pack may exceed the calculated balance. Nor does the chart establish that the sewing

line can convert the order by the required date. Dependent-demand explosion answers what the parent order calls for. Feasibility requires the subsequent netting, timing and capacity tests.

Version control closes the first test. The explosion record should carry the order revision, bill-of-material revision, size-colour breakdown, approved substitutions and effective dates. A late logo change can make an otherwise adequate label stock unusable, while a marker revision can alter fabric requirement. The planner should not overwrite the old structure silently. Retaining both versions shows whether a shortage arose from physical supply or from a technical change after purchase commitment.

8.2

Net requirements must recognise what can actually be consumed

Netting converts gross dependent demand into an additional requirement. GFT uses the conditional expression `N_ct = max(0, G_ct + A_ct - U_ct - F_ct)`, where `G` is the base requirement, `A` is a declared material allowance, `U` is accepted usable supply allocated to this requirement, and `F` is the qualifying firm receipt expected to become usable before the consuming operation. The formula is a proposed control convention built from material requirements planning, not a substitute for item-specific ledger rules. [S07]

Every term needs the same item identity, unit and date boundary. `U` cannot include quarantine, rejected stock, an unapproved substitute or a shade lot that the order cannot use. It also cannot include material already allocated to another requirement. `F` is not every open purchase order. Its quantity, specification, commitment status and expected usable-receipt date must satisfy the planning rule. A receipt due after cutting may help a later order, but it does not remove this order's shortage.

Receiving evidence supports that caution. Project Zipper examined brand-owner and retailer ship and receive data, providing a study-specific example of why order, shipment and receipt records require reconciliation. [S31] Its results do not calibrate a garment factory's receipt reliability, and this chapter uses none of its reported error rates. The operational lesson is narrower: keep purchase quantity, dispatched quantity, physically received quantity and accepted quantity as separate fields until their differences are resolved.

Allowances need equally visible definitions. If cutting loss is measured as a share of issued fabric, the required issue is base requirement divided by one minus that share. If the policy instead adds an allowance expressed as a share of base requirement, it multiplies base by one plus that allowance. These are different models. Applying both would count the same expected loss twice. Rework adds another distinction: it can consume machine and labour time without requiring the same additional material as first-pass production. The material plan and capacity plan should therefore carry linked but separate yield assumptions.

WORKED EXAMPLE: NET THE KIT BEFORE RELEASING IT

An illustrative order for 1,000 garments uses the base quantities in **Exhibit 31**, p. 123. Assume an additive fabric allowance of 4 percent of base and a thread allowance of 2 percent of base. Labels and zips each carry a 1 percent additive allowance. These are constructed assumptions, not measured factory losses. The resulting requirements are 1,560 metres of fabric, 122,400 metres of thread, 1,010 labels, 1,010 zips and 50 cartons. [S07]

Also assume 900 accepted metres of fabric and a further 400 metres with a qualifying usable date. For zips, 600 are accepted and 300 are qualifying receipts, while another 200 remain quarantined and are excluded. Accepted usable supply is 80,000 metres of thread plus a 30,000-metre qualifying receipt, 1,010 labels and 60 cartons. Net additional requirements are therefore 260 metres of fabric, 12,400 metres of thread, no labels, 110 zips and no cartons. The arithmetic demonstrates the netting rule only. [S07]

Exhibit 32, p. 125 preserves the worked inputs and outputs in one boundary. Negative balances are floored at zero because the model calculates additional need for this order, not an automatic transfer of surplus to another order.

EXHIBIT 32 / ILLUSTRATIVE COMPONENT UNITS SHOWN IN EACH ROW

Usable supply, not book stock, determines the additional requirement

Component and unit	Base	Allowance	Usable allocation	Firm receipt	Net need
Fabric, m	1,500	60	900	400	260
Thread, m	120,000	2,400	80,000	30,000	12,400
Main labels, pieces	1,000	10	1,010	0	0
Approved zips, pieces	1,000	10	600	300	110
Cartons, pieces	50	0	60	0	0

ILLUSTRATIVE [S07] [S31] Illustrative arithmetic for 1,000 garments. Net requirement = max(0, base + allowance - accepted usable allocation - qualifying firm receipt). Fabric: 1,500 + 60 - 900 - 400 = 260 m. Thread: 120,000 + 2,400 - 80,000 - 30,000 = 12,400 m. Labels: 1,000 + 10 - 1,010 = 0. Zips: 1,000 + 10 - 600 - 300 = 110, with a further 200 quarantined zips excluded. Cartons: 50 - 60 = 0 before the zero floor. S07 supplies the MRP method context; S31 supports ship-receive reconciliation, not any numerical input.

The worked kit is not complete. Surplus labels and cartons cannot compensate for missing fabric, thread or approved zips. This is why aggregate stock value is a weak release test. The planner must resolve each component exception or approve a feasible redesign. Releasing everything else into production would create work in progress that still cannot become a finished garment.

Across several orders, allocation becomes a priority decision. The same accepted fabric roll may satisfy more than one requirement, but not all of them at once. GFT would allocate against an approved order priority and retain the residual exception list rather than letting whichever order is entered first consume supply invisibly. When

priorities change, the system should recalculate every affected kit and show which previously covered order has become short. Netting is complete only when that inter-order consequence is visible.

8.3

Back-scheduling reveals late material before the line is loaded

A net requirement becomes actionable only when it carries a need date. Back-scheduling starts with the customer or stock-replenishment due date, places the routing operations in sequence and works backwards through their durations. Each component must be accepted and available before the operation that consumes it. Fabric and approved marker need to precede cutting; sewing thread and approved trims need to precede the relevant assembly steps; final labels and packaging need to precede finishing and packing. Time-phased material planning connects those requirements to planned order receipts and releases. [S04] [S07]

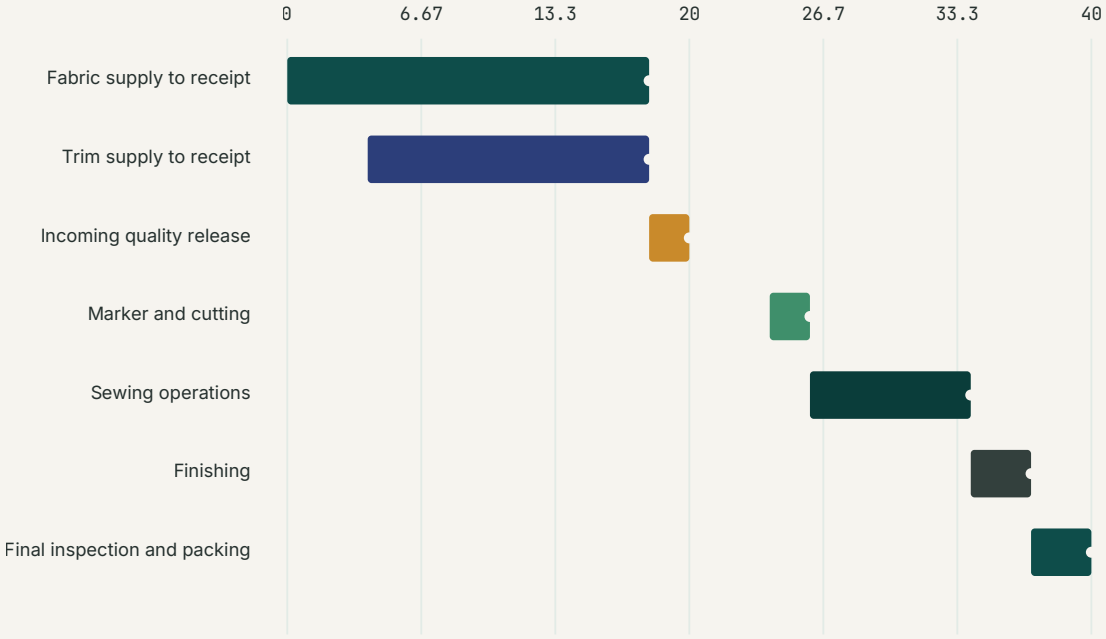
The clock must end at usable availability, not supplier dispatch. It includes order preparation, supplier processing, transport, receipt, inspection and quality release where those activities apply. The durations should come from measured scenarios or explicitly marked estimates. If the supplier quotes production time but not inspection and release, the planner cannot treat the quotation as the full material-ready interval. A component arriving physically early can still be late to the operation if its quality disposition remains open.

The scheduling calculation is straightforward but the calendar is not. For an operation with material need date `D`, a component with usable-receipt lead time `L` requires effective release no later than `D - L`, after converting both to a consistent working or calendar-day basis. Holidays, planned maintenance and non-working shifts can make those calendars differ. The record should show which calendar was used and who approved any compression. Subtracting an average duration without displaying these assumptions produces apparent precision rather than a defensible date.

Exhibit 33, p. 127 constructs a feasible-looking sequence for the same teaching order. Day 40 is an arbitrary due-date marker, not a promised delivery duration. Supplier and process times are illustrative; the exhibit does not calculate capacity or prove that the route exists. [S07]

EXHIBIT 33 / ILLUSTRATIVE PLANNING DAYS, WITH DAY 40 AS THE DUE MARKER

Material-ready dates must precede their consuming operations



ILLUSTRATIVE [S07] Illustrative backward schedule with constructed durations. Fabric and trims must complete quality release before cutting. Sewing follows cutting, finishing follows sewing, and final inspection and packing precede the day-40 due marker. The sequence shows timing logic only. It does not establish supplier lead times, resource availability, process capability or a real customer promise.

The chart exposes a useful slack interval between quality release and cutting, but slack is not proof of feasibility. Another order may occupy the cutting slot, the fabric receipt may slip, or inspection may reject a shade lot. Back-scheduling should therefore retain both the latest acceptable dates and the current planned dates. The difference between them tells the planner where a delay can be absorbed and where it immediately threatens the promise.

A calculated release date in the past is an exception, not an instruction to rewrite history. The planning engine should flag it and preserve the original need date. The planner can then evaluate an approved substitute, a later customer date, a different feasible route or an explicitly costed expedite. Automatically moving the planned release to today while leaving the due date untouched would hide the very infeasibility that the calculation revealed. Time-varying planning requires the shortage to remain visible in the period where action is needed. [S04] [S07]

Rescheduling also needs restraint. Every small forecast or receipt revision should not generate a cascade of unreviewed order changes. GFT would classify exceptions by their effect on complete kits and consuming operations, then route them to a named role. A receipt delay that remains inside approved slack has a different consequence from one that misses cutting. The review record should retain the old date, new date, cause, decision and affected parent orders so planners can distinguish a genuine material problem from unstable parameters.

8.4

Minimum lots and capacity can make an MRP plan infeasible

The net requirement is an unconstrained quantity until supplier and factory constraints are applied. A supplier may require a minimum order, sell in fixed packs or limit production within the required window. The factory may lack the machine hours, approved process capability or inspection capacity needed to convert the kit. EOQ extensions address lot and replenishment trade-offs, while MRP time-phases requirements; neither fact makes every calculated quantity and date physically achievable. [S03] [S07]

Order rounding comes first. In the illustrative net kit, a 300-metre fabric pack would turn the 260-metre net need into a 300-metre purchase and leave 40 metres before later consumption. A 200-piece zip minimum would turn the 110-piece need into 200 pieces and leave 90 pieces. If thread comes in 5,000-metre order multiples, the 12,400-metre need would round to 15,000 metres and leave 2,600 metres. These residuals are arithmetic consequences of teaching assumptions, not recommended purchase quantities. They need an owner, an age and a credible route into later compatible orders. [S03]

Capacity reconciliation then asks whether the acquired kit can be converted in the planning window. For each process, required work is the sum of order quantity multiplied by the applicable work content, plus explicitly modelled setup and rework. Available work time is the approved calendar time after breaks, planned maintenance and existing commitments, using the same process boundary. Detailed line balancing belongs to WP15. The material-plan test here is coarser but essential: do not release a customer promise whose aggregate process load already exceeds feasible hours. [S07]

Consider an illustrative capacity check for the 1,000-garment order. Assume cutting work content of 0.8 minutes per garment, sewing of 24 minutes and finishing of 4 minutes, with setup and ordinary allowances already included. Required time is therefore 800 cutting minutes, 24,000 sewing minutes and 4,000 finishing minutes. After other commitments, suppose the window offers 720 cutting minutes, 22,800 sewing minutes and 4,800 finishing minutes. Cutting is short by 80 minutes, sewing by 1,200 minutes and finishing has 800 minutes unallocated. The material kit could be numerically complete while the order remains infeasible in its current slot. [S07] [S19]

That result calls for a decision, not an average utilisation percentage. Cutting and sewing cannot borrow finishing's spare minutes unless a genuinely qualified alternative process exists. Possible responses include moving to another approved slot, splitting the order across feasible deliveries, using an approved alternate resource, resequencing other orders or renegotiating the date. Overtime is a separate operating and labour decision, not free capacity that the planning engine may insert automatically. Each option changes cost, other orders or risk and should appear in the approval record.

Lot size also interacts with flow. A larger batch may reduce repeated setup, but it can occupy a process for longer and delay other complete kits. Splitting a batch may improve delivery sequencing while adding setups and handling. Time-phased lot-sizing methods provide a way to compare these choices over a fixed horizon, but the objective must include due dates and terminal stock rather than setup cost alone. [S04] The planner should test the actual candidate lots against process calendars rather than assume that the supplier's rounded quantity is also the best production batch.

Supplier capacity deserves the same test. An accepted purchase order for the full net requirement does not prove that the supplier can produce and release it by the required usable date. Partial deliveries may cover an initial operation while leaving later operations exposed. The plan should therefore retain dated receipt tranches and their accepted quantities. A feasible promise uses the sequence of usable supply and conversion capacity, not the order total alone.

8.5

Work in progress is time as well as money

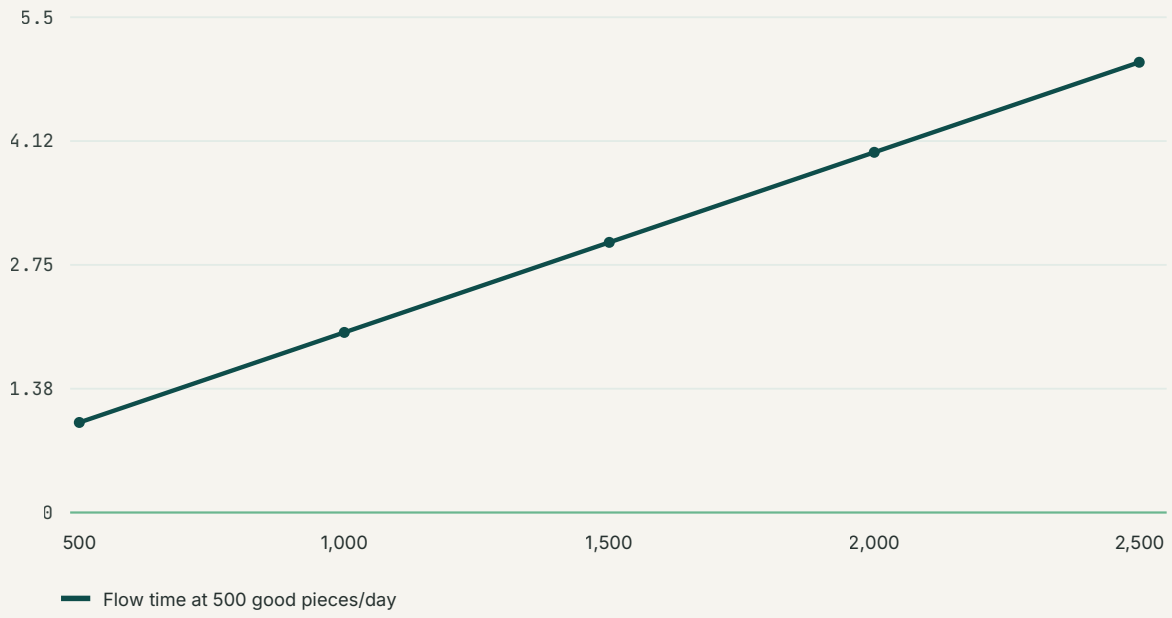
Once material is released, it becomes work in progress within a defined production boundary. That stock represents committed material and conversion cost, but it also occupies delivery time. Little's Law links the average work in progress L, average throughput λ and average flow time W through $L = \lambda W$. The relationship applies to a stable system with consistent units and boundaries. MIT's lean teaching also separates processing time from waiting time, which is essential when diagnosing where flow time accumulates. [S28]

Consistency matters more than algebraic familiarity. If L counts all started but unfinished garments across cutting, sewing and finishing, then λ should be completed good garments leaving that same boundary per day. W is then the average elapsed time a garment spends from entry to exit. Using cutting output as throughput while counting factory-wide work in progress would mix boundaries. Using pieces per hour beside flow time in days would mix units. Quality-held items must have an explicit inclusion rule because they can remain inside the physical boundary while being unavailable to the next process.

Exhibit 34, p. 130 holds illustrative throughput at 500 good pieces per day and varies average work in progress over the same boundary. The calculated flow times follow directly from Little's Law. They are not measured lead times and do not demonstrate that throughput would remain unchanged in a real factory. [S28]

EXHIBIT 34 / ILLUSTRATIVE AVERAGE FLOW TIME IN DAYS

A growing queue consumes delivery time



ILLUSTRATIVE [S28] Illustrative application of W = L/λ with throughput λ fixed at 500 good pieces/day and average WIP L varying from 500 to 2,500 pieces. Flow times are respectively 1, 2, 3, 4 and 5 days. All values use one stable process boundary and consistent units. The chart is an arithmetic teaching case, not a measured relationship or a forecast of queue behaviour.

At average work in progress of 1,500 pieces and throughput of 500 good pieces per day, the illustrative average flow time is `1,500 / 500 = 3` days. Doubling work in progress to 3,000 pieces would imply 6 days only if the same stable throughput and boundary actually held. The equation does not instruct the factory to double throughput, nor does it say that every individual garment waits exactly the average. Its value is to make the time implication of an inventory position inspectable. [S28]

During start-up, severe disruption or another unstable interval, averages can conceal accumulation. A direct flow balance is then more informative: closing work in progress equals opening work in progress plus releases, less completed output, scrap and any other defined exits. For example, an illustrative opening balance of 500 pieces, 600 releases, 500 completions and 20 scrapped pieces gives a closing balance of 580 pieces. This accounting identity explains the change without pretending that the system has reached a steady state.

The diagnosis should distinguish processing, waiting, blocking and quality hold. A high work-in-progress total may reflect long operation queues, a downstream stoppage or batches awaiting disposition. Little's Law cannot identify which mechanism caused the queue. It supplies a reconciliation between stock, throughput and time; the event record must supply the cause. That distinction prevents a planner from treating every work-in-progress reduction as an improvement when it may simply starve a later operation.

8.6

A small buffer may protect a bottleneck while a large queue hides failure

A production buffer has a different purpose from replenishment safety stock. Safety stock protects a defined material service outcome against demand and supply uncertainty. A production buffer decouples operations whose processing times, failures or setups would otherwise stop one another immediately. Both are inventory, but they sit at different boundaries and require different evidence. Moving fabric into the sewing hall does not prove that more protection has reached the customer.

Three events should remain distinct in the line record. **Starvation** occurs when an available operation has no released input to work on. **Blocking** occurs when an operation cannot pass completed work onward because the next buffer or process cannot accept it. **Quality hold** occurs when material or work cannot advance because disposition is unresolved. A single idle-time total mixes these causes and can lead to the wrong remedy. More upstream release may relieve starvation, worsen blocking and increase the amount of work exposed to a quality problem.

A deliberately small buffer before a genuine bottleneck can protect that resource from short upstream interruptions. Its purpose is to keep scarce conversion capacity doing useful work, not to maximise work in progress everywhere. The buffer should have a defined location, unit, maximum quantity, release authority and review rule. Its size needs observed variability and recovery behaviour. Without those inputs, GFT would treat any proposed quantity as illustrative and test it before granting it policy status.

Finite-buffer simulation offers one way to conduct that test. NIST's Simantha method reference explicitly represents asynchronous production lines with finite buffers and includes machine degradation and maintenance functions. [S19] This supports the modelling choice, not a claim that a garment line has been simulated here. A credible model would need the actual route, processing-time distributions, setups, breakdown and repair behaviour, rework loops, release rule and buffer capacities. It would also need enough replications to distinguish a policy effect from random variation.

Controlled release and batch release should be compared under those common conditions. Controlled release admits work when a defined line or buffer condition permits it. Batch release admits a larger order quantity according to the production lot. Neither rule is universally superior. A large batch may be appropriate when setup cost dominates and downstream capacity is available, while controlled release may limit queues when disruption and long waiting dominate. The comparison must hold demand, route, material availability and failure scenarios constant so the release rule is the actual difference being tested.

An illustrative scenario can show why raw loading is a poor objective. Suppose a batch release places 1,500 pieces inside the boundary while a controlled alternative produces a measured average of 1,000 pieces, and both later demonstrate stable throughput of 500 good pieces per day. Little's Law would imply average flow times of 3 days and 2 days respectively. Those outcomes are constructed for arithmetic, not predicted effects of either rule. The test is whether a real pilot can sustain service and throughput with less waiting, not whether the line initially looks busier. [S28]

A large queue can also conceal failure for longer. Work waits without an immediate shortage signal, quality defects may be discovered only after more pieces have been produced, and order identity becomes harder to preserve as batches move. Buffer caps and age alarms make those conditions visible earlier. They do not remove the need to repair the underlying process. GFT would review starvation, blocking, quality-hold age and completed orders together before changing a buffer, because improving one local measure can damage the whole promise.

8.7

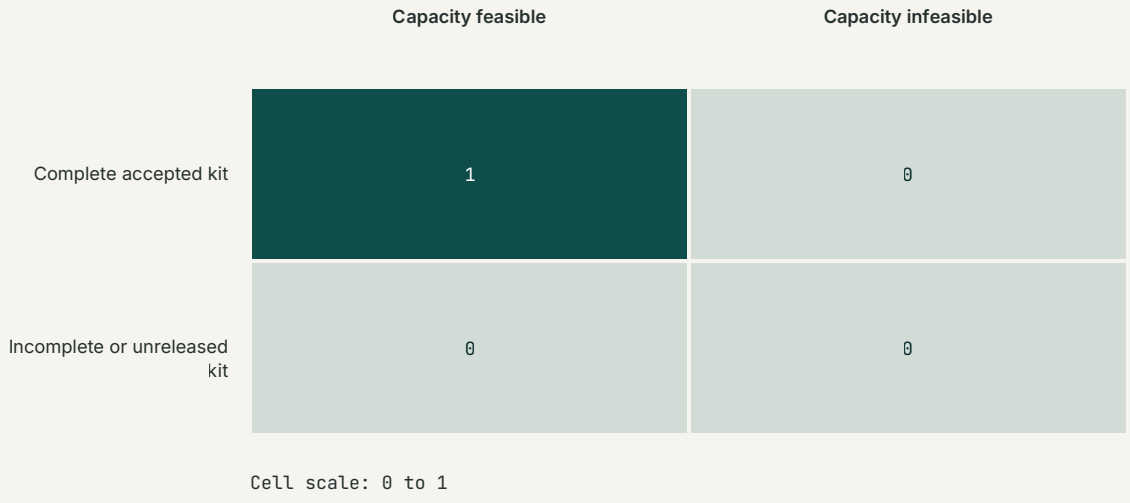
A complete-kit gate can improve the delivery promise

GFT's proposed complete-kit gate joins the preceding tests at the point of production commitment. Release is eligible only when the current order and specification are approved, every required component is accepted in the required quantity and compatible lot, dated receipts no longer conceal a shortage, and the routing fits reserved process capacity. Packaging, work instructions and quality requirements must also be ready. A kit can be physically complete but commercially wrong if its label revision or customer packaging is obsolete.

Exhibit 35, p. 132 expresses the two decisive dimensions as independent gate checks. A value of one means eligible to proceed to the remaining release review, not automatic production. A value of zero means hold and resolve. The states are a proposed GFT control design, informed by dependent-demand planning and finite-capacity modelling; they are not probabilities or measured factory outcomes. [S07] [S19]

EXHIBIT 35 / GATE STATE (0 HOLD, 1 REVIEW)

A material plan can be complete while capacity is infeasible



ILLUSTRATIVE [S07] [S19] Illustrative decision matrix. Material completeness requires every accepted component, current approved specifications and packaging readiness. Capacity feasibility requires sufficient qualified process time in the required slot. Only a complete kit with feasible capacity receives state 1, meaning eligible for the remaining human release review. All other combinations hold. Values encode a proposed gate, not frequencies, scores or observed performance.

The matrix prevents two common forms of false confidence. A full warehouse does not cure an overloaded sewing slot, and an idle line does not cure a missing approved zip. The gate must evaluate both conditions at the same order date. It also prevents an aggregate factory average from hiding order-level failure: total spare hours elsewhere cannot be assigned to this order unless the calendar, skills, machine capability and other commitments genuinely permit it.

> **GFT NOTE**
>
> Run the complete-kit gate before material issue, then repeat it at any order, specification, receipt or capacity change that affects feasibility. Keep the first result and every later exception. If the gate fails, require the planner to choose among approved substitution, a feasible new slot, a revised quantity or a renegotiated date. Do not release an incomplete kit merely to raise line loading. Where data are insufficient, mark the input unknown and hold the affected promise rather than treating an empty field as zero delay or zero requirement. [S07]

The operating sequence can remain deliberately simple. Planning freezes the candidate order version and produces the component exception list. Stores confirms physical identity, quantity and allocation. Quality confirms acceptance and any approved concession. Production confirms the route and feasible slot. The authorised release role reviews unresolved exceptions and records the decision. These are role assignments, not invented staff or evidence of an installed GFT system.

Exceptions should preserve the reason for failure. A missing-material exception needs the component, net quantity, required date and latest feasible action. A quality exception needs the held lot, disposition owner and decision deadline. A capacity exception needs the overloaded process, required and available time in the same unit, and affected orders. A data exception needs the missing or conflicting field. Combining all of them into a generic late-order flag would make recurrence impossible to analyse.

The gate's value must be judged at order completion, not at release. GFT would compare material-caused stoppage, work-in-progress age, first-pass acceptance, completed-order delivery performance and exception frequency. Flow-time and throughput measures need the same boundary established earlier, while queue causes should remain separated. [S28] Ship and receive discrepancies should remain linked to their original records so a later correction does not erase the reason a kit was held. [S31]

There is also a worker and supplier safeguard in this measurement plan. A late customer change or unrealistic release should not be absorbed through unrecorded overtime or emergency subcontracting as a matter of routine. Those responses belong in the exception record with their cost and approval. The purpose of the gate is to expose infeasibility early enough for a fair decision, not to move the same pressure to a less visible part of production.

A complete-kit gate does not guarantee on-time delivery. Demand can change, equipment can fail and accepted material can later prove defective. It instead establishes that the factory had a coherent basis for its commitment at release and a trace-

able account of every deviation. That is the appropriate bridge from inventory optimisation to WP15 line balancing and WP02 performance management: material-ready constraints inform the ramp plan, and starvation, blocking and quality hold become interpretable line signals rather than unexplained idle time. [S07] [S28]

Release only what the complete kit and the line can finish, and let every exception improve the next promise.

CHAPTER SOURCE NOTES

S07 MIT OpenCourseWare. 2006.

S19 National Institute of Standards and Technology. 2022-01-05.

S28 MIT OpenCourseWare. 2012.

S31 GS1 US and Auburn University RFID Lab. 2018.

S04 MIT OpenCourseWare. 2006.

S03 MIT OpenCourseWare. 2006.

GREEN FUTURE TECH
LIMITED / HONG KONG SAR

09

PART III: RESOLVE CONSTRAINTS ACROSS THE SUPPLY CHAIN / NETWORK CHOICES

Pool material and demand only where specifications and risks permit it

Centralisation and postponement can reduce duplicated buffers while increasing transit time, concentration and conversion risk.

CHAPTER 09
DECISION TEST

Centralisation and postponement can reduce duplicated buffers while increasing transit time, concentration and conversion risk.

Material and demand should be pooled only where the materials are genuinely interchangeable, the combined uncertainty is smaller, and the network can still convert and deliver within the customer promise. A common warehouse or a common fabric code does not establish those conditions. Green Future Tech Limited (GFT), Hong Kong SAR, therefore treats pooling as a conditional network design, not a standing instruction to centralise. The decision must combine demand dependence, technical compatibility, conversion capacity, usable-receipt time, shared disruption exposure and residual-stock ownership. [S01] [S05] [S07] [S17]

This chapter compares factory, hub, distributor and consignment positions without presenting any of them as existing GFT assets. It derives the pooling equation, tests common-input postponement, reconciles local and network stock boundaries, prices transport consolidation and examines suppliers that fail together. Every numerical case is illustrative because no factory demand history, route sample, supplier quotation or warehouse ledger has been supplied. The recommendation is to prove each proposed pool against end-customer service and whole-chain inventory before moving either stock or ordering authority. [S03] [S08] [S12] [S30]

9.1

A stock location is a service decision

A stock location earns its place by protecting a defined promise at a defined point of use. Fabric held at a sewing factory may support immediate cutting. The same fabric at a central hub may serve several factories but require a transfer, a receiving check and another allocation decision. Stock at a regional distributor may be close to a customer while remaining too far from the production operation that needs it. Consignment changes the ownership arrangement, not automatically the physical distance. Distribution requirements planning connects locations and replenishment timings, but it does not choose the service promise on the planner's behalf. [S07]

The comparison should therefore begin with the event that the buffer must protect. A make-to-order garment producer may need accepted fabric at the cutting room before a scheduled operation. A repeat-stock seller may need a complete size assortment at a dispatch point before the confirmed collection. A distributor may need saleable finished goods rather than material that still requires finishing. These promises have different stock objects and different clocks. Combining them in one network total would conceal the very shortage that the location decision is meant to prevent. [S01] [S07]

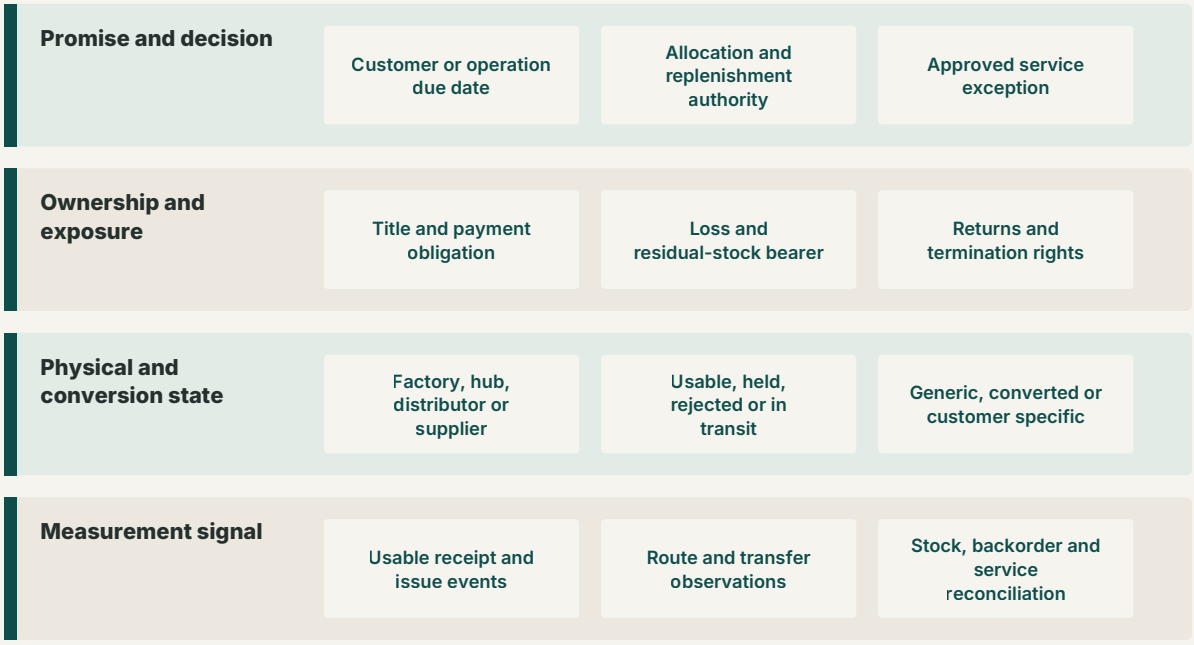
Four boundaries make the choice inspectable. The physical boundary records where the material is and whether it is usable. The service boundary records the point and time at which a customer or internal operation may rely on it. The conversion boundary records what work remains before differentiation is complete. The ownership boundary records who carries the cost, loss and residual-risk exposure. A centralised design can improve one boundary while weakening another. For example, moving generic fabric upstream may reduce duplicated factory buffers but lengthen the interval between a confirmed style requirement and approved fabric at the cutting table.

A hub also introduces an allocation rule. When several factories or channels draw from one pool, the planner must decide which requirement receives stock after a delayed receipt. An aggregate availability report cannot answer that question. The policy needs a priority rule, an exception owner and a record of the orders that lose coverage when priorities change. Otherwise, centralisation can report good network service while repeatedly starving the same location. This is a governance requirement proposed by GFT, not a feature that a warehouse automatically supplies.

Exhibit 36, p. 137 separates these responsibilities before any quantities are compared. It is an illustrative design, not a map of an implemented network. Its purpose is to stop a change of address or title being mistaken for a change in service or physical stock. [S01] [S07] [S08]

EXHIBIT 36 /

Inventory boundaries differ from ownership boundaries



ILLUSTRATIVE [S01] [S07] [S08] Illustrative GFT network-design schema. The layers are separate questions for every proposed pool. No supplier, factory, hub or distributor is asserted to exist, and no service or stock reduction is claimed. S01 supports differentiated supply-chain policies, S07 location-based planning and S08 explicit contractual allocation.

The layered view also prevents a false choice between one central warehouse and many independent stores. A network may centralise information while leaving stock near use, or centralise a generic input while distributing finished variants. It may also hold a small operational buffer locally and a contingency pool upstream, provided each buffer has a distinct protection period. The relevant design variable is not the number of buildings. It is where an irreversible commitment and a measurable service risk are placed. [S01] [S05]

Actual route and channel evidence must settle the final choice. The planner needs measured transfer time to usable receipt, handling and inspection requirements, shipment frequency, channel returns, stock visibility and the cost of an emergency alloca-

tion. A location drawn on a diagram supplies none of these inputs. GFT would leave them unresolved until dated observations or explicitly approved scenarios are available, rather than insert a nominal journey time into a business case. [S12] [S30]

> **GFT NOTE**
>
> Create an item-location service record before approving a pool. Name the stock object, approved substitutes, point of use, protection period, allocation authority, transfer evidence and residual-stock owner. Reject the proposal if any of those fields has no accountable answer. A lower network stock total is not sufficient approval evidence when the promise or failure boundary has changed.

9.2

Demand correlation determines pooling potential

Pooling combines uncertain requirements before a buffer is sized. Its statistical benefit comes from imperfectly correlated deviations, not from adding average demand together. For two demand or forecast-residual series D_1 and D_2, with standard deviations σ_1 and σ_2 and correlation ρ, the variance of their sum is $\mathrm{Var}(D_1 + D_2) = \sigma_1^2 + \sigma_2^2 + 2\rho\sigma_1\sigma_2$. This is arithmetic applied to the probabilistic inventory framework, with both series measured in the same unit over the same protection period. [S05]

The covariance term is decisive. If deviations in one location tend to offset those in another, the combined series can be less variable than the sum of separate exposures. If the same promotion, season, tender or disruption raises demand in both locations together, positive covariance removes much of that benefit. At perfect positive correlation between equal-variance demands, pooling does not reduce the standard-deviation total at all. Negative correlation can produce a larger reduction, but a planner should demonstrate a stable mechanism rather than assume that an attractive historical coefficient will persist. [S05]

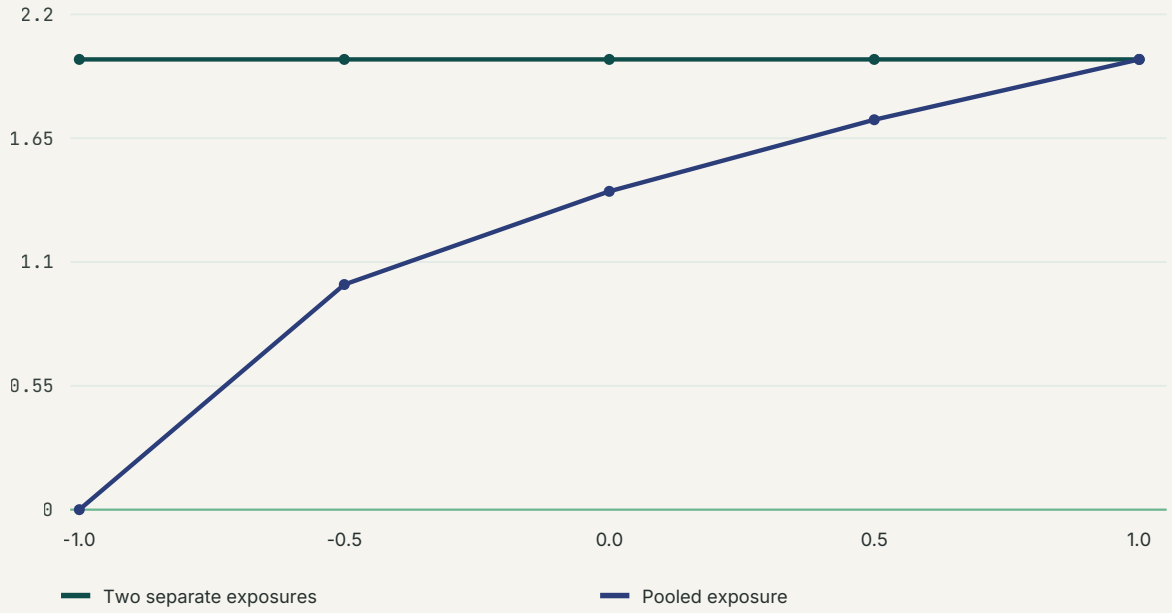
The familiar square-root result is a special case. For n comparable, independent demands with equal standard deviation σ, the pooled standard deviation is $\sigma\sqrt{n}$, while the sum of separate standard deviations is $n\sigma$. A buffer proportional to standard deviation would therefore grow with $\sqrt{n}$ rather than n under that construction. The result does not say that total inventory must fall by the same proportion. Cycle stock, pipeline stock, minimum lots, local response stock, conversion losses and different service targets sit outside the simplified comparison. [S03] [S05]

Three further conditions matter. First, the input should be the forecast residual relevant to the replenishment decision, not automatically raw sales. Second, the pooled stock must be able to reach each requirement before its due date. Third, one service definition must govern the comparison. A central buffer that offers a lower unit fill rate or a later delivery cannot be declared more efficient merely because it holds fewer units. GFT would compare the policies on common demand paths and record any unmet or late demand separately. [S05] [S07]

Exhibit 37, p. 139 shows the equal-variance boundary for a pair of demands. Both separate locations have one standard-deviation unit each, so their separate standard-deviation total is two. The pooled line follows the variance equation across constructed correlation scenarios. These are teaching values, not estimated apparel demand parameters. [S05]

EXHIBIT 37 / ILLUSTRATIVE STANDARD-DEVIATION UNITS

Pooling benefits shrink as demands move together



ILLUSTRATIVE [S05] Illustrative equal-variance pair with sigma 1 = sigma 2 = 1. Separate total = sigma 1 + sigma 2 = 2. Pooled standard deviation = sqrt(1 + 1 + 2 rho), for correlation rho from -1 to 1. The chart compares uncertainty exposure only. It excludes means, cycle stock, pipeline stock, transfer time, capacity, cash and service differences. Correlation scenarios are not measured observations.

The horizontal separate-exposure line is not a recommended buffer. It is the benchmark that makes the covariance effect visible. Likewise, the zero at perfect negative correlation is a mathematical boundary for this equal-variance pair, not an operational promise of zero stock. Real policies still need expected demand coverage, feasible replenishment and protection against model error. The chart cannot establish that centralisation is worthwhile once transfer time and concentration are included.

WORKED EXAMPLE: COVARIANCE CHANGES THE BUFFER, NOT THE MEAN

An illustrative comparison uses protection-period forecast residuals with standard deviations of 120 units and 90 units, correlation $\rho = 0.25$, and a common normal-approximation service factor of $z = 1.645$. Separate standard-deviation-based safety stocks total $1.645 \times (120 + 90) = 345.45$ units. The pooled variance is $120^2 + 90^2 + 2 \times 0.25 \times 120 \times 90 = 27{,}900$, so pooled standard deviation is 167.03 units and pooled safety stock is $1.645 \times 167.03 = 274.77$ units. [S05]

The illustrative difference is 70.68 units, or 20.46 percent of the separate safety-stock total, before integer rounding. It is not a measured saving. The case holds expected demand, service factor, unit and protection period constant, and excludes every additional network cost. If the common pool adds transfer delay or the two channels require different service targets, this arithmetic alone cannot approve it. [S03] [S05] [S07]

A usable empirical study would retain aligned residual series and explain missing observations. It would test correlation across ordinary and disrupted periods, rather than report one coefficient for the full history. A positive shift during a shared selling event may be predictable and removed by the forecast, while the remaining residuals may behave differently. Conversely, a calm historical sample may contain no common port closure, supplier failure or coordinated promotion. WP04 owns the demand-estimation method; the network decision here consumes that evidence without treating a correlation estimate as a permanent physical property. [S05] [S24]

The practical output is a range of pooling scenarios with their conditions. The planner should preserve the separate-stock comparator, show the covariance assumption, and add route and capacity constraints before interpreting the stock difference. Where the dependence structure is unstable or the sample is too small, a scenario range is more honest than a single optimised network. The pooling equation remains valid arithmetic, but its input credibility determines the strength of the recommendation. [S05]

9.3

Common fabric can pool risk before differentiation

Demand pooling and material pooling are related but not identical. Two styles can have imperfectly correlated demand yet still require incompatible fabric. Two styles can use the same approved fabric yet face demand that rises together, leaving little statistical pooling benefit. The strongest candidate combines technical commonality with useful demand behaviour, but it must also survive the conversion and delivery tests. A shared item name is not evidence of interchangeability. Material planning relies on the product structure and the exact component requirement. [S05] [S07]

A technical commonality review should compare fibre and construction requirements, width, weight, finish, colour or shade rules, performance tests, customer approvals and permitted substitutions. It should also retain the bill-of-material version. An unprinted fabric may be common at one process stage and become non-interchangeable after dyeing, printing, coating or cutting. A trim may be physically similar

but commercially unique because a customer has approved only a named specification. GFT would pool at the latest verified common stage, not at the earliest convenient purchasing description. [S07]

Postponement can preserve that common stage until better demand information arrives. Instead of holding finished stock for every style, the network holds a generic input and converts the required variant after an order or a more reliable allocation signal. The statistical buffer is then attached to the common input. This can reduce duplicated variant buffers, but it creates work that must still be performed. Dyeing, printing, embroidery, cutting, finishing, inspection and transfer all consume time and may require minimum batches or setups. [S07] [S17]

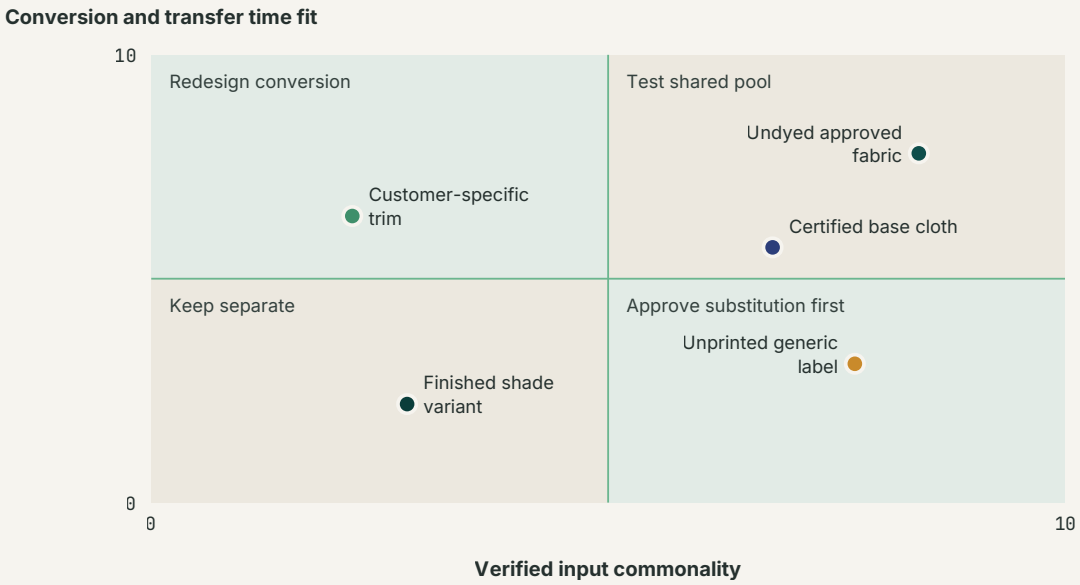
The Japanese apparel evidence supports a conditional rather than promotional reading. The opened study reports that postponed manufacturing was associated with lower inventory, while postponed planning increased inventory when apparel firms externalised their retail function. It is a wholesaler survey in Japan, not a causal forecast for a Cameroon garment producer. Its useful warning is that delaying a decision can weaken control or information quality even when delaying physical differentiation helps. [S17]

A feasible postponement interval must fit inside the remaining promise window. If T_p is the time left after the demand signal, T_c the conversion and quality-release time, and T_t the transfer to the point of use, the simple feasibility condition is $T_c + T_t \leq T_p$. The terms must use the same calendar and end at usable availability. Setup queues, failed inspection and capacity contention require scenarios or measured distributions, not a single optimistic duration. A mathematical pool that cannot be converted in time is only stock in the wrong form. [S07] [S19]

Exhibit 38, p. 142 places illustrative candidates against commonality and time fit. The positions are proposed screening judgements, not measured factory scores. They show why technical approval and process timing must be assessed together before a shared buffer is approved. [S07] [S17] [S19]

EXHIBIT 38 / ILLUSTRATIVE SCREENING SCORE FROM 0 TO 10

Postponement is feasible only when conversion and promise times fit



ILLUSTRATIVE [S07] [S17] [S19] Illustrative GFT screening model. The horizontal axis combines verified technical and commercial interchangeability. The vertical axis represents fit among conversion, quality release, transfer and remaining promise time. Coordinates are constructed teaching judgements, not observations or statistical estimates. S07 supports specification-based requirements, S17 the conditional postponement evidence and S19 finite-capacity process modelling.

A candidate in the test quadrant still needs an operating trial. The trial should record conversion setup, first-pass acceptance, queue time, transfer to usable receipt and service by originating requirement. Finite-buffer simulation can represent queues, breakdowns and maintenance where those mechanisms matter, but a model does not validate its own input distributions. No such factory simulation has been performed for this chapter. [S19]

The economic comparison must also include new cycle stock. A common dye or print process may require a minimum run larger than the immediate variant requirement. Pooling can then reduce safety stock while increasing batch stock and residual variants. Setup cost and lost capacity may outweigh the uncertainty benefit, particularly when differentiation is inexpensive at a later stage closer to use. GFT would compare total relevant cost across the whole protection and conversion cycle rather than count only the reduction in variant safety stock. [S03] [S07]

The decision is reversible only before differentiation becomes irreversible. That makes approval of the common stage a technical and commercial gate. If a late specification change invalidates the pool, the record should identify which commitments consumed the old version and who owns the residual material. The aim is not maximum commonality. It is the latest common form that can still meet the promise without creating a larger conversion or obsolescence exposure. [S07] [S17]

9.4

Multiple echelons can count the same uncertainty more than once

A network can duplicate protection without duplicating a warehouse code. A hub may size its buffer from factory orders, while each factory sizes its own buffer from end demand and uncertainty about the hub. If those decisions are made independently, both stages can protect against the same downstream deviation. The factory orders also contain batching and replenishment logic, so they are not an independent demand process that can simply be added to end-customer variability. A coordinated comparison needs one network boundary and a shared service definition. [S05] [S07]

Installation stock answers a local question: what accepted quantity is held at this location under the stated status rules? Pipeline stock is moving between locations and is not yet usable at the destination. Customer backorders are unfilled obligations, not negative physical stock on a shelf. Echelon stock looks across a defined downstream network boundary so that inventory and commitments can be coordinated without treating each location as an isolated business. Distribution planning provides the location-and-time structure; the precise accounting convention must still be declared. [S07]

For a proposed GFT echelon view, physical stock within the boundary is counted once, whether it sits at the hub, a downstream site or an internal transfer buffer. Firm inbound receipts are shown separately until they qualify under the planning rule. Backorders reduce the echelon inventory position under the selected convention but never turn quarantined or rejected goods into usable supply. Returns await inspection rather than being added immediately to sellable stock. These are control definitions for comparison, not claims that one formula fits every contract or accounting system. [S07] [S09]

METHOD

Reconcile an echelon before comparing policies:

- Draw the physical boundary and list every included location, transfer and ownership state.
- Reconcile accepted local stock to the network total, eliminating internal transfers counted at both ends.
- Separate firm inbound supply, pipeline quantity, quality holds and customer backorders.
- Apply one end-customer service measure and one protection-period convention to both the local and coordinated policies.
- Retain local exception reports so a network average cannot hide a persistently short location. [S05] [S07]

The comparison should then run the same demand and supply paths through both policies. Independent local buffers form the comparator, not an assumed inefficient straw case. The coordinated policy may place a common buffer upstream, preserve a

small local response quantity and allocate against visible downstream positions. It wins only if end-customer service, total usable stock, cash exposure and exception outcomes improve under the agreed measures. A reduction at the hub accompanied by larger factory shortages is not a network improvement. [S05] [S07]

A decentralised policy can remain preferable. Long or unreliable transfers, weak stock visibility, different local promises and frequent urgent requirements may make upstream coordination too slow. Local ownership can also make discrepancies easier to investigate. The purpose of echelon analysis is not to prove that central control is superior. It is to reveal where uncertainty is protected, where inventory is counted and which stage has authority when the plan becomes infeasible. [S01] [S07]

9.5

Transport consolidation trades freight cost against stock and time

Consolidation can reduce freight per unit while increasing the inventory needed to support the same promise. A larger shipment may lower the transport rate or reduce handling events, but waiting to fill the load can lengthen the replenishment interval. That interval creates cycle stock and may delay the usable receipt that triggers a reorder. A central hub can add another transfer and inspection before material reaches production. The transport quote is therefore one component of the network policy, not its total cost. [S03] [S07]

A useful comparison holds the delivered service boundary constant. For each shipment design, record order frequency, minimum and maximum load, planned dispatch interval, transit to physical arrival, clearance, receiving, quality release and emergency recovery. The inventory side includes expected cycle stock, pipeline quantity, uncertainty protection, damage or rejection, ageing and residual goods. The cost side includes freight, handling, carrying, expected expediting and service failure where those can be estimated without duplicating charges. Total-relevant-cost thinking is more informative than minimising a single rate. [S03] [S30]

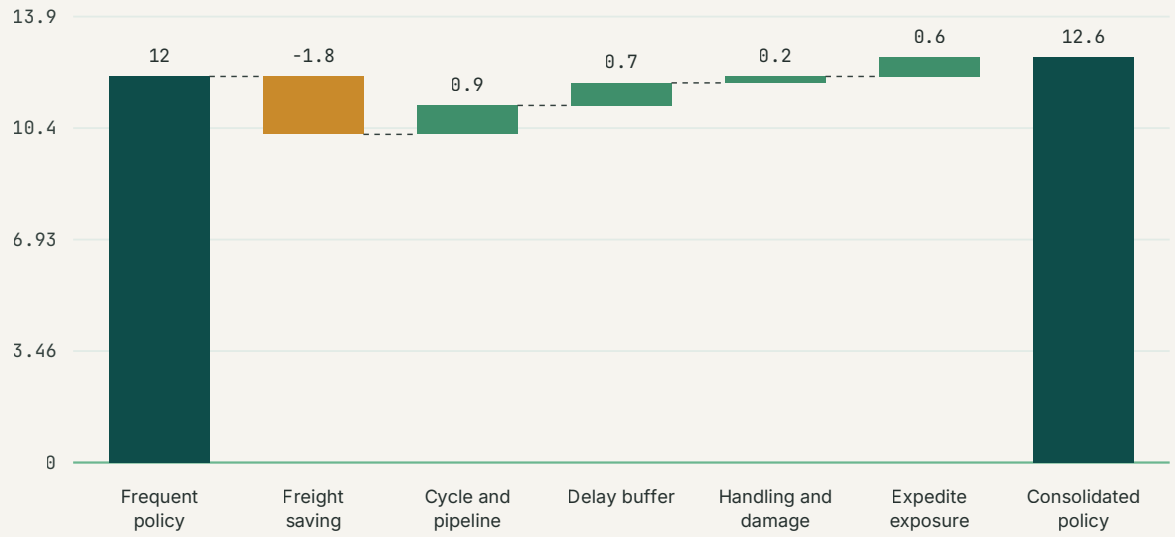
Contract terms must be read alongside the physical route. The selected Incoterms rule and named place help allocate carriage tasks, costs and risks, but they do not make the journey instantaneous and do not by themselves settle every question of title or payment. The model should retain the party responsible for each stage and the actual elapsed time whoever pays for it. Legal interpretation of a live contract requires professional review; this chapter is not legal advice. [S29]

Evidence quality is especially important for route comparisons. The World Bank's revised logistics indicators use shipment-level operational data to examine connectivity, speed and reliability, reinforcing the value of observed movements rather than perception alone. That international source does not provide a factory-specific Cameroon transit distribution. The World Customs Organization guide instead supports process mapping, sampling and segmented time analysis around release. Together they justify collecting actual route events, not adopting a generic corridor average. [S12] [S30]

Exhibit 39, p. 145 constructs a case in which a freight saving is more than offset by inventory and recovery effects. All amounts are illustrative annual teaching assumptions in millions of CFA francs (XAF). They are not quotations, tariffs, measured savings or a recommendation against consolidation. [S03] [S29]

EXHIBIT 39 / ILLUSTRATIVE XAF MILLION PER YEAR

Freight savings can be consumed by slower replenishment



ILLUSTRATIVE [S03] [S29] [S30] Illustrative whole-chain annual cost bridge. Opening frequent-shipment policy cost = 12.0. Consolidation saves 1.8 in freight, then adds 0.9 cycle and pipeline carrying cost, 0.7 additional delay-buffer cost, 0.2 handling and damage exposure and 0.6 expedite exposure. Closing consolidated cost = 12.0 - 1.8 + 0.9 + 0.7 + 0.2 + 0.6 = 12.6. Every input is a teaching assumption. No Cameroon quote, tariff, lead time or realised cost is asserted.

In the illustration, a procurement view sees an annual freight reduction of XAF 1.8 million, while the whole-policy comparison shows an increase of XAF 0.6 million. Both statements follow from the constructed inputs, but only the second includes the inventory and recovery consequences. The result could reverse with reliable frequent departures, low-value materials, short transfer times or a genuine reduction in duplicated handling. The waterfall is a reconciliation template, not a universal direction of effect. [S03] [S30]

The decision should also expose service outcomes that are difficult to price honestly. If a late consolidated shipment stops a line or misses a customer collection, the planner should show the affected order and recovery action rather than bury it in a speculative goodwill charge. Where cash limits prevent the extra stock, the consolidated design may be infeasible even when its modelled annual cost is lower. The alternative is then to change shipment frequency, renegotiate the minimum load or revise the service promise explicitly. [S03] [S05]

Before approval, GFT would request dated carrier and handling quotations, actual transit and release samples, packaging and damage evidence, and a statement of which party bears each cost. Open shipments must remain visible in the sample so

completed journeys do not hide the tail. Until those inputs exist, the network case should present scenarios with unresolved fields, not a precise freight or journey claim. [S12] [S30]

9.6

Two suppliers may still share one disruption

Supplier count is not the same as independent supply. Two vendors may buy from the same upstream mill, use the same port, depend on the same power or communications infrastructure, or subcontract to the same specialist process. A garment network can therefore appear diversified at the first contractual tier while remaining concentrated at the point that actually determines usable receipt. The OECD notes that subcontracting is common in garment and footwear supply chains, which makes indirect stages a necessary part of the dependence map. [S25]

The same caution applies to pooled inventory. A single hub may serve several demand points efficiently during ordinary variation but concentrate exposure to a site outage, system failure, transport interruption or quality hold. Separate local stocks can contain some events, while a common pool can deplete everywhere after one upstream shock. The choice is not between resilience and efficiency as slogans. It is between explicit failure modes, quantities and recovery times under each design. [S05] [S24]

Probability notation makes the missing evidence visible. If `A` and `B` are the events that two sources are usable when required, the probability that at least one is usable is `P(A or B) = P(A) + P(B) - P(A and B)`. Independence would simplify the joint term, but it is precisely the assumption that a shared mill, route or subcontractor can invalidate. The planner needs evidence about the intersection, not just each supplier's separate on-time history. [S05] [S25]

Stochastic lead-time research provides another warning against convenient assumptions. The opened manuscript analyses models with independent and identically distributed lead times while also discussing their limitations and the need to study dependence between lead times and demand. It does not establish that any garment supplier has independent delays. GFT would use such models as scenario structures and test their assumptions against usable-receipt observations. [S24]

A stress design should contain at least three distinct paths: ordinary separate delays, a shared upstream or corridor event, and loss of the pooled location or conversion process. The first path tests routine buffers. The second tests whether apparent source diversification survives common dependence. The third tests concentration created by the network itself. Each path should record which customer promises fail, which substitute is technically approved, how much stock remains usable and who may authorise emergency allocation. No probability needs to be invented to make those actions inspectable.

Qualification must precede the emergency. An alternate source needs the same specification, test-lot, capacity and usable-receipt checks as any other source, together with a map of its upstream dependencies. A second contact in a procurement data-

base is not a qualified recovery route. If both qualified sources share an unavoidable dependency, the policy should say so and consider a different material, route, stock position or customer promise rather than assign false independence. [S07] [S25]

This analysis can also change the demand-pooling result. A common disruption may raise demand for substitutes at the same time that it restricts supply, producing adverse correlation between requirements and availability. The calm-period covariance estimate would then overstate pooling protection precisely when the buffer is most needed. GFT would therefore report ordinary-demand pooling and shared-shock stress separately, with no claim that one historical correlation coefficient covers both. [S05] [S24]

9.7

The network decision must state who carries residual stock

Pooling changes incentives as well as quantities. A supplier-owned buffer, a factory-owned hub, distributor stock and consignment arrangements allocate ordering authority, payment, loss and obsolete material differently. Supply-chain contract analysis shows why the allocation of overage risk matters to ordering incentives and why a buyback or return right must be assessed across the chain. A contract can move exposure without removing the physical units or the uncertainty that created them. [S08]

The decision record should state who forecasts, who may release replenishment, when quantity becomes binding, who owns stock at each stage, who pays for carriage and storage, and who may return or redirect residual units. It should also define damage, shrinkage, quality rejection, specification change, termination and dispute treatment. Visibility is a separate field: a party that owns stock may still lack timely quantity and status information, while a party without title may carry the operational consequence of a shortage. [S07] [S08]

Consignment is a clear example of why factory stock and whole-chain stock must be reported separately. Moving material to a supplier's balance sheet may reduce the factory's recorded inventory or change its cash timing, but the material still occupies a location and may still be reserved for the same customer promise. Conversely, shifting finished goods to a distributor does not prove that end demand has been served. The network review should count physical quantities once and show ownership as an additional dimension rather than using ownership to erase stock. [S07] [S08]

Residual value requires equal care. IAS 2 states that inventories are measured at the lower of cost and net realisable value, which provides accounting context for write-down exposure. It does not tell the network which pooling policy is operationally preferable, and applicable reporting treatment needs professional confirmation. A pooled generic input may retain more recovery options than differentiated variants, but that advantage depends on actual alternative uses, approval rules and selling or recovery costs. [S09]

A whole-chain review should also ask how the apparent improvement was obtained. The ILO's global supplier survey motivates attention to purchasing practices and working conditions, while its age, voluntary response and international scope prevent any inference about current Cameroon prevalence. Within this chapter, the implication is narrow: do not describe lower factory inventory as a clean gain if it depends on unexamined supplier pressure, unrealistic changes or unpaid recovery work. [S10]

GFT NOTE

Approve a pooling pilot only with six artefacts: a verified common-specification record, aligned demand-residual evidence, a conversion and route feasibility test, an echelon stock reconciliation, a shared-failure map and a residual-ownership schedule. Evaluate the pilot against end-customer service and whole-chain usable stock, with the separate-location policy retained as the comparator. Stop the pilot if the pool improves the reported factory balance while worsening service, total exposure or exception control. [S05] [S07] [S08]

The final network choice should present a frontier rather than a universal answer. A central generic-material pool may suit compatible inputs, offsetting demand and reliable conversion. Local buffers may dominate where promises are short, routes unreliable or failure containment matters. A mixed design may be appropriate when local response stock and upstream contingency stock protect different horizons. In each case, the recommendation remains conditional on measured service, feasible conversion, visible stock and an accountable owner for what remains. [S01] [S05] [S07]

Pool compatible requirements, but never hide a shared failure inside a shared buffer.

CHAPTER SOURCE NOTES

S01 MIT OpenCourseWare. 2006.

S05 MIT OpenCourseWare. 2006.

S07 MIT OpenCourseWare. 2006.

S17 Japan Marketing Academy, Quarterly Journal of Marketing. 2020-03-31.

S03 MIT OpenCourseWare. 2006.

S08 MIT OpenCourseWare. 2006.

S12 World Bank. 2025.

S30 World Customs Organization. 2025.

S24 arXiv, Michna, Nielsen and Nielsen. 2015-03-20.

S19 National Institute of Standards and Technology. 2022-01-05.

S09 IFRS Foundation. n.d..

S29 International Chamber of Commerce. 2020.

S25 OECD. 2026-02-26.

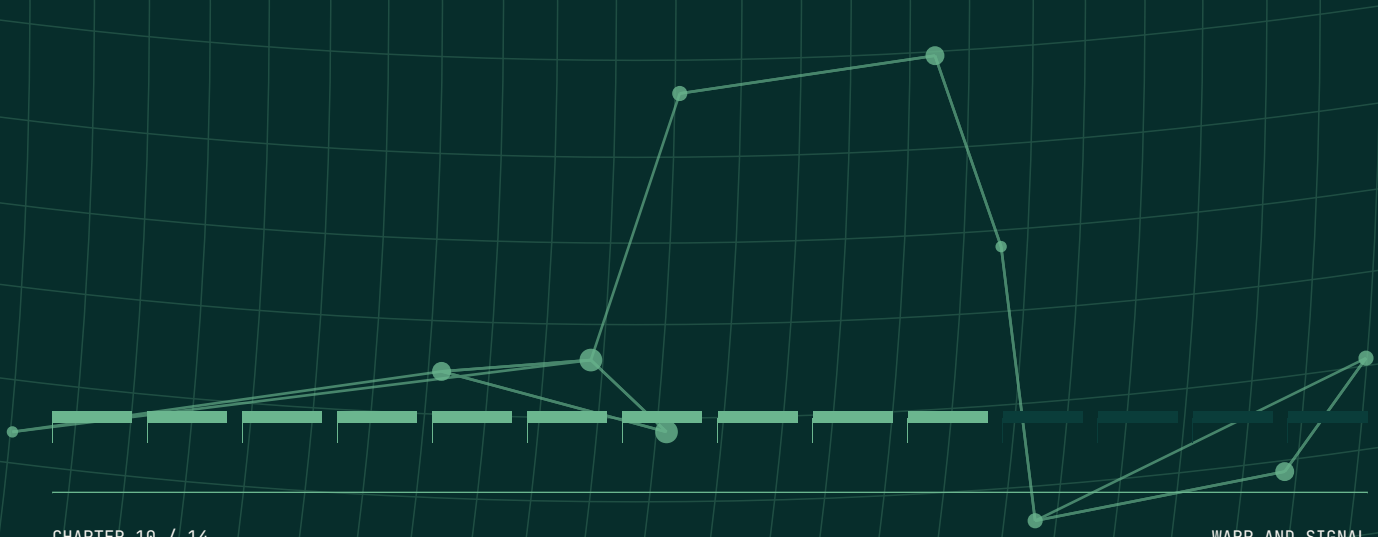
S10 International Labour Organization. 2017-06-09.

10

PART III: RESOLVE CONSTRAINTS ACROSS THE SUPPLY CHAIN / PURCHASING CHOICES

Supplier reliability and fair contracts change how much stock is needed

Procurement should price usable delivery and shared risk rather than reward low quotes that create expensive disruption.



CHAPTER 10
DECISION TEST

Procurement should price usable delivery and shared risk rather than reward low quotes that create expensive disruption.

A garment factory cannot size its material buffers independently of how it selects, contracts with and pays its suppliers. When a purchasing rule rewards the lowest nominal quote without pricing late arrival, shade rejection, missing documentation or uncompensated schedule changes, the cost reappears on the factory floor as larger safety stock, starved sewing lines, emergency freight or harmful pressure on workers. [S03] [S10] [S20] Green Future Tech Limited (GFT), Hong Kong SAR, treats supplier reliability and contractual risk-sharing as primary parameters of the inventory decision. In GFT's Cameroon textile design context, where national policy seeks to raise local cotton processing from roughly 5 percent against 95 percent raw export toward 50 percent by 2030, a garment maker buys yarn, fabric, thread, interlinings, zips and packaging across multiple counterparties. [R06] [R07] [S07] This chapter specifies how an apparel maker should measure usable delivery, structure purchasing gates, diagnose order amplification, separate forecasts from commitments, evaluate flexibility contracts under cash constraints, prevent purchasing changes from harming workers and qualify alternate sources before disruption strikes. All worked figures and thresholds are illustrative design proposals, not operating results from an installed GFT facility or legal, tax or investment advice.

10.1

An on-time dispatch can still be an unusable delivery

A supplier shipment that leaves the mill or reaches the factory gate on the promised date does not protect production if the material cannot be released to the cutting table. In apparel manufacturing, a roll of fabric with an unapproved shade variation, a carton of zips with the wrong slider specification or a dye lot lacking chemical compliance documentation immobilises dependent components just as effectively as an empty truck. [S07] [S20] Inventory models that treat physical arrival as equivalent to usable supply understate the protection horizon and overstate available stock. [S05] [S30] GFT defines a supplier delivery outcome at the usable-receipt boundary: the timestamp and quantity at which an inbound lot passes physical count reconciliation, specification verification, shade-band matching and document clearance.

Empirical work on supply-chain data exchange shows why physical receipt and usable reconciliation must remain distinct events. In the Project Zipper study by GS1 US and the Auburn University RFID Lab across eight brand owners and five retailers, researchers examined shipment and receipt records to diagnose data-exchange and order-accuracy discrepancies across trading partners. [S31] That retail and brand study does not supply a transferable error rate for a Cameroon garment factory, nor does it prove that automated scanning alone resolves non-conformance. [S31] Its lesson for GFT is structural: a shipping notice, a carrier arrival scan and a verified usable receipt are distinct informational states, and collapsing them into one timestamp conceals where quality holds and documentation disputes consume buffer time. [S30] [S31]

10.2 SEPARATE SHIPMENT PUNCTUALITY FROM USABLE SUPPLIER SERVICE

To prevent averages from hiding line-stopping failures, GFT decomposes supplier delivery into four measures at the purchase-order-line boundary: dispatch punctuality against agreed carrier handover, physical arrival punctuality at the receiving gate,

first-pass quality acceptance at initial inspection before rework, and supplier on-time in-full usable receipt (usable OTIF), which requires the complete ordered quantity within governed tolerances to achieve approved usable status on or before the committed release date. [S05] [S07] [S30] [S31]

Recording these events separately changes how replenishment uncertainty is calculated. As Chen and colleagues observe in their research manuscript on inventory systems with uncertain supplies, both demand and supply observations can be censored when truncated by availability or capacity limits. [S22] Excluding overdue open orders at month-end removes the longest delay tails through right censoring. [S22] [S30] Similarly, when a supplier delivers an illustrative 1,000 metres of shirting fabric in two split dispatches of 600 metres on time and 400 metres twelve working days late, or when 150 metres fail shade inspection and enter quarantine, recording the order as delivered on the first arrival date fabricates usable coverage that the cutting room never possessed. [S07] [S22]

METHOD

GFT specifies an effective usable-supply record for each item-supplier pair over a declared observation window:

1. **Order-line completeness boundary.** Every order line records ordered quantity, promised usable-release date, first physical arrival date, final usable-release date when the last accepted unit clears inspection, and accepted usable quantity after subtracting rejected, quarantined and undocumented units. [S07] [S30] [S31]
2. **Right-censored open orders.** Any order line whose promised usable date has passed without full usable release by the evaluation cutoff remains in the sample with its elapsed delay flagged as right-censored. [S22] [S30]
3. **Effective usable yield.** First-pass usable yield equals accepted usable quantity divided by ordered quantity. If an illustrative supplier exhibits a mean first-pass usable yield of 0.92 with a standard deviation of 0.06, an order of 1,000 metres provides an expected usable receipt of 920 metres before lot-to-lot yield variance. [S05] [S22]
4. **Exception classification.** Split deliveries, shade mismatches, width shortages, missing certificates and invoice disputes are logged as separate failure codes. [S20] [S30]

Evaluating suppliers solely on invoice price and dispatch date forces the planner to carry an unattributed buffer to absorb the inspection gap and rejected fraction of each lot. [S03] [S05] Tracking usable OTIF and first-pass yield charges a vendor whose shipments require 100 percent manual re-sorting with the working capital and schedule risk its performance creates. [S07] [S20]

10.3

Supplier scoring needs visible trade-offs

Procurement scorecards often collapse price, lead time, quality and social compliance into a single weighted-average number. In a linear weighted sum, a low quoted price can mathematically offset a severe human-rights risk, unverified subcontracting or a

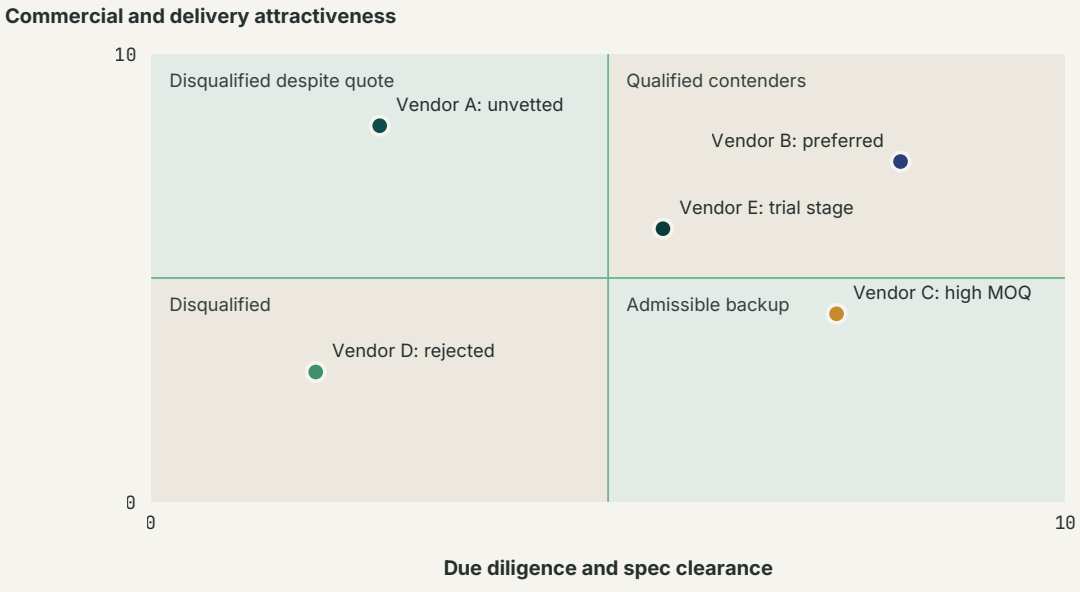
high rate of shade rejection that halts assembly. [S20] [S25] The 2018 OECD Due Diligence Guidance for Responsible Supply Chains in the Garment and Footwear Sector, aligned with the UN Guiding Principles on Business and Human Rights, states that due diligence should be integrated into supplier selection and purchasing management systems. [S20] Furthermore, the 2026 OECD spotlight on garment and footwear due diligence notes that subcontracting is common practice in the sector, meaning that an unvetted low-cost bidder may pass orders to unauthorised workshops outside the buyer's visibility. [S25]

GFT rejects single-index weighted scoring in favour of a two-stage architecture: non-compensatory qualification gates followed by an explicit economic trade-off frontier among admissible suppliers. A candidate must first clear hard thresholds on legal identity, absence of forced or child labour indicators, disclosed manufacturing sites and subcontracting boundaries, minimum technical capability and documentary integrity. [S20] [S25] No price discount can override a failure on these gates. Only after a supplier enters the qualified pool does the planner compare price, usable lead-time distribution, first-pass yield, lot-size minimums and payment terms. [S03] [S05] [S20]

The governance structure in **Exhibit 40**, p. 152 illustrates how GFT separates non-negotiable due-diligence and technical admissibility from commercial comparison. Its coordinates are illustrative ordinal placements showing which candidates may receive an order and which are blocked regardless of unit price. [S20] [S25]

EXHIBIT 40 / ILLUSTRATIVE ORDINAL SCORE (0 TO 10)

A supplier score cannot conceal a disqualifying adverse impact



ILLUSTRATIVE [S20] [S25] Illustrative GFT two-stage supplier evaluation design. Candidates below x = 5 fail mandatory due-diligence, subcontracting-disclosure or specification gates and cannot be selected regardless of commercial attractiveness on the vertical axis. Coordinates are teaching constructs based on S20 and S25 principles, not measured Cameroon vendor scores.

As **Exhibit 40**, p. 152 shows, Vendor A offers an attractive quote (y = 8.4 on this illustrative scale) but relies on undisclosed subcontracting or fails safety verification (x = 2.5), placing it in the disqualified zone. [S20] [S25] A 60/40 weighted average of price and compliance would score Vendor A at 6.04 ahead of compliant Vendor C at

5.52 in the same illustrative arithmetic. Under GFT's rule, Vendor A is blocked until its site disclosure and labour safeguards pass the admission boundary, while admissible contenders Vendor B and Vendor C are compared on total relevant cost. [S03] [S20] [S25]

10.4 PRICE THE BUFFER AND REWORK PENALTY OF AN UNRELIABLE QUOTE

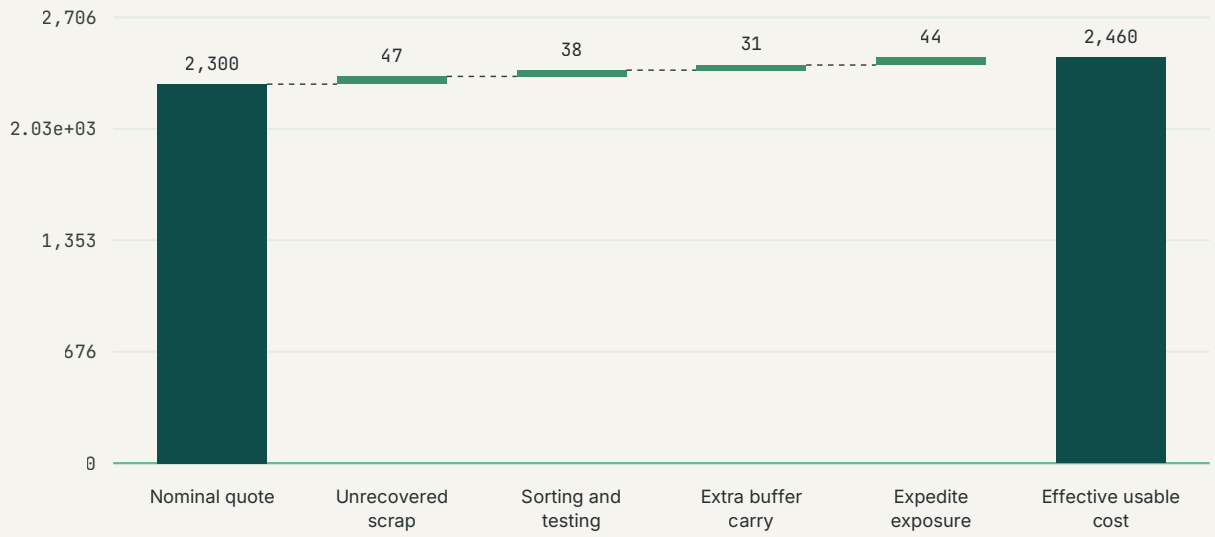
Among qualified suppliers, procurement should evaluate quotes using landed usable cost rather than ex-works invoice price. MIT's material on EOQ extensions and discounts shows that evaluating a quote requires adding ordering, holding and lot-size costs to the purchase price. [S03] When reliability varies across qualified bidders, GFT adds four supplier-driven inventory costs: unrecovered defective units, sorting and re-inspection labour, carrying cost of extra safety stock needed for a wider lead-time tail, and expected expedite exposure. [S03] [S05] [S07]

Consider an illustrative comparison between two qualified cotton-twill suppliers for a workwear programme requiring 24,000 usable metres per year (100 usable metres per working day across 240 days). Supplier 1 quotes XAF 2,300 per metre, with a 96 percent first-pass usable yield (2 percent net unrecovered loss) and a usable-receipt lead time of mean 20 working days and standard deviation 6 days. Supplier 2 quotes XAF 2,420 per metre, with a 99.5 percent first-pass usable yield (0.25 percent net loss) and a usable-receipt lead time of mean 14 working days and standard deviation 2 days. [S03] [S05]

Using Chapter 5's conditional variance formula with illustrative daily demand mean of 100 metres, standard deviation of 20 metres, cycle-service factor z = 1.645 and annual holding rate of 20 percent, lead-time demand standard deviation is sqrt(E[L] × Var(D) + E[D]^2 × Var(L)). [S05] [S24] Supplier 1's illustrative standard deviation is 606.63 metres, requiring 998 metres of safety stock, whereas Supplier 2's is 213.54 metres, requiring 351 metres, so Supplier 1 requires 647 extra metres of buffer stock. [S05] [S24] The waterfall bridge in **Exhibit 41**, p. 154 reconciles Supplier 1's nominal XAF 2,300 per metre quote to its illustrative effective cost per usable metre. [S03] [S05]

EXHIBIT 41 / ILLUSTRATIVE XAF PER USABLE METRE

Scrap, sorting and buffer costs can erase a nominal price discount



ILLUSTRATIVE [S03] [S05] [S24] Illustrative GFT cost bridge for Supplier 1 on 24,000 usable metres/year. Opening invoice quote is XAF 2,300/m; unrecovered scrap loss adds XAF 47/m; 100 percent re-inspection and shade sorting adds XAF 38/m; extra safety-stock and pipeline carrying cost adds XAF 31/m; expected expedite/stoppage allowance adds XAF 44/m, reaching an effective XAF 2,460/m versus Supplier 2's effective XAF 2,445/m. All figures are teaching assumptions.

If Supplier 2's quote were XAF 2,600 per metre, Supplier 1 would remain cheaper even after pricing the larger buffer. [S03] What **Exhibit 41**, p. 154 establishes is a decision rule: procurement may not claim a purchasing saving by comparing nominal unit quotes while transferring the resulting buffer, inspection and disruption costs into the production and working-capital budgets unpriced. [S03] [S05]

10.5

Order amplification is a system outcome

When an apparel factory places erratic, large-lot orders on its fabric mills and trim vendors, upstream suppliers experience demand variability far exceeding actual garment consumption. In supply-chain theory, this upstream variance amplification is known as the bullwhip effect. [S07] [S24] For an emerging-market manufacturer negotiating lead times and prices with spinning or weaving partners, order amplification is destructive: a mill confronted with long silences followed by urgent surges either holds expensive speculative capacity, quotes defensive lead times or rations supply during peak periods. [S08] [S24]

GFT cautions against treating the bullwhip effect as an automatic local diagnosis. In their 2015 paper on stochastic lead times and the bullwhip effect, Michna, Nielsen and Nielsen demonstrate that assuming constant lead times is unrealistic and that random replenishment lead times interact with forecasting and ordering rules to amplify order variance even when end-item demand is stationary. [S24] Alongside stochastic lead times, classical mechanisms include order batching to meet container

or dye-vat minimums, repeated updating of demand forecasts multiplied by lead-time factors, price-discount buying and shortage gaming when buyers inflate orders to secure scarce mill capacity. [S03] [S07] [S24]

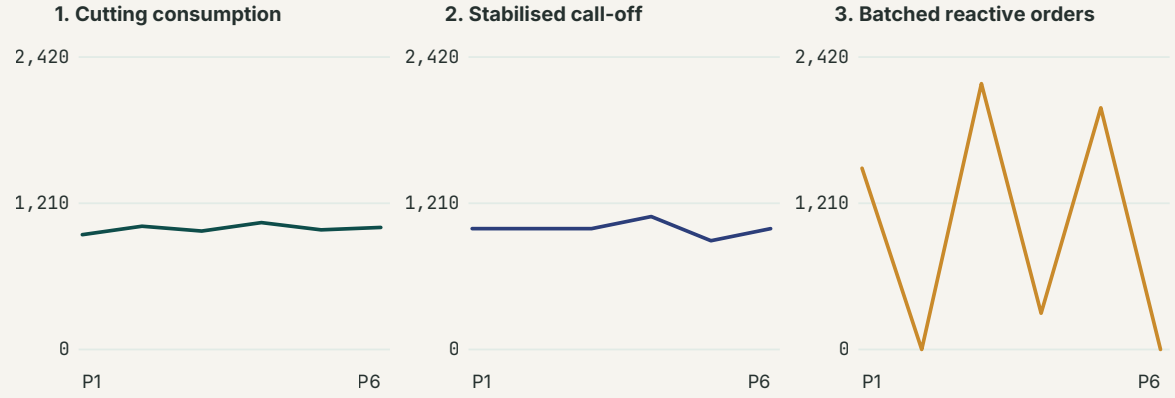
10.6 ALIGN PERIODS, UNITS AND STAGES BEFORE MEASURING AMPLIFICATION

Before a planner claims that order amplification is occurring or recommends a policy change to dampen it, the measurement boundary must be consistent. Comparing weekly garment sales in pieces against monthly fabric orders in kilograms mixes time aggregation, physical conversion factors and bill-of-material netting. [S07] GFT specifies that any empirical test of upstream amplification must express both series in the same material unit, aggregate both over identical calendar buckets, and separate exogenous seasonal shifts from policy-induced oscillation. [S07] [S24]

Suppose, as an illustrative diagnostic, that an apparel maker consumes an average of 1,000 metres per bi-weekly period of standard pocketing fabric across six periods, fluctuating between 950 and 1,050 metres. If the purchasing rule combines lot batching with reactive manual adjustment of the safety factor whenever a shipment arrives late, the order series placed on the mill alternates between zero and surges of 1,500 to 2,200 metres. [S03] [S05] [S24] The constructed sequences in **Exhibit 42**, p. 155 show this mechanism across three panels on a shared scale of 0 to 2,500 metres per bi-weekly period. [S03] [S24]

EXHIBIT 42 / ILLUSTRATIVE METRES PER BI-WEEKLY PERIOD

Batching and delayed information can amplify ordering swings



ILLUSTRATIVE [S03] [S05] [S24] Illustrative constructed sequences over six bi-weekly periods (P1 to P6), not an observed Cameroon factory dataset. All three panels share the same cumulative consumption of 6,000 metres across six periods. Panel 1 is net cutting consumption; Panel 2 applies a stabilised call-off rule; Panel 3 combines lot batching with reactive parameter jumps after a stochastic lead-time delay, illustrating the amplification mechanism analysed in S24.

The operational consequence of **Exhibit 42**, p. 155 is that a garment maker can often reduce its required safety buffer by changing its own ordering grammar rather than demanding that the supplier hold more stock. [S08] [S24] Separating the framework price agreement (covering 6,000 illustrative metres over twelve weeks) from the physical release instruction (dispatching 1,000 metres every two weeks) removes batching amplification at the release boundary. Freezing reorder-point lead-time

parameters against single-shipment noise unless a verified structural shift occurs prevents reactive parameter updating from turning one late truck into a multi-month oscillation. [S05] [S24]

10.7

Shared demand information needs an agreed meaning

Sharing forecasts with suppliers is frequently promoted as a cure for order amplification and shortages, yet unstructured information sharing breeds commercial conflict. Buyers and suppliers routinely attach incompatible meanings to the same spreadsheet. As the 2017 ILO INWORK Policy Brief on global purchasing practices notes, a buyer may treat sample approval or a tentative projection as a routine planning step while a supplier interprets it as the signal to commit raw material, or a buyer may delay confirming material details while holding the delivery deadline fixed. [S10] [S1002] [S1003] In the 2016 ILO Global Survey of 1,454 suppliers across 87 countries, 35 percent of surveyed suppliers reported operating with some unwritten contracts, and among written contracts, only 45 percent specified who was responsible for costs incurred when orders changed. [S10] [S1002] Those survey figures reflect a voluntary multi-country sample from 2016 and do not establish current prevalence in Cameroon, where informal employment across sectors is roughly 90 percent in the inherited register. [R16] [S10] [S1002] They nonetheless highlight a universal engineering hazard: an undefined forecast is an invitation to unallocated loss.

GFT proposes an explicit four-state commitment ladder for all upstream supplier communication. Each state must carry a documented horizon, a permitted revision band, an authorised sender role and an explicit rule stating who owns the financial exposure if downstream demand vanishes. [S08] [S1002]

10.8 DEFINE THE FOUR CONTRACTUAL STATES FROM PROJECTION TO RECEIPT

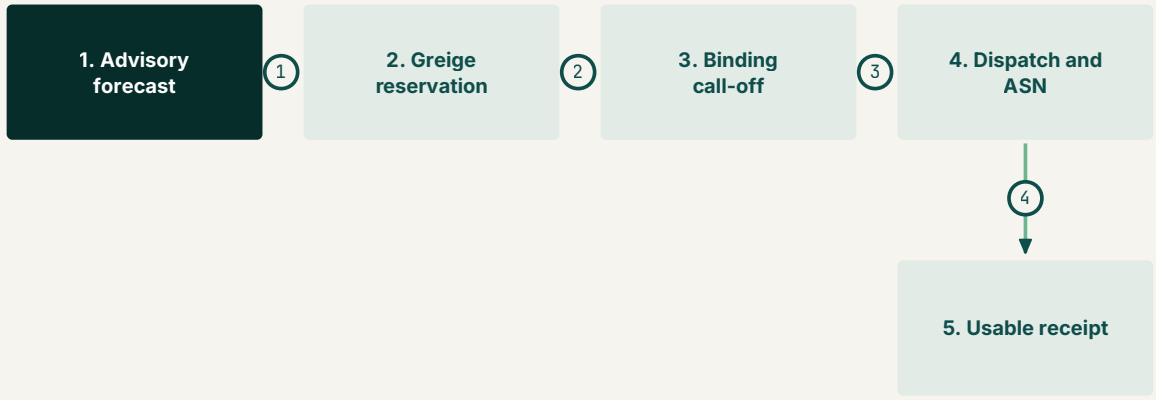
1. **Advisory rolling forecast (non-binding visibility).** The buyer shares a time-phased projection over an extended horizon (for example, an illustrative 12 to 24 weeks) solely to help the supplier plan aggregate mill capacity and fibre procurement. [S07] [S08] It authorises zero material conversion or financial commitment at the buyer's expense.
2. **Capacity or greige reservation (bounded liability).** The buyer reserves a corridor of spinning, weaving or dyeing capacity, or authorises the supplier to hold undyed greige fabric or yarn, in exchange for a stated reservation fee or bounded cancellation liability. [S08] [S1001]
3. **Binding production call-off (firm purchase commitment).** The buyer issues an authorised order specifying item code, shade lot, quantity, packaging, delivery window and price. [S07] [S08] Any buyer-initiated reduction, revision or schedule compression inside the frozen window triggers a pre-agreed cost-reimbursement and lead-time-extension formula. [S1002]

4. **Verified usable receipt (reconciled settlement event).** Physical arrival, quantity verification, quality release and document matching close the delivery obligation and establish the invoice acceptance date for payment or supply-chain finance. [S18] [S31]

The process sequence in **Exhibit 43**, p. 157 maps these four states and the verification gates that separate them. Each transition changes commercial liability and inventory ownership; no planner or script may skip a gate or treat a forecast upload as a binding call-off. [S08] [S31]

EXHIBIT 43 / PROPOSED CONTRACTUAL STAGES

A forecast becomes a commitment only at a defined contractual event



1. Approve reservation cap
2. Confirm shade and lot
3. Inspect pre-ship lot
4. Verify count and quality

ILLUSTRATIVE [S08] [S31] [S1001] [S1002] Illustrative GFT purchasing-communication protocol. Progression from non-binding visibility to capacity reservation, binding call-off and usable receipt requires explicit contractual triggers and cost-allocation rules informed by S08, S31, S1001 and S1002.

If a supplier shares its available-to-ship stock balance, the garment planner must know whether that figure represents physical finished rolls reserved exclusively for the buyer or unallocated mill stock that could be sold to another customer before the call-off is issued. [S07] [S31] GFT's digital-strategy boundary requires that shared supplier balances carry a timestamp, reservation status and expiry window before they are counted as firm upstream supply. [S07] [S31]

10.9

Flexibility contracts allocate rather than remove risk

When demand is uncertain across styles, colours or tender awards, a rigid wholesale contract forces one party to carry disproportionate overage or underage exposure. As Yossi Sheffi's MIT OpenCourseWare lecture on supply contracts demonstrates, under a simple wholesale-price contract in a single-period setting, the buyer bears the entire overage risk of unsold stock while earning only a fraction of the total channel margin, causing the buyer to order less than the quantity that maximises combined supply-chain profit. [S08] [S1001] Supply-chain contracting literature analyses mechanisms such as buyback contracts, revenue-sharing agreements, quantity-flexibility clauses and real-option contracts that redistribute overage and underage risk between supplier and buyer. [S08] [S1001]

For an emerging-market apparel factory, GFT insists on two practical qualifications. First, a flexibility contract never eliminates physical inventory risk; it reallocates who pays when actual demand diverges from the initial reservation and changes the stage where materials are committed. [S08] [S1001] Second, a contractual promise of a buy-back refund, restocking credit or reserved surge capacity is worth no more than the counterparty's liquidity, solvency and physical ability to perform when the option is exercised. [S08] [S18]

10.10 COMPARE WHOLESALE, BUYBACK AND CAPACITY-OPTION STRUCTURES

In the MIT lecture formulation, a buyback contract pairs a wholesale price with a buy-back credit per unused unit returned, raising the buyer's effective salvage value and encouraging a larger initial order while transferring part of the overage loss back to the supplier. [S08] [S1001] In garment procurement, physical return of cut or custom-dyed fabric is rarely economic once transport, customs and shade-lot specificity are considered; buyback clauses are therefore practical mainly for standard undyed yarns, core pocketing, sewing threads and generic packaging. [S07] [S08]

By contrast, a real-option or capacity-reservation contract separates the upfront commitment to reserve undyed material or machine time from the final exercise decision to dye and finish specific colours. [S08] [S1001] In the MIT real-options model, the buyer purchases `Q` call options at a unit reservation price `w`, the supplier prepares `Q` units of upstream capacity or base material, and as demand materialises, the buyer exercises up to `Q` units for an additional exercise price `E`. [S1001] This maps directly onto postponed fabric differentiation: the garment maker pays `w` for greige fabric held at the mill and pays `E` only for the dyed shade breakdown confirmed at call-off. [S08] [S1001]

WORKED EXAMPLE: COMPARING RIGID WHOLESALE AGAINST A GREIGE CAPACITY OPTION

Consider an illustrative uniform programme with fabric demand across three scenarios: Low (4,000 metres, illustrative probability 0.25), Medium (6,000 metres, probability 0.50) and High (8,000 metres, probability 0.25), giving expected demand of 6,000 metres. [S08] [S1001] Accepted fabric yields an illustrative net contribution (before fabric cost) of `R = XAF 4,500` per metre, unused dyed salvage is `S = XAF 800` per metre, and mill cost is XAF 1,400 per metre in greige plus XAF 700 for dyeing (`C = XAF 2,100` total, with unused greige recoverable at XAF 1,100 per metre). [S08] [S1001]

Under **Contract 1 (Rigid Wholesale)** at `W = XAF 2,800` per metre paid upfront, underage cost is `XAF 1,700`, overage cost is `XAF 2,000` and the critical fractile is `0.459`, making 6,000 metres optimal (`XAF 16,800,000` upfront). [S08] [S1001] Illustrative net contribution is `XAF 2,800,000` in Low and `XAF 10,200,000` in Medium and High (leaving 2,000 High-scenario metres unserved), for an expected net contribution of `XAF 8,350,000`. [S08] [S1001]

Under **Contract 2 (Greige Option and Call-Off)**, the buyer reserves 8,000 illustrative metres of greige at `w = XAF 450` per metre (`XAF 3,600,000` upfront) and pays `E = XAF 2,450` per dyed metre called off. [S08] [S1001] Illustrative net contribution is `XAF 4,600,000` in Low (zero obsolete dyed stock), `XAF 8,700,000` in Medium and `XAF 12,800,000` in High, raising expected net contribution to `XAF 8,700,000` while cutting upfront cash from `XAF 16,800,000` to `XAF 3,600,000`. The mill's expected profit also rises from `XAF 4,200,000` to `XAF 4,250,000` because unexercised greige retains `XAF 1,100` recovery without wasting the `XAF 700` dyeing cost, provided post-exercise dyeing fits the delivery window and reserved greige is physically held. [S08] [S18] [S1001]

10.11 TEST COUNTERPARTY SOLVENCY AND WORKING-CAPITAL TIMING

Contract design cannot be separated from supply-chain finance and counterparty verification. As the International Finance Corporation's Global Trade Supplier Finance (GTSF) programme documentation explains, post-shipment supplier finance depends on the acceptance of receivables by approved buyers after goods are shipped and accepted. [S18] A post-shipment receivables facility does not automatically finance a supplier's upstream purchase of raw cotton or yarn before manufacture, nor does it guarantee that an undercapitalised supplier can honour a buyback credit at the end of a weak season. [S08] [S18]

GFT requires three safeguards whenever a contract includes deposits, reservation fees, rebates or buybacks: upfront fees must match verified physical progress such as tagged greige rolls in an audited warehouse; buyback credits must settle via escrow, bank guarantee or payable offset rather than unsecured refund promises; and while WP13 owns the credit structure, the planner must stress-test buffers under counterparty default. [S08] [S18] [S31]

10.12

A late change must not be financed by harmful working practices

In apparel supply chains, inventory buffers, purchasing timelines and workplace conditions are tightly coupled. When a buyer compresses lead times, issues late technical specifications, changes colourways after production slots are booked or demands price reductions that ignore statutory wage increases, the buyer transfers its own planning deficit onto the supplier's workforce. [S10] [S11] [S1002] [S1003] GFT's position is unequivocal: an inventory policy that achieves lower buyer stock by forcing uncosted overtime, casualised surge labour, unauthorised subcontracting or below-cost production at the supplier is an engineering failure, not an efficiency gain. [S20] [S25]

International institutional research documents these mechanisms while requiring attention to each study's sample boundaries. In the 2017 ILO INWORK Policy Brief reporting on the 2016 Global Survey of 1,454 suppliers across 87 countries, only 17 percent of surveyed suppliers considered their orders to have sufficient lead time, while the majority reported that more than 30 to 50 percent of their orders carried insufficient lead times. [S10] [S1002] In the same survey, 72 percent stated that inaccurate technical specifications from buyers led to additional sampling time and cost, 66 percent reported increased production costs and 50 percent reported financial losses. [S1002] Furthermore, 55 percent reported accepting orders below production cost (81 percent of those in textiles, clothing, leather and footwear doing so to secure future orders), only 25 percent of customers adjusted prices for statutory minimum-wage increases, and suppliers waited 12 weeks on average before incorporating such increases into prices. [S1002]

Complementing that survey, Mark Starmanns's 2017 ILO working paper examined 14 garment brands and retailers and documented how internal silos undermine fair purchasing: a brand's buying department frequently exerts pressure for faster delivery and lower Free on Board (FOB) prices while its compliance team asks suppliers to raise wages and respect working-hour limits. [S11] [S1003] Similarly, the OECD's 2018 guidance and 2026 due-diligence spotlight emphasise that enterprise due diligence must address the enterprise's own purchasing practices and subcontracting risks. [S20] [S25]

The evidence summary in **Exhibit 44**, p. 161 records these four institutional sources alongside their exact dates, sample scopes and analytical limitations so that planners use them honestly: they justify hard purchasing safeguards without pretending to be a current statistical census of Cameroon garment workshops. [S10] [S11] [S20] [S25] [S1002] [S1003]

EXHIBIT 44 / PUBLISHED INSTITUTIONAL STUDIES

Purchasing evidence supports safeguards without proving local prevalence

Source ID	Publisher and date	Empirical or normative scope	Operational finding for purchasing design	Limitation and boundary
S10 / S1002	ILO (2017-06-09)	2016 survey of 1,454 suppliers in 87 countries	35% had unwritten contracts; 45% allocated order-change costs; 17% had sufficient lead time	Voluntary SEDEX/ETI sample; cross-sectional associations, not current Cameroon prevalence
S11 / S1003	ILO (2017-05-15)	Qualitative study of 14 garment brands and retailers	Buying teams pressure speed and FOB cost while compliance teams ask for higher wages	Small qualitative SME brand sample; illustrates governance conflict, not causal effect sizes
S20	OECD (2018-03-07)	Sector due-diligence guidance aligned with UN principles	Enterprises must assess how purchasing, pricing and lead-time choices create adverse impacts	Normative institutional framework; requires local operational rules and legal review
S25	OECD (2026-02-26)	Garment and footwear due-diligence spotlight	Subcontracting is common practice; visibility must extend beyond direct tier-one suppliers	Sector overview; does not map specific Cameroon subcontractors or certify local vendors

SOURCED [S10] [S11] [S20] [S25] [S1002] [S1003] Sourced comparison of institutional evidence on purchasing practices and garment supply-chain due diligence. Each source's sample, publication date and scope limitation are preserved explicitly.

10.13 ENFORCE CHANGE-CONTROL, WAGE PASS-THROUGH AND OVERTIME BOUNDARIES

In Cameroon, the inherited evidence register records that the statutory minimum wage (SMIG) rose from XAF 41,875 per month in 2023 to XAF 43,969 per month from 23 February 2024 (up 5 percent), and that the standard working week is 40 hours with a 48-hour maximum and overtime paid at 1.2 times the normal rate, within an economy where informal employment is about 90 percent. [R14] [R16] Whether a Cameroon garment maker buys local textile services or manufactures for an institutional client, GFT proposes binding purchasing controls so inventory miscalculations are never balanced on workers' backs. [R14] [R16] [S20]

Every supply contract must be written and state the frozen window, the cost schedule for buyer revisions, and an automatic day-for-day delivery extension whenever a buyer delays sample, shade or trim approval. [S1002] [S1003] If a buyer requests a late change without extending delivery, the planner must verify regular-hour capacity within the 40-hour standard and 48-hour maximum week (or applicable law), and fund the statutory overtime premium (1.2 times in Cameroon's baseline) or authorise a split delivery. [R16] [S20] Framework contracts must also index labour-linked prices to statutory minimum-wage revisions immediately rather than after the 12-week

delay reported in the ILO survey, and Chapter 13 scorecards must penalise buyers whose unit-price discounts trigger late payments, unapproved subcontracting or quality holds. [R14] [S20] [S1002] [S1003]

> **GFT NOTE**
>
> Never allow a purchasing KPI to reward unit-price reductions in isolation from lead-time realism, first-pass usable yield and labour safeguards. GFT recommends three mandatory clauses in every garment-input contract: (1) a day-for-day delivery extension for any buyer delay in sample, shade or credit approval; (2) explicit buyer funding of verified re-sampling, material write-off and statutory overtime premiums when the buyer alters a confirmed call-off; and (3) immediate automatic adjustment of labour-linked conversion rates when statutory minimum wages change. [R14] [R16] [S20] [S1002] This is an operational governance recommendation, not legal advice; contract clauses require review by qualified counsel.

10.14

Alternate sourcing needs qualification before an emergency

When a primary fabric mill or trim vendor suffers a breakdown, port delay or financial dispute, procurement teams often point to a list of alternative vendor names as proof of resilience. GFT draws a strict line between a contact list and a qualified alternate source. Switching a fabric or trim supplier mid-programme without prior technical and operational qualification produces shade mismatches, shrinkage variations during washing, needle breakage, missing provenance records or unverified subcontracting. [S07] [S20] [S25] Until an alternate supplier has passed specification matching, physical test-lot conversion, due-diligence screening and lead-time measurement, its effective emergency capacity during a short-horizon disruption is zero. [S20] [S25] [S30]

Even individually qualified suppliers do not automatically provide independent risk reduction. As Chapter 9 established for network pooling and as the OECD's 2026 spotlight underscores for garment subcontracting, two apparent competitors may rely on the same upstream spinning mill, dyehouse, grid feeder or port corridor. [S24] [S25] In Cameroon, where policy aims to expand domestic cotton processing from roughly 5 percent toward 50 percent by 2030 alongside announced restructuring and integrated textile projects, local suppliers may share upstream fibre, ginning, spinning or power dependencies, while imported alternatives share shipping lanes and customs bottlenecks. [R06] [R07] [R08] [R09] [S30]

10.15 QUALIFY SECOND SOURCES ACROSS FIVE OPERATIONAL GATES

To determine whether a secondary supplier can reduce the safety buffer held against primary-supplier disruption, GFT proposes a five-gate readiness protocol:

1. **Technical equivalence.** The alternate item matches approved construction, yarn count, weight, width, shrinkage tolerance, colourfastness and shade standard, with formal substitution approval on the bill of materials. [S07]

2. **Physical trial lot and line run.** The supplier has delivered at least one paid trial lot passing incoming inspection, cutting, sewing and wash testing, with first-pass yield and usable-receipt timestamps recorded. [S07] [S30] [S31]
3. **Traceability and due-diligence clearance.** Facilities and permitted subcontractors are disclosed, screened under OECD due-diligence principles and linked to WP05 traceability requirements. [S20] [S25]
4. **Common-cause independence mapping.** Upstream yarn and fabric sources, energy dependencies, transport routes and customs entry points are mapped against the primary supplier to identify shared disruption modes. [S24] [S25] [S30]
5. **Commercial and activation readiness.** A valid framework agreement, pack size, payment mechanism and activation lead time to usable receipt are active in the current planning cycle. [S08] [S18]

> **WARNING**
>
> Do not reduce primary safety stock on the basis of an unqualified backup vendor or a dual-source split where both vendors share the same bottleneck. If Primary Supplier A and Backup Supplier B both import greige fabric through the same port queue or subcontract dyeing to the same finisher, a corridor delay or dyehouse outage halts both streams simultaneously. [S24] [S25] [S30] Safety-stock reductions from dual sourcing are admissible only for the portion of lead-time and yield variance proven independent between the two qualified sources. [S05] [S24]

10.16 CONNECT PURCHASING TERMS TO THE GFT TOOL AND PAPER ARCHITECTURE

These purchasing rules feed directly into GFT's tools and companion papers. In `unit-economics`, **Exhibit 41**, p. 154 requires entering first-pass usable yield, re-inspection cost and buffer carry alongside nominal quotes. [S03] [S05] In `safety-stock`, usable-receipt lead-time distributions and censored-order flags set the protection horizon, refusing buffer reductions for unqualified backups. [S05] [S22] [S24] In `risk-matrix`, **Exhibit 40**, p. 152 and common-cause checks prevent labour or subcontracting exposures from being averaged away. [S20] [S25] Across the library, WP05 owns traceability; WP09 owns supervisor and inspector training; WP13 owns letters of credit, deposits and receivables finance (S18); and Chapter 11 next examines how working-capital ceilings, deposit dates and currency exposures constrain the feasible order plan. [S08] [S18] [S20]

A low price that hides a late, defective or harmful delivery is only a loan from the factory's buffer and its workers.

CHAPTER SOURCE NOTES

S03 MIT OpenCourseWare. 2006.

S10 International Labour Organization. 2017-06-09.

S20 OECD. 2018-03-07.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S07 MIT OpenCourseWare. 2006.

S05 MIT OpenCourseWare. 2006.

S30 World Customs Organization. 2025.

S31 GS1 US and Auburn University RFID Lab. 2018.

S22 arXiv, Chen, Jiang, Zhang and Zhou. 2023-10-31.

S25 OECD. 2026-02-26.

S24 arXiv, Michna, Nielsen and Nielsen. 2015-03-20.

S08 MIT OpenCourseWare. 2006.

S1002 International Labour Organization. 2017-06-09.

S1003 International Labour Organization. 2017-05-15.

R16 WageIndicator. Publication date unrecorded; inherited source-text check.

S1001 MIT OpenCourseWare. 2006-11-09.

S18 International Finance Corporation. 2026.

S11 International Labour Organization. 2017-05-15.

R14 WageIndicator. Publication date unrecorded; inherited source-text check.

R08 Business in Cameroon. Publication date unrecorded; inherited source-text check.

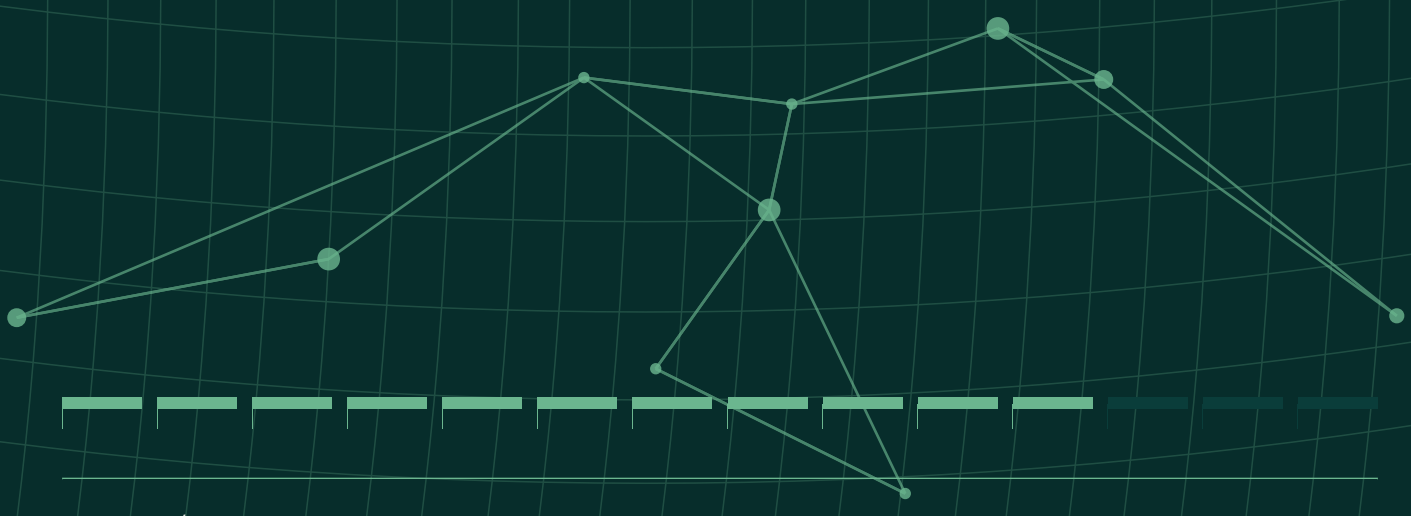
R09 Ecofin Agency. Publication date unrecorded; inherited source-text check.

11

PART III: RESOLVE CONSTRAINTS ACROSS THE SUPPLY CHAIN / WORKING CAPITAL

The feasible stock plan is bounded by cash, not just by statistical service

A decision that reduces expected shortages may still be unaffordable once deposits, transit stock and receivables are included.



CHAPTER 11
DECISION TEST

A decision that reduces expected shortages may still be unaffordable once deposits, transit stock and receivables are included.

A stock plan is feasible only when the business can pay every obligation on the date it falls due while preserving an approved cash reserve. A statistically justified buffer can fail that test. It can require a supplier deposit before production, a balance payment while fabric is still in transit, conversion expenditure before shipment and further waiting before the customer pays. Green Future Tech Limited (GFT), Hong Kong SAR, therefore treats cash timing as a constraint on the order itself, not as a finance review after planning has finished. The feasible policy must show the material quantities, the customer promise, the payment calendar, the unmet requirement and the approval for any exception on one decision record. This chapter sets out that method for an apparel maker. No factory ledger, supplier quotation, bank offer or funding approval was supplied. All monetary cases and thresholds below are illustrative teaching models, not GFT operating results, current market rates or investment, legal, accounting or tax advice.

11.1

Map cash dates across the material-to-payment cycle

The cash position of an order differs from its inventory position. Inventory position asks which usable stock and firm receipts can support demand. Cash position asks which payments have already occurred, which are committed and which receipts can legally and practically be collected. A fabric roll can consequently appear as inbound supply before it is usable, while the deposit and supplier balance attached to it have already left the business. Conversely, an accepted customer invoice can appear as an asset without yet providing cash for the next cutting order. Planning against only one of these positions creates a false feasibility result.

GFT would construct a dated obligation ledger before releasing a material order. Each row should identify the commercial commitment, counterparty, amount, currency, due date, expected settlement date, related purchase-order line, stock state and accountable role. The ledger should distinguish a supplier deposit, the supplier balance, carriage, insurance where applicable, duties and clearance charges where applicable, receiving and inspection costs, conversion wages and overheads, customer invoicing, invoice acceptance and final collection. Estimated dates must remain labelled as estimates until an authorised event confirms them.

Contract terms help allocate these events, but they do not replace the payment calendar. ICC Academy explains that the Incoterms rules clarify delivery, transfer of risk and allocation of costs between seller and buyer. [S1101] That makes the rule edition and named place necessary inputs to landed-cost planning. It does not tell the factory when a deposit is due, when title passes or when a bank will settle a foreign-currency payment. Those terms require separate contractual and banking evidence. An order can also carry transit risk before the material reaches the factory, so physical location, commercial responsibility and cash exposure must remain separate fields. [S1101]

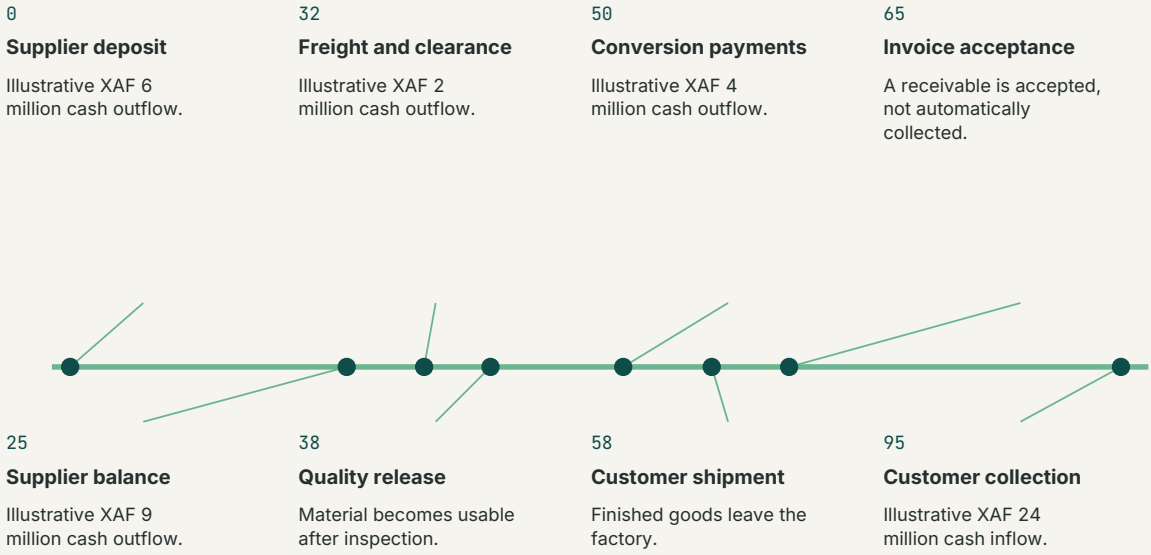
Goods in transit and quality-held stock deserve particular attention. IAS 2 describes inventory cost as including purchase, conversion and other costs incurred to bring inventories to their present location and condition. [S09] The Project Zipper study separately examined ship and receive records and found discrepancies between trading-

partner records, reinforcing the need to reconcile informational and physical events rather than treating a shipping notice as a usable receipt. [S31] Neither source supplies a Cameroon apparel cash cycle. Their useful boundary lesson is that cash can already be attached to material that cannot yet be promised to production.

Exhibit 45, p. 167 constructs this sequence for a uniform order. Every date and amount is a teaching assumption. The example begins with an approved order, not with a claim that funding is available.

EXHIBIT 45 / ILLUSTRATIVE CALENDAR DAYS FROM ORDER APPROVAL

Cash leaves before the garment can generate a receivable



ILLUSTRATIVE [S18] [S1101] Illustrative teaching sequence. Deposit XAF 6 million on day 0; supplier balance XAF 9 million on day 25; freight and clearance XAF 2 million on day 32; conversion payments XAF 4 million on day 50. Shipment occurs on day 58, invoice acceptance on day 65 and customer collection of XAF 24 million on day 95. Total pre-collection outflows are XAF 21 million. Dates, prices and customer terms are invented to demonstrate timing, not observed lead times or quotes.

WORKED EXAMPLE: THE CASH PEAK IS NOT THE PURCHASE PRICE

In the illustrative sequence, cumulative cash outflows before collection are XAF 6 million + XAF 9 million + XAF 2 million + XAF 4 million = XAF 21 million. The supplier invoice represented by the first two payments is XAF 15 million, but the funding need is wider because freight, clearance and conversion also precede collection. The deposit-to-collection interval is an illustrative 95 calendar days. If the business begins with only XAF 18 million of uncommitted cash available to this order and requires no reserve, the sequence still has an illustrative XAF 3 million funding gap before any other factory obligation is considered. [S18]

The important calculation is the maximum cumulative net outflow by date, including the reserve that management will not spend. Average monthly cash balances can miss the deposit-day breach. A weekly or monthly reporting total can also hide the order in which obligations occur. GFT would recalculate the curve after every material change to price, payment terms, promised shipment, acceptance status or collection expectation, while retaining the previous version and the reason for change.

11.2

Stock value, cash released and cost saved are different outcomes

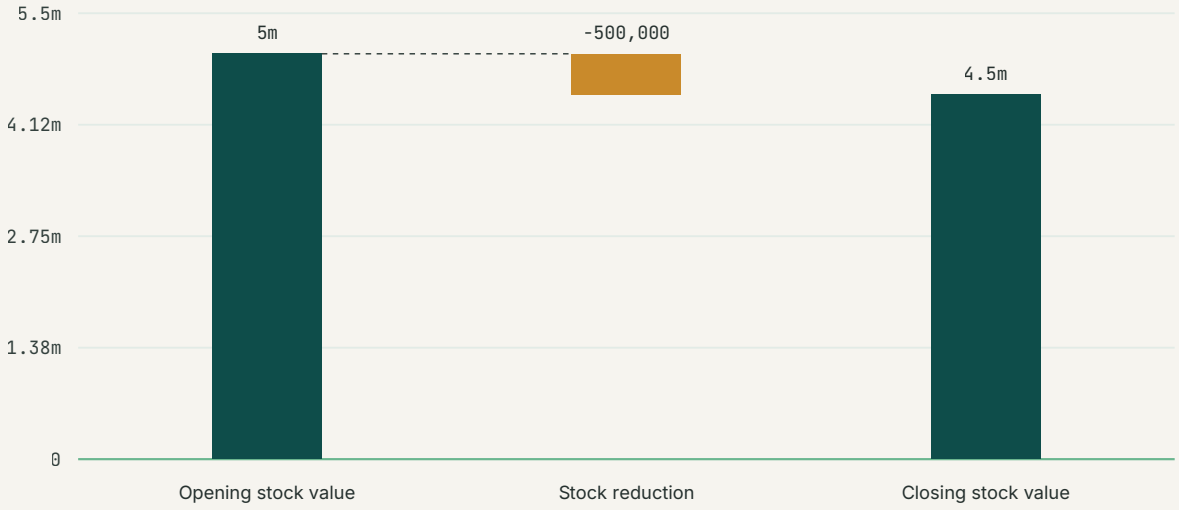
A reduction in inventory can produce four different reported outcomes that should never be added together as if they were the same benefit. Balance-sheet stock reduction changes the carrying amount of an asset. Cash release changes available funds when stock is consumed, sold, returned or otherwise converted into cash under actual payment terms. Recurring carrying-cost reduction lowers selected future costs associated with holding stock. A write-down recognises an expense when the carrying amount exceeds net realisable value under the applicable accounting treatment. IAS 2 requires inventories to be measured at the lower of cost and net realisable value and recognises write-downs and losses as expenses in the period in which they occur. [S09]

The decision-model carrying charge has another purpose. MIT's inventory teaching separates purchase, ordering, holding and shortage costs in total relevant cost. [S02] A local holding estimate can include the opportunity cost of capital, storage handling, insurance, deterioration and obsolescence where these genuinely vary with the stock decision. It should not import an accounting stock balance and call the whole balance an annual expense. Nor should it count the same financing charge, warehouse cost or expected loss twice under two labels.

Exhibit 46, p. 169 isolates the asset movement before discussing annual savings. It uses an illustrative reduction of 100 units carried at XAF 5,000 each. The arithmetic is `100 × XAF 5,000 = XAF 500,000` of stock value. If the teaching carrying-rate assumption is 20 percent a year, the associated annual carrying exposure is `XAF 500,000 × 20 percent = XAF 100,000` before tax, service, capacity, terminal-stock and implementation effects. [S02] [S09] Neither amount is automatically cash released, and the annual amount is not an addition to the closing stock balance.

EXHIBIT 46 / ILLUSTRATIVE XAF CARRYING VALUE

A stock reduction is not equal to a recurring profit gain



ILLUSTRATIVE [S02] [S09] Illustrative reconciliation: opening stock of 1,000 units at XAF 5,000 has carrying value XAF 5,000,000. Removing 100 units at the same assumed value reduces stock by XAF 500,000 to XAF 4,500,000. A separate illustrative 20 percent annual carrying assumption on the reduction gives XAF 100,000 per year, not a step in this asset-balance waterfall. Actual cash release depends on sale, consumption, payables and collection events. All inputs are teaching assumptions.

The cash effect depends on what happens next. Consuming existing fabric instead of placing a new order can avoid a future payment. Selling finished goods can create a receivable rather than immediate cash. Returning material can create a supplier credit whose settlement date and enforceability need confirmation. Writing down obsolete fabric changes its carrying amount but does not by itself recover money. These are distinct events with distinct evidence. A saving claim should identify the avoided invoice, the collection or the reduced future cost that forms its basis.

Changed ownership or payment terms require the same discipline. A supplier may retain ownership of stock at the factory, offer a buyback or permit later payment. Supply-chain contract teaching shows why such terms change incentives and allocation of overage risk; they do not make physical stock or demand uncertainty disappear. [S08] GFT would therefore report physical quantity, ownership, cash paid and cash committed as separate measures. Moving stock to a supplier's balance sheet can improve the factory's reported inventory while increasing the supplier's financing need or price. That is a contractual redistribution, not proof of higher material productivity.

A credible comparison also needs a common counterfactual. The baseline and candidate should use the same demand path, customer promise, stock boundary, observation window and treatment of residual material. The candidate cash curve should include implementation expense and any additional expediting or handling. The comparison should then state four results independently: change in stock carrying value, change in peak cash requirement, change in periodic relevant cost and change in accounting expense where known. Combining them into one headline would conceal the mechanism that produced the result.

11.3

A cash-constrained allocation must reveal unmet demand

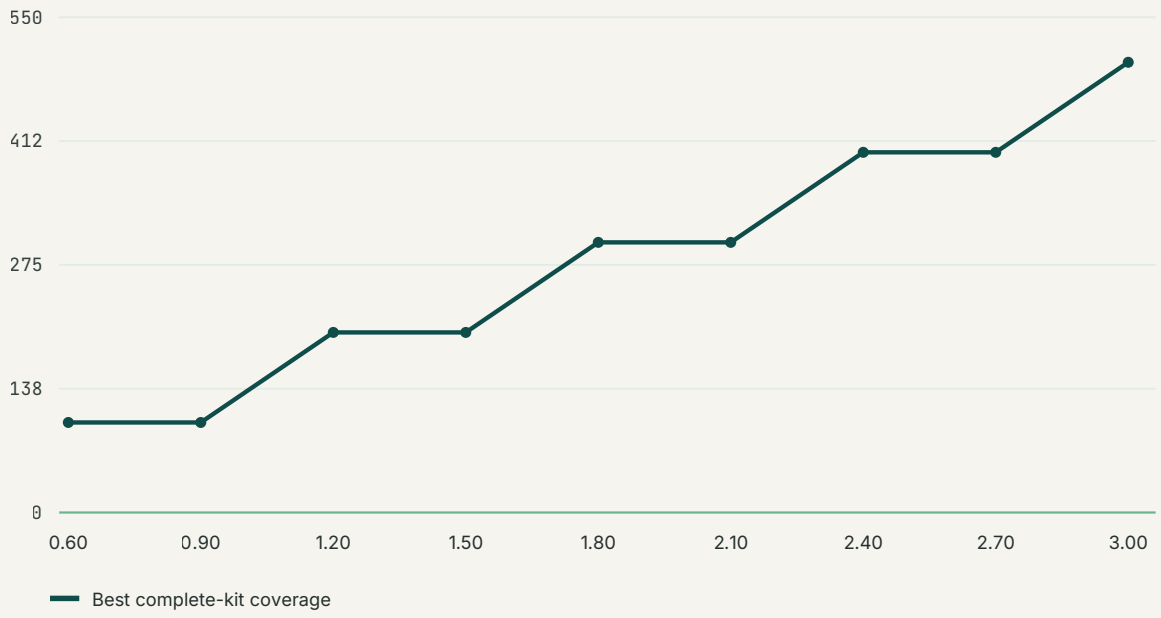
A statistical service calculation can recommend additional stock without determining which other material order must be deferred to pay for it. That omission matters when several products share one cash pool. The decision is then an allocation problem across competing orders, not an independent reorder calculation for each item. The appropriate output is a feasible combination, the demand that remains uncovered and the constraint responsible. A solver should not silently reduce the customer promise to make its answer look feasible.

GFT would use integer order quantities where supplier packs, minimum lots or indivisible production commitments require them. For each material `i`, let `q_i` be a non-negative integer number of packs, `c_i` the cash cost per pack at the relevant payment date, `p_i` the storage positions consumed per pack and `u_i` the complete-kit coverage contributed per pack. With budget `B` and storage limit `P`, the simple teaching model maximises covered kits `K` subject to `Σ c_i q_i ≤ B`, `Σ p_i q_i ≤ P` and `K ≤ u_i q_i` for every required component. A minimum service requirement `K ≥ K*` can be added as a constraint rather than assumed away. Lot-sizing and service literature supplies the cost and policy foundations; the cash, space and complete-kit rules here are GFT's proposed operational extension. [S03] [S05]

> **WORKED EXAMPLE: PACKS MAKE THE SERVICE FRONTIER LUMPY**
>
> Consider an illustrative confirmed order for 500 uniform kits requiring three components. Fabric pack A costs XAF 300,000 and covers 100 kits. Trim pack B costs XAF 120,000 and covers 200 kits. Interlining pack C costs XAF 180,000 and covers 100 kits. Every kit needs all three components, so coverage is `K = min(100q_A, 200q_B, 100q_C)`. These are constructed teaching values, not supplier quotations, bills of material or a forecast. A cash-only complete enumeration from zero to twenty packs per component gives the frontier in **Exhibit 47**, p. 171. [S03] [S05]

EXHIBIT 47 / MAXIMUM COMPLETE KITS FROM ENUMERATED ILLUSTRATIVE PACKS

A tighter cash ceiling can force a service trade-off



ILLUSTRATIVE [S03] [S05] Illustrative cash-only enumeration. Labels are cash ceilings in XAF millions: 0.60, 0.90, 1.20, 1.50, 1.80, 2.10, 2.40, 2.70 and 3.00. Pack A costs XAF 300,000 for 100 kits; pack B costs XAF 120,000 for 200 kits; pack C costs XAF 180,000 for 100 kits. Coverage is the minimum component coverage. Complete enumeration from zero to twenty packs per component was used. No statistical service probability or realised saving is claimed.

At an illustrative XAF 2.40 million ceiling, the best affordable combination is four A packs, two B packs and four C packs, costing XAF 2.16 million and covering 400 complete kits. Spending the remaining XAF 240,000 cannot buy a feasible A or C pack, so aggregate budget headroom does not translate into another complete kit. Against the illustrative 500-kit commitment, the plan therefore has 100 kits unmet. At XAF 2.70 million, coverage remains 400 because the least-cost combination reaching 500 requires five A packs, three B packs and five C packs at XAF 2.76 million. [S03] [S05]

The result changes again when space is binding. In the same illustrative model, assign one storage position to every pack and impose a limit of twelve positions. The least-cost 500-kit combination needs thirteen positions: five A, three B and five C. No reduction in any pack count can retain 500-kit coverage, so the saved constraints are infeasible. The correct output is not a 500-kit plan with a warning in small print. It is an exception stating that cash and storage cannot support the complete promise, together with the nearest inspected alternatives and their uncovered quantities. [S03]

A planner can respond to that exception in several ways, but each response changes an explicit business assumption. Management can increase the dated cash limit, approve additional storage, negotiate smaller packs, split delivery dates, approve a technically valid substitution or renegotiate the customer quantity or date. It can also accept a lower covered quantity and record the commercial consequence. None of these is a mathematical repair performed without authority. The selected action needs an owner, an effective date and a revised cash curve.

The frontier also prevents a common misreading of budget value. Moving from an illustrative XAF 0.60 million ceiling to XAF 0.90 million adds no complete-kit coverage in this constructed model, while moving from XAF 2.40 million to XAF 2.76 million adds 100 kits. This does not establish a universal return on working capital. It shows why an incremental cash request should identify the particular pack or constraint that changes the feasible set. The value of extra cash is discontinuous when orders are indivisible.

> **METHOD**
>
> For an operational allocation, GFT would save the candidate quantities, payment-dated cash cost, pack definitions, storage use, component-coverage equations, customer minimum, solver or enumeration status and every unmet requirement. Independent checks should re-add the cash by date, verify each pack multiple and recalculate complete-kit coverage. If the model is infeasible, the exception should name the conflicting constraints. If it is feasible, the report should still show service below the original request and any unused cash that cannot buy a valid pack. [S03] [S05]

11.4

Currency exposure belongs in scenario costs

A foreign-currency purchase creates a dated translation exposure as well as a material commitment. The order record should preserve the invoice currency, invoice amount, payment dates, reporting currency and the source and date of every conversion rate. Reporting in CFA franc (XAF) does not turn a dollar or euro obligation into a fixed XAF cash requirement before the applicable settlement terms are known. Bank charges, intermediary fees, settlement delays and any difference between a reference rate and the rate actually available also belong in the cash scenario.

The inherited evidence register records source-annex context that the CFA franc is pegged to the euro at 655.957 XAF per euro. [R32] That entry is not a current US dollar conversion and does not establish the rate, spread or fee a factory can obtain from its bank. GFT would not reuse it to price a dollar payable. An operational plan needs dated supplier terms and banking inputs. Where those inputs are absent, the analysis should present explicitly labelled scenarios rather than a precise but unsupported landed-cost figure.

For an illustrative sensitivity only, suppose a supplier invoice is USD 25,000 and management tests invented settlement rates of XAF 550 and XAF 605 per US dollar. The corresponding XAF equivalents are `USD 25,000 × 550 = XAF 13,750,000` and `USD 25,000 × 605 = XAF 15,125,000`, a difference of XAF 1,375,000 before any bank charges. These rates are teaching assumptions, not observed prices, forecasts or a recommended hedge. The inherited euro-peg context does not support either dollar scenario. [R32]

That difference can change the order rather than merely the reported cost. If the illustrative foreign-currency scenario adds XAF 1,375,000 to a plan whose remaining approved headroom is XAF 1,000,000, the plan exceeds its limit by XAF 375,000 before fees. The response must be a new feasible decision: approve more cash,

change the quantity or payment schedule, renegotiate the commercial terms or decline the commitment. Presenting the higher scenario as an explanatory note beneath an unchanged purchase order would hide the breach.

Contractual delivery terms do not settle this currency question. Incoterms guidance separates delivery, risk and cost responsibilities between buyer and seller, but the sales contract and banking records still need to establish invoice currency, payment schedule and settlement mechanics. [S1101] GFT would link each rate scenario to the relevant payment event rather than translate the entire order once at an arbitrary review-date rate. Deposits and balances paid on different dates can face different available rates and fees.

This chapter makes no investment or hedging recommendation. Its narrower requirement is an inspectable cash record: original currency and amount, scenario or confirmed rate, rate source and date, charges, XAF equivalent, payment date and approver. A treasury specialist should assess any actual currency-risk instrument, its accounting treatment and its suitability. Operations needs the resulting payment-dated cash constraint before it can certify that the stock plan is feasible.

11.5

Receivables finance does not automatically fund the first fabric buy

The timing of a finance product determines which cash gap it can address. IFC's Global Trade Supplier Finance description states that the programme provides post-shipment finance to suppliers based on acceptance of receivables by select buyers approved by IFC. It also describes purchasing and discounting invoices accepted for payment by pre-approved participating buyers. [S18] The described triggering events are therefore shipment and buyer acceptance of the receivable. They occur after the supplier has acquired inputs and performed the work needed to reach shipment.

That structure cannot be cited as automatic funding for the initial fabric deposit. A supplier deposit finances a commitment before production. Pre-shipment working capital finances materials, freight and conversion before a receivable exists. Supplier credit defers a payment to the input provider. Post-shipment receivables discounting addresses an accepted invoice after shipment. These mechanisms can complement one another, but they have different counterparties, evidence requirements, risks and dates. Treating them as interchangeable leaves the earliest part of the cash cycle unfunded.

Exhibit 48, p. 174 maps the opened programme description to the illustrative cash events used in this chapter. The mapping is derived from the post-shipment and accepted-receivable language in the source. It does not say that IFC offers no other product, and it does not establish eligibility for GFT or any proposed venture. [S18]

EXHIBIT 48 / EVENT COVERAGE DERIVED FROM THE OPENED PROGRAMME DESCRIPTION

Finance products cover different points in the cash cycle

Cash event	Position in the cycle	Opened GTSF description	Planning consequence
Supplier deposit	Before inputs are produced	Before the described post-shipment event	Fund separately before order release
Materials, freight and conversion	Before customer shipment	Before the described post-shipment event	Show pre-shipment cash requirement
Accepted customer receivable	After shipment and buyer acceptance	Matches the described programme basis	Assess actual eligibility and terms
Unaccepted or disputed invoice	Acceptance not established	Described acceptance condition is not met	Do not count as available finance
Customer collection at maturity	Cash settlement of the receivable	Distinct from earlier discounting	Reconcile finance repayment and residual cash

DERIVED [S18] Derived event mapping from S18, which describes post-shipment finance based on receivables accepted by select buyers approved by IFC. The deposit, pre-shipment and disputed-invoice rows are illustrative needs in this chapter, not statements about every IFC product. No GFT or venture eligibility, approval, pricing or funding commitment is asserted.

The illustrative timeline makes the sequence concrete. By the assumed invoice-acceptance event on day 65, the business has already paid XAF 21 million of deposits, logistics and conversion costs. Even if accepted-receivable finance became available at that point, it would replenish cash after those outflows rather than explain how the initial XAF 6 million deposit was funded. Collection remains a separate event on the illustrative day 95. The fees, recourse, advance proportion, eligibility and settlement mechanics of any real facility would need dated documentary evidence. [S18]

GFT would therefore divide the funding plan at the first shipment event. Before shipment, the plan should identify own cash, supplier credit, any confirmed pre-shipment facility and the reserve that cannot be spent. After shipment, it can separately assess whether an accepted receivable qualifies for discounting under actual terms. An expression of interest, a generic programme webpage or an expected customer relationship is not a confirmed funding input. The stock plan should use only the cash and credit whose status the decision record discloses.

11.6

Ageing should trigger recovery decisions before value disappears

Cash exposure does not end when material reaches the warehouse. Every lot that remains unused continues to connect invested cash to a narrowing set of recovery options. The useful ageing record therefore combines lot age, current stock state, remaining selling window, specification version, customer restrictions and estimated recovery cost. A single days-on-hand average can conceal a small group of shade lots, logos or size variants whose practical recovery is falling even while aggregate stock appears healthy.

The trigger for action should be an event that changes recoverability, not merely an arbitrary birthday. Relevant events include a cancelled order, a design revision, a failed quality release, the last feasible production date, the end of a selling window, a customer return and a missed recovery review. Each event should assign an accountable role and a deadline for deciding between rework, approved substitution, return, resale, donation, reuse, recycling or write-down. The legal and accounting treatment of the selected route requires professional review.

IAS 2 defines net realisable value as estimated selling price in the ordinary course of business less estimated completion costs and costs necessary to make the sale. [S09] GFT would make that calculation lot-specific where the route to recovery differs. Expected proceeds should not be confused with cash already recovered. Sorting, repair, repacking, freight, commissions and disposal can reduce the amount available, while an uncertain outlet price should remain a scenario rather than a confirmed receipt.

Returns add time and cost to this boundary. The European Environment Agency reports that returned and unsold textiles are expensive to handle and explains that complex return processes can reduce resale potential, particularly for seasonal products. [S26] This is European evidence, not a Cameroon return rate or a recovery-cost parameter. Its operational lesson is to measure inspection delay, handling, resale eligibility and realised proceeds rather than recording every returned garment at its original expected value.

For an EU-facing scenario, the recovery route may also carry regulatory conditions. The European Commission states that from 19 July 2026 large companies across the EU are prohibited from destroying unsold clothes, clothing accessories and footwear, with medium-sized companies subject to the same rules from 2030. The announcement describes exemptions and conditions that require legal interpretation. [S21] It is not a blanket obligation on a Cameroon manufacturer. It is a reason to establish the customer, entity, product scope and applicable recovery duties before choosing disposal.

> **GFT NOTE**
>
> Create a recovery queue beside the replenishment queue. For each at-risk lot, record cash already committed, physical state, last useful production or selling event, possible recovery routes, estimated completion and selling costs, decision owner and next review. Compare routes on expected net cash recovery and service effect, not on original purchase price alone. Keep accounting write-down, physical disposition and actual cash receipt as three separate outcomes. [S09] [S26]

11.7

Finance and operations must approve one feasible policy

The final stock policy needs one shared approval boundary. Operations should not approve a service promise without the payment-dated cash curve. Finance should not cut an order by an aggregate percentage without seeing complete-kit coverage, pack effects and the customer exception that results. GFT would put both reviews on

the same item-location and order record, with authority limits for cash at risk, service shortfall and material readiness. The approved output is a feasible compromise or an explicit decision not to commit, never two inconsistent departmental plans.

A useful review packet begins with the original customer promise and the candidate quantities. It then shows the cash curve by obligation date, the reserve, the foreign-currency scenarios, available credit by evidential status, storage and pack constraints, complete-kit coverage and residual stock exposure. The final page should rank exceptions by the promise they affect. Each exception needs the missing input or decision, its owner, the latest safe decision date and the consequence of doing nothing.

The packet should present a small service, cost and cash frontier rather than one opaque optimum. One candidate may provide higher complete-kit coverage but require an earlier deposit. Another may reduce peak cash while increasing terminal stock or emergency freight. A third may preserve the promise only after a supplier agrees to split delivery. Each candidate should use the same demand basis and customer date so the trade-off is visible. The review should also identify dominated candidates: those that require equal or more cash and relevant cost while delivering equal or less coverage. Removing dominated options simplifies the discussion without concealing feasible alternatives. The chosen point on the frontier remains a management decision because the sources do not assign a universal monetary value to a missed garment order or a unit of cash headroom. [S02] [S05]

Cash-at-risk limits should distinguish committed cash from uncommitted headroom. Committed cash includes paid amounts and obligations the business cannot avoid without an approved contract change. Headroom is the amount remaining under the approved limit after those commitments and the protected reserve. A supplier quote can increase committed cash before stock arrives. A customer acceptance can remain unavailable for spending until collection or confirmed finance. These distinctions stop the same expected receipt from supporting several orders at once.

> **GFT NOTE**
>
> Approve a cash-constrained stock plan through four gates. First, operations confirms usable-material and complete-kit coverage. Second, finance confirms every dated obligation and the reserve under agreed scenarios. Third, both roles review uncovered demand, ageing exposure and funding status. Fourth, an authorised role accepts the plan, changes the promise or stops the order. Any later change to quantity, price, currency assumption, payment date or customer date reopens the affected gates. [S05] [S09]

There is a legitimate contrary case to cash minimisation. Management may deliberately fund additional stock for a strategic customer, a fragile supply route or a selling window whose expected contribution justifies the exposure. GFT's method does not prohibit that choice. It requires the choice to show its maximum cash requirement, downside recovery and affected promises. A conscious exception with a named owner is different from an unnoticed breach created by an item-by-item reorder rule.

The same record supports later learning. After receipt, production and collection, the team should compare expected and actual payment dates, usable quantities, service shortfall, recovery proceeds and peak cash. It should preserve the assumptions that were wrong rather than replacing them with hindsight. The next policy can then change a measured parameter, renegotiate a pack or revise a customer promise. Detailed solver authority belongs to the following chapter, and financing structures belong to the working-capital paper. This chapter supplies the feasibility boundary they must respect.

A feasible plan is therefore not the cheapest statistical stock calculation, nor the largest buffer the warehouse can hold. It is the best inspected commitment that fits material, service, space and cash at the dates on which consequences occur. Where no such commitment exists, the honest result is infeasibility and a ranked set of decisions, not an order that assumes the money will appear.

Cash does not merely pay for the stock plan. It decides which promise the plan can keep.

CHAPTER SOURCE NOTES

S1101 ICC Academy. 2026-01-20.

S09 IFRS Foundation. n.d..

S31 GS1 US and Auburn University RFID Lab. 2018.

S18 International Finance Corporation. 2026.

S02 MIT OpenCourseWare. 2006.

S08 MIT OpenCourseWare. 2006.

S03 MIT OpenCourseWare. 2006.

S05 MIT OpenCourseWare. 2006.

R32 Wise; XE. Publication date unrecorded; inherited source-text check.

S26 European Environment Agency. 2024.

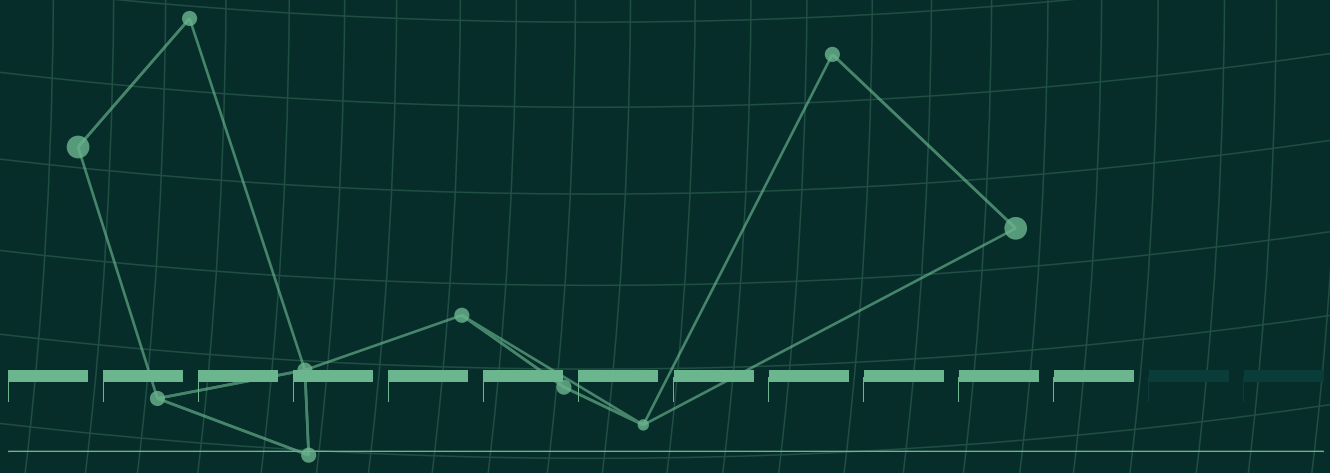
S21 European Commission, Directorate-General for Environment. 2026-07-17.

12

PART IV: PROVE THE POLICY BEFORE WIDENING ITS AUTHORITY / DECISION SYSTEMS

Optimisation earns authority through reproducible decisions and safe fallback

A small transparent policy engine is more useful than an opaque optimiser that cannot explain missing data or infeasible orders.



CHAPTER 12
DECISION TEST

A small transparent policy engine is more useful than an opaque optimiser that cannot explain missing data or infeasible orders.

An inventory decision engine should earn the right to recommend before it earns the right to commit. GFT proposes a reproducible calculation, a visible constraint check, an accountable approval and a fallback that preserves unresolved obligations. More sophisticated optimisation can widen the choices examined. It should not silently widen purchasing authority. This chapter specifies that boundary for an apparel maker; it describes a proposed operating design, not an installed GFT system or evidence of realised savings.

12.1

The baseline should be executable without machine learning

GFT would begin with an item-location policy that a planner can calculate independently. It should identify the stock object, the customer promise, the review event, the replenishment rule and the owner of exceptions. A policy also needs its effective date, expiry condition and permitted material substitutions. An engine that cannot state these fields has no stable instruction to execute, even if its forecast is statistically elaborate.

The foundational teaching material distinguishes order-point policies from periodic review and defines inventory position separately from physical stock. Those distinctions support an executable baseline: a trigger, an order calculation and a review calendar can be specified without a learning algorithm. Their suitability still depends on the demand and replenishment process. [S05] GFT proposes keeping this baseline available as a comparator and as a documented manual procedure, subject to current input checks.

The forecast interface should accept an identified forecast vintage, item-location scope, horizon, demand unit, time basis and uncertainty representation. It should also describe whether the demand represents customer requests, observed sales or component consumption. GFT would reject an interface that sends a bare number with no definition. Forecast estimation belongs to the forecasting paper; this chapter owns the contract by which an estimate becomes a candidate replenishment action.

Calendars require equal care. GFT would distinguish the factory consumption calendar, supplier dispatch calendar and delivery acceptance calendar. A supplier closure should change the feasible receipt date before a recommendation is approved. The planner should be able to inspect the calendar version used in the calculation and see any assumption about working days. A recommendation recalculated after a calendar change should supersede its predecessor rather than appear as an additional requirement.

The exception queue is part of the baseline. GFT proposes grouping exceptions by the decision they block: uncertain availability, incompatible specification, unaffordable quantity, late receipt or missing approval. Each entry should identify an accountable role, the affected promise and the evidence needed to clear it. The North Carolina cycle-counting procedure separates count entry from approval and explicitly addresses discrepancies; it is a procedural reference, not proof of apparel-factory performance. [S27]

The proposed chain in **Exhibit 49**, p. 180 keeps data capture, calculation and commitment distinct. Its layers are design responsibilities, not claims about deployed software. A small factory might initially perform several responsibilities in a controlled workbook. The separation remains useful because the planner must know which action changes a record and which action changes a commercial obligation.

EXHIBIT 49 / PROPOSED RESPONSIBILITIES

A recommendation needs an inspectable path to approval



ILLUSTRATIVE [S07] [S27] [S31] Illustrative GFT architecture. The cited sources inform material planning, count approval and shipment reconciliation; they do not establish this design as installed.

GFT would preserve a specimen decision packet that another authorised planner can reproduce from the saved extract. Reproduction means the same accepted inputs produce the same policy calculation, rounding and constraint results. It does not mean that the next shipment will arrive as predicted. The packet should distinguish reproducibility of a calculation from reliability of the commercial assumptions on which it depends.

A baseline also needs a declared response when no replenishment is proposed. GFT would distinguish a valid decision to wait from an absent calculation, an expired policy and a blocked input. Each state should have a reason visible to the planner. Silence should never be interpreted as confirmation that the customer promise is protected. The review record should show the next event that can change the decision, such as a usable receipt, a revised requirement or an approved parameter update.

12.2

Optimisation must expose its objective and constraints

A solver needs an economic question rather than a general instruction to improve inventory. GFT proposes quantities by material and release date as decision variables, with explicit choices for approved substitutions where permitted. The objective should identify relevant purchase, ordering, holding, shortage and expediting costs over a stated horizon. The treatment of remaining stock must be visible. A terminal assumption can otherwise make buying beyond the selling window appear attractive without revealing the residual exposure.

The MIT material on lot-sizing extensions provides a basis for examining purchase discounts alongside inventory costs. It does not establish local supplier prices or factory constraints. [S03] GFT would therefore bind the objective to dated commercial inputs and retain the input currency and cost boundary. A discount schedule from a supplier quotation should be distinguished from an assumed teaching schedule. Cost coefficients should never imply precision that the underlying shortage or recovery evidence cannot support.

GFT proposes separate constraints for spending dates, order multiples, minimum lots, receiving space, available conversion capacity and material compatibility. Service commitments need their own mathematical definition. A promised garment quantity cannot be supported by an aggregate fabric total if required components are unavailable. MIT's material-planning lecture explicitly raises capacity and uncertainty as issues beyond an initial requirements calculation; a requirements schedule alone does not prove feasibility. [S07]

For implementation, GFT would write the inventory balance as opening usable supply plus accepted receipts, less consumption and other authorised outflows, equals closing usable stock. The arithmetic must use the same material state throughout. A receipt under quality hold should enter the physical ledger without automatically entering usable supply. The finance constraint should evaluate payment obligations on their actual assumed dates. A plan passing an aggregate budget check may still fail at a deposit date.

The engine should return a status alongside its proposed quantities. GFT's proposed statuses in **Exhibit 50**, p. 182 separate independently checked feasibility from a claim of optimality. A time limit may leave an admissible candidate, an inadmissible candidate or no candidate. The planner needs to know which occurred. The status should preserve termination reason, objective value, any available bound, the reported gap definition and the result of independent constraint checks.

EXHIBIT 50 / PROPOSED DECISION STATES

Solver status determines the permitted next action

State	Interpretation	Permitted response
Checked feasible	All saved constraints pass	Review within authority
Time limit, candidate	Feasible; optimum unproved	Review candidate and bound
Time limit, no plan	No feasible candidate saved	Escalate; no order output
Proved infeasible	Saved constraints conflict	Approve changed promise
Invalid or stale input	Problem is not admissible	Repair and recalculate
Unbounded model	Objective lacks a valid limit	Block and inspect model

ILLUSTRATIVE [S03] [S07] Illustrative GFT control contract. These labels are proposed business states, not the native status names of a selected solver. Feasibility is checked against the saved model.

A reported optimum remains conditional on the submitted formulation. GFT would independently recalculate pack divisibility, cash timing, material balances and capacity consumption from the proposed quantities. Integer rounding after optimisation should trigger those checks again. If costs can take values close to zero, the report should explain how a relative gap is defined rather than presenting an unstable percentage as an intuitive measure of economic importance.

Infeasibility is a business finding when the formulation is correct. GFT would present the conflicting obligations and controlled alternatives: change the delivery promise, split the order, qualify another source or obtain additional approved funding. A diagnostic model may allow labelled service shortfalls to identify a trade-off. Its output must remain a diagnostic until the authorised owner accepts the changed promise. A penalty coefficient should not quietly purchase permission to breach a hard commitment.

12.3

Uncertainty belongs in scenarios that a planner can inspect

A deterministic run answers what the selected model would recommend if its point inputs occurred. It does not describe the consequences of a delayed trim delivery or an unobserved demand surge. GFT proposes saving the deterministic case because it is easy to audit, then testing the same candidate quantities against named scenario paths. Each path should specify demand, usable-receipt timing, accepted quantity, costs and capacity events on a common calendar.

Maitra's research manuscript combines stochastic inventory simulation with optimisation and discusses both random and conditional sampling. It is a method illustration, not evidence that a particular policy will outperform in Cameroon. [S23] GFT would

retain the generated paths and random seeds where sampling is used. Another planner should be able to inspect the path that caused a shortage, rather than seeing only an average objective or a favourable selected run.

Scenario probabilities require evidence. Where no local distribution has been calibrated, GFT would label adverse paths as illustrative stress cases and report their consequences separately from any probability-weighted result. A late-delivery case might be important even when its probability is unknown. Assigning it an arbitrary small probability would create an apparently precise expected cost while concealing the uncertainty the exercise is meant to expose.

Dependence is another input. Michna and colleagues examine models with stochastic lead times and discuss differing demand and lead-time assumptions; their analysis does not supply a local apparel-factory distribution. [S24] GFT proposes retaining joint episodes where demand pressure and supply delay occur together. If independent resampling is used, the report should identify that assumption and compare it with a deliberately dependent stress path. The comparison tests sensitivity to the assumption without claiming that either path has been observed locally.

Observed sales are also an incomplete foundation for some scenarios. Chen and colleagues describe censoring of both demand and supply in their research problem. [S22] For the proposed factory interface, GFT would preserve stockout flags, unfilled requests, partial receipts and unresolved open orders. Those records should accompany any fitted input. A period with no sale during unavailability must not automatically become evidence of no demand.

A conservative uncertainty set offers another representation. GFT would define admissible ranges and test whether a candidate remains feasible across their specified combinations. This can reveal exposure without assigning probabilities, but the set itself is a judgement requiring approval. Combining every adverse bound may produce an unaffordable plan. That result should prompt a discussion of service and recovery options, with the assumptions visible, rather than automatic expansion of every buffer.

The operational report should therefore separate expected outcomes under a stated distribution, outcomes on named stress paths and failures outside the admitted model. GFT would show which customer promise is lost, when cash becomes binding and whether an approved fallback remains possible. The next chapter owns the full comparative evaluation and field-pilot protocol. Here the requirement is narrower: the decision engine must retain enough detail for those evaluations to be performed honestly.

An illustrative stress calculation clarifies the distinction. Suppose a component needs 40 units per working day and a delay adds 3 working days without replacement supply. Additional exposure is 40 × 3 = 120 units before any change in demand or initial availability. This is teaching arithmetic, not a recommended buffer. The decision packet should retain the assumed rate, calendar and delay together; changing any of them should produce a new scenario result.

Scenario reports should retain the timing of a failure as well as its total magnitude. GFT would identify whether a shortage occurs before a production release, during conversion or after the promised dispatch. The same missing quantity can have different consequences at these boundaries. A feasible recovery before release might preserve the promise, while a late arrival can leave an otherwise complete garment order stranded. The proposed engine should show the affected requirement and the recovery assumption explicitly, leaving commercial acceptance to the authorised planner.

12.4

An input error should stop an order recommendation

GFT proposes admission checks before any calculation can produce an approvable purchase recommendation. The checks should distinguish missing, invalid, stale and disputed inputs. A missing lead time is unknown. A negative lead time is invalid. A quotation beyond its accepted review window is stale. A receipt that conflicts with the physical count is disputed. These states need different remedies, even when they all block order release.

The item definition should carry the purchase unit, consumption unit and governed conversion. GFT would forbid conversion between rolls and metres without an accepted roll-length basis, and between mass and length without the relevant specification. Apparent numerical compatibility is insufficient. Material revision, shade restrictions and approved substitution should be checked before netting supply against demand. The engine must not manufacture interchangeability by collapsing different stock identifiers.

For time inputs, GFT would require event definitions and calendars. Promised dispatch, physical arrival and usable acceptance are different events. An imported timestamp without a time-zone interpretation should not silently determine which review window contains the receipt. Future-dated movements and out-of-sequence releases should enter an exception queue. The data owner can then explain a recording delay or correct a genuine error while preserving the original event.

Count discrepancies require controlled resolution. The opened North Carolina procedure explicitly covers entering results and approving counts. [S27] GFT would use that separation to design a local discrepancy workflow: record the observation, investigate relevant movements, approve an adjustment and recalculate affected recommendations. The recommendation engine should not overwrite the ledger to make its material balance fit. Model feasibility cannot adjudicate a disputed physical quantity.

Shipment reconciliation is similarly distinct from successful data transfer. Project Zipper compared barcode and RFID information across shipping and receiving stages; its study-specific findings should not become a guaranteed accuracy target for a different factory. [S31] GFT proposes reconciling expected, dispatched, physically received and accepted quantities as separate states. Duplicate messages should not create additional stock, and a successfully uploaded shipping notice should not become a usable receipt.

The failure response should remain proportionate to the affected boundary. GFT would block recommendations using an unresolved item-location record and flag any dependent garment commitments. Unrelated records may continue only if their independence has been established. A shared unit-conversion error may invalidate a wider scope than a local count dispute. The interface should explain the blocked scope and the rule that caused it, making recovery an accountable decision.

GFT would also apply numerical admission checks: finite values, valid probability ranges, consistent horizons and adequate precision for the declared units. Zero variability should be accepted only when it is a deliberate input supported by the chosen method, rather than a replacement for a missing estimate. Extremely large proposed quantities should be compared with approved bounds before display as a recommendation. Validation should reveal uncertainty rather than hide it behind defaults.

Finally, freshness should follow the decision's consequence. GFT would set review rules by input class and record the last accepted observation. A stock extract can become stale after an important movement even if little time has elapsed. A stable material specification may remain usable longer than a delivery promise under disruption. The system should evaluate these explicit rules, retain their versions and state which changed input would make recalculation necessary.

12.5

Integration needs an idempotent decision and event log

GFT's proposed integration contract separates the identity of a physical event from the identity of a recommendation. A receipt event should have a stable identifier that survives retry. A recommendation should have an identifier tied to its decision window, scope and revision. A technical retry must reuse the original identifier; a genuine recalculation must explicitly supersede the earlier recommendation. Otherwise the receiving planner cannot distinguish a duplicate from an additional requirement.

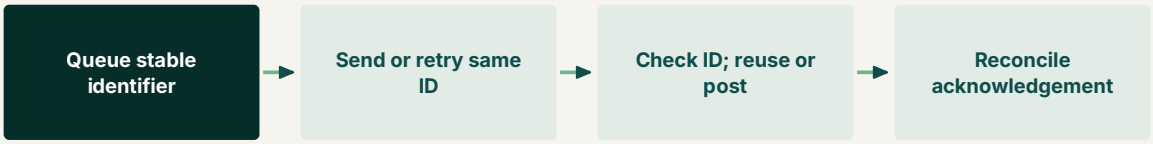
The decision log should retain the source-extract fingerprint, forecast vintage, calendar and policy versions, solver configuration, quantities before and after rounding, constraint results and approval state. If a planner overrides a quantity, GFT would save both the calculated proposal and the authorised replacement, with its reason. The log should record whether the replacement passed the same checks. An override should not acquire apparent mathematical approval merely because it appears in the same screen.

Reproducibility also requires preserving ordering and tie-breaking rules where they affect results. GFT would record the software version and random seed when relevant, together with tolerances used in feasibility checks. Equivalent objective values may permit different feasible plans. The packet should distinguish an exact replay from a materially equivalent result and state the equivalence criteria. Commercial approval should bind to the selected quantities and dates, not an unspecified family of alternatives.

The reconnect sequence in **Exhibit 51**, p. 186 is an illustrative business contract. It requires an acknowledgement or reconciliation before a transmission is treated as complete. The architecture paper owns transport and storage implementation, while the data-governance paper owns permissions and retention. This chapter specifies the inventory consequence: uncertain delivery of an instruction must remain uncertain until the receiving system's state has been checked.

EXHIBIT 51 / PROPOSED RECONNECT SEQUENCE

A retry must preserve the original business identifier



ILLUSTRATIVE [S31] Illustrative GFT transaction design. Shipment reconciliation motivates the business control; the source does not establish idempotent transport or an exactly-once guarantee.

Consider an illustrative failure in which a purchasing system creates an order but the response is lost. GFT would retain the submission as acknowledgement pending and query the receiving record using the stable identifier. Reissuing it as a new business request would be prohibited. If the receiving system cannot support this lookup, the fallback should require manual reconciliation before another submission. A local log alone cannot prove that a remote order was never created.

GFT would require the receiving control to bind the identifier to the submitted content. A retry with matching content should return the recorded result; the same identifier with changed quantities should be rejected for reconciliation. The uniqueness check and posting must share a protected transaction boundary so concurrent retries cannot both post. Where the receiving system cannot provide this behaviour, execution authority should remain restricted and the manual reconciliation requirement should be explicit.

The same principle applies to reversals. GFT would treat cancellation and correction as linked events rather than deleting the original record. An order cancellation request is distinct from supplier acceptance of that cancellation. A material receipt reversal is distinct from physical removal of the goods. Each event should retain its own acknowledgement and authority so that the stock and commitment ledgers can be reconstructed without guessing what an edited field previously meant.

Concurrent decisions need a further check. GFT proposes revalidating stock reservations, remaining spending authority and policy validity immediately before release. Another planner may already have committed the shared component or used the available cash. If the saved snapshot no longer supports approval, the recommenda-

tion should return for recalculation. Approval of an earlier snapshot should not reserve every changing resource indefinitely unless a specific reservation mechanism establishes that right.

Acceptance criteria should include failed transactions as explicit cases. GFT would ask the implementation team to demonstrate a repeated receipt message, a retry after a lost response, a cancelled proposal reopened by an old browser session and an order acknowledgement received after a manual fallback purchase. The expected result should be stated before each exercise. These demonstrations would check the proposed control logic; passing them would not establish operating reliability under every outage. Any unresolved receiving-system behaviour should remain a release restriction.

12.6

Automation authority should remain narrower than model coverage

GFT proposes starting with read-only calculations and named planner approval. A model may cover materials, sites or disruption cases that have never been accepted for automated action. Coverage is useful for analysis. Execution authority should be separately listed by item-location, supplier, quantity, spending date, permitted action and validity window. Any request outside that list should return to review, even when the calculation is mathematically valid.

The conceptual matrix in **Exhibit 52**, p. 188 is deliberately illustrative. Its ordinal coordinates organise a governance discussion; they are not calibrated evidence scores or an automatic promotion rule. GFT would require the decision owner to approve both the supporting evidence and the permitted reversibility of the action. A move towards bounded execution should be a documented authorisation, with the former approved state available for restoration.

EXHIBIT 52 / ILLUSTRATIVE ORDINAL PLACEMENT

Authority expands only within an approved envelope



ILLUSTRATIVE Illustrative GFT governance map. Coordinates are teaching positions on ordinal axes, not measured scores. Moving a point does not authorise an order.

The release envelope should include cumulative exposure. GFT would require a cap that applies across related recommendations, outstanding submissions and accepted orders. Splitting a large quantity into smaller requests should not bypass that cap. The rule should specify when exposure is released, for example after a reconciled cancellation or an approved completion event. Merely losing the connection should not restore spending authority already consumed by an uncertain submission.

> **GFT NOTE**
>
> Keep policy approval separate from order approval. A planner may have permission to run a calculation without permission to change a service target or commit cash. GFT recommends an explicit expiry, a visible stop control, an acknowledged manual fallback and a reconciled commitment ledger before any authority expands. This is a proposed control design; this chapter authorises no vendor order.

An emergency stop should prevent new submissions and make the in-flight state visible. GFT would distinguish stopping the engine from cancelling orders already accepted elsewhere. Operators need a queue showing which actions are pending, acknowledged or commercially committed. Reopening execution should require reconciliation of that queue and confirmation that the admission checks still pass. A restart should not simply replay every saved instruction as a fresh request.

Manual fallback also needs a boundary. GFT would retain an approved rule card, the last trusted stock snapshot, unresolved movements and a register for temporary decisions. If stock truth is unavailable, the manual process may require a physical

check and finance confirmation before buying. The objective is controlled continuity of necessary decisions. It is not permission to reproduce stale quantities on paper or to promise service that the known constraints cannot support.

A return from fallback should include reconciliation of temporary orders, supplier acknowledgements, receipts and reservations. GFT would then regenerate affected recommendations from the reconciled state and explicitly retire superseded proposals. The next chapter defines how operational evidence supports a wider pilot. The present chapter supplies the control conditions that make such evidence interpretable: the reader must know which policy acted, under whose authority and with which exceptions.

12.7

The safety-stock tool must show what it does not calculate

A safety-stock output should expose the chosen service measure and admitted assumptions beside the answer. Under stationary independent daily demand and fixed lead time, the teaching formulation uses mean lead-time demand plus a safety factor times its standard deviation. Cycle service and unit fill rate are different measures in the MIT material. [S05] GFT's proposed interface should therefore refuse to label a cycle-service input as a guaranteed fill-rate result.

The worked case below is GFT's illustrative arithmetic for a deliberately narrow model. It does not repeat the full derivation from the safety-stock chapter. Its purpose is to show what a decision packet must reveal: assumed inputs, an equation, unit handling, rounding, a constraint check and the limits of the calculated output. No factory data, supplier quotation or measured local service outcome is implied.

WORKED EXAMPLE: A REORDER TRIGGER CANNOT APPROVE AN ORDER

An illustrative repeat trim has mean demand of 40 units per working day, daily residual standard deviation of 10 units and fixed usable-receipt lead time of 9 working days. Assume stationary independent daily errors, an approximately normal lead-time total and a chosen cycle-service factor z = 1.645. Safety stock is z × 10 × sqrt(9) = 49.35 units. Mean lead-time demand is 40 × 9 = 360 units. The unrounded reorder trigger is therefore 409.35 units. Rounding upwards gives a displayed trigger of 410 units. These assumptions are teaching inputs, with the probability framework informed by the foundational lecture. [S05]

In this illustrative case, assume usable on-hand of 250 units, firm outstanding supply of 100 units and disjoint unfilled commitments of 80 units. Inventory position is 250 + 100 - 80 = 270 units under this stated convention. Suppose the separately approved fixed order lot is 200 units. The trigger is reached, so the candidate is that fixed lot. The tool must not subtract inventory position from the reorder trigger and call the difference the order quantity. Trigger and lot size are separate policy parameters. [S05]

Assume an illustrative purchase price of XAF 500 per unit and an available commitment ceiling of XAF 80,000. The candidate costs 200 × 500 = XAF 100,000 before any other charges and exceeds the assumed ceiling by XAF 20,000. The correct output is a blocked candidate with a finance exception. Substituting a smaller order would change the approved policy and needs a separate decision. The calculation supplies no evidence that the chosen lot or service factor is appropriate locally.

The sensitivity table in **Exhibit 53**, p. 190 changes only assumed lead time. The model's trigger rises with both mean consumption during replenishment and its uncertainty term. It does not account for additional variability in lead time itself. The table is useful precisely because its narrow basis is visible; it cannot establish the buffer needed under correlated delays, uncertain acceptance or a seasonal selling deadline.

EXHIBIT 53 / WORKING DAYS AND TRIM UNITS

A longer assumed lead time raises the illustrative trigger

Lead time	Mean demand	Safety stock	Rounded trigger
4	160	32.9	193
9	360	49.4	410
16	640	65.8	706

ILLUSTRATIVE [S05] Illustrative arithmetic: mean daily demand 40, independent daily residual standard deviation 10, fixed L, z 1.645. Safety stock = 1.645 ℴ sqrt(L); trigger = 40 * L + safety stock. Displayed trigger rounds upwards to a whole unit. No observed demand or local probability calibration.

GFT would display exclusions next to this result. The narrow calculation does not coordinate multiple stock locations, price seasonal obsolescence, infer censored demand, model correlated shocks, predict quality release or enforce cash. Those functions should appear as implemented capabilities only after their logic and interfaces have been verified. Linking to another analytical tool should not imply that its constraints were actually included in this calculation.

Rounding deserves a visible distinction between a mathematical threshold and a commercial quantity. GFT would retain fractional intermediate results, round a displayed stock threshold under its declared convention and apply purchase packs only to the candidate order. A pack conversion should not alter the meaning of the service parameter. If a rounded candidate exceeds cash, space or maximum-age constraints, the interface should show the conflict. Hiding the conflict by silently rounding down would substitute a different policy without informing its owner.

A sensitivity result should also identify what remains fixed. GFT would keep mean demand, residual variability, calendar and service factor visible while the user changes lead time. If a user instead changes the time unit, the interface should require consistent conversion of all relevant inputs. A numerical output that looks plausible after an inconsistent conversion is still invalid. The tool should permit the planner to recover the last valid input set and explain why the rejected set cannot support the selected equation.

A useful export should retain accepted inputs, rejected fields, the equation identifier, unrounded values, rounding convention and unresolved constraints. GFT would present a blocked state as a complete calculation outcome, with the next accountable decision visible. That is the practical connection to the programme's feasibility, performance-management and digital-strategy design: commercial commitments, measured events and digital recommendations should share an inspectable boundary. This chapter is an operating-method proposal, not investment, legal or tax advice.

A decision engine earns authority by preserving the reasons it must stop.

CHAPTER SOURCE NOTES

S05 MIT OpenCourseWare. 2006.

S27 North Carolina Office of the State Controller. 2023-10-24.

S07 MIT OpenCourseWare. 2006.

S31 GS1 US and Auburn University RFID Lab. 2018.

S03 MIT OpenCourseWare. 2006.

S23 arXiv, Sarit Maitra. 2024-06-27.

S24 arXiv, Michna, Nielsen and Nielsen. 2015-03-20.

S22 arXiv, Chen, Jiang, Zhang and Zhou. 2023-10-31.

GREEN FUTURE TECH

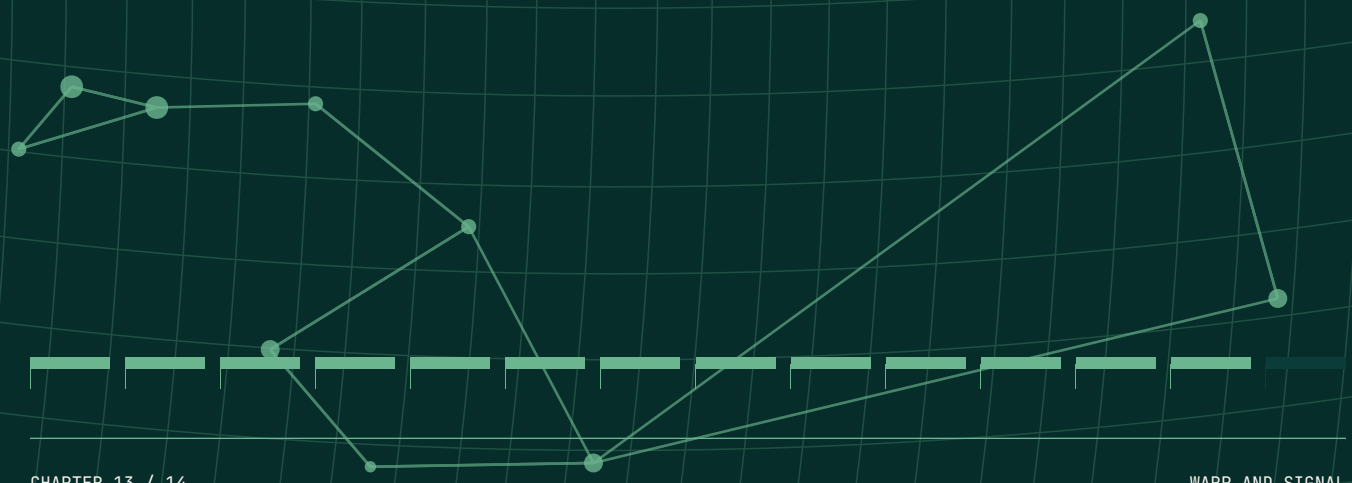
LIMITED / HONG KONG SAR

13

PART IV: PROVE THE POLICY BEFORE WIDENING ITS AUTHORITY / PROOF OF VALUE

A policy improves the factory only when measured service and cash improve together

Backtests, simulations and a controlled operating pilot answer different questions and must be reported separately.



CHAPTER 13
DECISION TEST

Backtests, simulations and a controlled operating pilot answer different questions and must be reported separately.

An inventory policy earns wider use when a credible comparison shows that it protects the customer promise, reduces avoidable cash exposure and respects supplier and worker safeguards. A favourable simulation cannot establish that result. GFT proposes an evaluation contract that preserves what was known when decisions were made, makes missing observations visible and refuses expansion after a failed gate. The contract below is a proposed method for the Cameroon textile programme context. It reports no GFT factory pilot, deployed replenishment system or measured inventory saving.

13.1

A replay must use only information available at the decision date

GFT recommends beginning evaluation with a precisely specified question: would the candidate rule have made a better feasible commitment using the information then available? Specify the material family, location, customer promise and authority being tested. A trial of planner recommendations should not quietly become a trial of supplier selection or customer pricing. Record the incumbent rule, including its exceptions, before interpreting the candidate's result. Otherwise the comparison risks testing a documented new process against an imaginary version of the old process.

Research on lost-sales inventory distinguishes observable transactions from underlying demand and supply. Research on simulation-based inventory optimisation compares policies within particular modelling assumptions. Neither source establishes that a historical factory ledger contains every input needed for a counterfactual decision. GFT therefore recommends a replay admission check before scoring policies: the evaluator should demonstrate the availability, meaning and timing of the required records. [S22] [S23]

For each review event, preserve the forecast vintage, usable stock by status, reservations, outstanding purchase commitments, released receipts, current cost basis, cash availability and applicable constraints. Keep both physical event time and recording time. A receipt physically delivered but not recorded until later was not necessarily visible to the planner. The replay should reproduce the agreed information boundary rather than granting the candidate a cleaner view of the past. Corrections discovered afterwards belong in a separate reconciliation layer.

Supplier observations need the same discipline. GFT proposes estimating a decision-date lead-time distribution only from information available by that date. Open orders remain open observations with elapsed time and promised dates; they are not completed deliveries with invented arrival dates. Retain partial receipts and rejected quantities. The censored-data problem identified in inventory research is a reason to expose these gaps, rather than silently fill them with convenient averages. [S22]

Separate tuning from evaluation. GFT recommends retaining an earlier calibration window, a subsequent validation window and a final untouched evaluation window wherever the available history permits that division. Choose boundaries around the actual ordering and selling cycle. Freeze parameter-selection criteria before reading

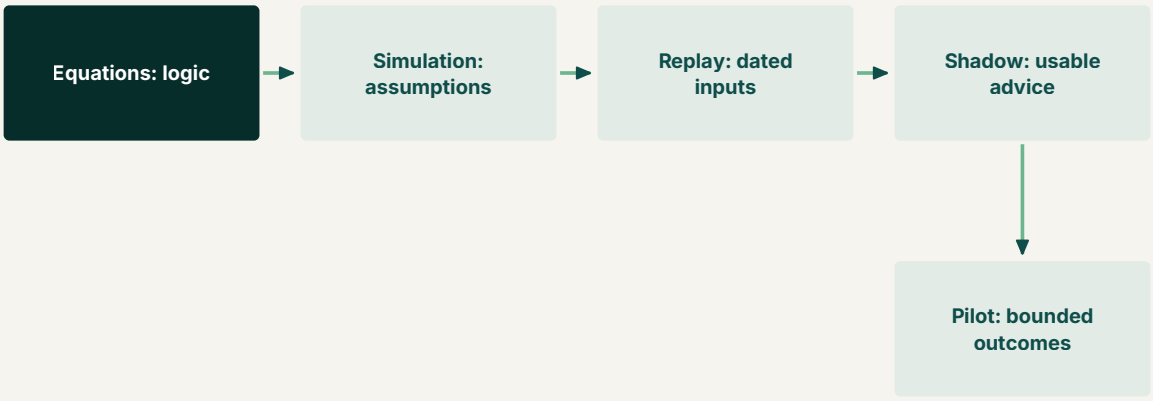
final outcomes. Report when a short or interrupted history prevents an honest holdout. A replay with repeatedly adjusted parameters should be labelled exploratory, even when its final score looks favourable.

A recorded receipt path also has limits. A different order quantity, order date or supplier may have changed dispatch priority, partial fulfilment or payment requirements. Holding historical receipts fixed is therefore a conditional experiment, not a reconstruction of an alternative reality. GFT recommends reporting a fixed-receipt replay separately from scenarios that model supplier response. Where that response cannot be estimated, retain an explicit assumption and a sensitivity range rather than calling the alternative result observed.

The evidence ladder in **Exhibit 54**, p. 194 is a proposed sequence of questions, not a numerical ranking of credibility. Equation checks establish implementation consistency; later stages expose additional failure modes. None automatically substitutes for the next.

EXHIBIT 54 / PROPOSED EVALUATION SEQUENCE

Each evidence stage permits a narrower claim than factory-wide success



ILLUSTRATIVE [S19] [S22] [S23] Illustrative GFT evidence ladder, informed by simulation and censored-observation methods. No stage is reported as completed. Arrows denote additional checks, not automatic promotion.

GFT would retain a replay manifest containing input extracts, clock convention, parameter version, exclusion reasons and the exact decisions scored. A reviewer should be able to reconstruct an unfavourable recommendation as readily as a favourable one. If stock states or purchase approvals cannot be reconstructed, the appropriate finding is limited evaluability. That finding supports improved capture before a stronger claim; it does not establish that the candidate policy has failed operationally.

13.2

Observed sales do not reveal every missed request

Inventory research explicitly identifies censoring in both demand and supply observations. When a customer cannot obtain an unavailable item, a sales record may contain no transaction for the missed request. Equally, a short delivery need not reveal the supplier's full available capacity. These are information limits, not permission to invent the missing quantities. The cited research manuscript develops a specialised learning model; its theoretical guarantees are not evidence of local factory service. [S22]

GFT proposes a request ledger that records the requested item and specification, quantity, request time, confirmed due date, offer, acceptance, substitution, cancellation and final fulfilment. Link repeated enquiries to the same commercial request where possible. A refused quotation, a cancelled confirmed order and a late shipment should remain distinguishable. They answer different questions about demand, the promise and execution. Preserve the original confirmed date when a later date is negotiated.

Do not make unavailable items disappear from the denominator. For a proposed immediate unit-fill measure, count eligible requested units and identify those supplied immediately under the agreed definition. For order delivery, retain the complete confirmed order and due-date boundary. GFT recommends disclosing cancellations and customer-approved substitutions alongside the principal result. A policy should not receive an apparent service improvement merely because difficult requests were removed or a different material was shipped without equivalent approval.

Use bounds when request capture is incomplete. In an illustrative stockout interval, assume recorded sales of 80 garments, documented unfilled requests for 10 distinct additional garments, and an explicitly assumed further unrecorded-demand range of zero to 20 garments. Eligible requests would lie between 90 and 110 garments. With all recorded sales counted as immediately supplied, the illustrative fill range is 80 divided by 110 to 80 divided by 90, approximately 72.7 to 88.9 percent. The range expresses an assumed missing-demand bound, not a confidence interval.

That illustration deliberately avoids treating stockout duration as a demand estimate. GFT recommends using availability timestamps to identify exposure and seeking corroboration from enquiries, backorders or comparable in-stock periods. Comparable periods still require a judgement about season, price, customer mix and selling opportunity. If the ordering decision changes customer willingness to request, historical demand is itself an imperfect comparator. Report that dependence rather than presenting a single adjusted sales series as measured truth.

A useful sensitivity test asks whether the policy ranking survives plausible missing-demand bounds. GFT would calculate the service and cost differences under the same declared assumptions for both policies. If the ranking reverses, the decision is sensitive to information that has not been observed. Prioritise better request capture or a bounded field comparison. Do not choose the assumption that makes the preferred candidate win and then describe the result as a backtested improvement.

Supply uncertainty requires an equally visible register. Distinguish ordered, confirmed, dispatched, received and quality-released quantities. A request for more material may expose a supplier capacity limit that the historical smaller orders never tested. GFT would therefore reject a replay that assumes the alternative order could always have been fulfilled merely because its simulated warehouse had room. The proposed record should identify whether evidence, a contractual commitment or an illustrative response model supports each alternative receipt.

13.3

Simulation isolates policy logic but does not validate local inputs

A simulation can represent a decision process that a static spreadsheet misses. NIST describes discrete-event modelling of manufacturing lines with finite buffers and machine degradation and maintenance. That establishes a relevant modelling capability, not a calibrated model of this programme. GFT proposes including a physical constraint only when it changes the decision being evaluated, and documenting where a simplified representation could change the result. [S19]

For an apparel-material trial, begin with stock conservation, request arrival, purchase release, partial receipt, quality release and issue to production. Add line starvation, capacity interruption or restricted storage when material availability depends on those events. Test boundary cases such as no demand, delayed receipts and all stock under quality hold. GFT recommends reconciling every simulated unit to an opening balance, accepted receipt, issue, loss or closing state before interpreting a cost or service score.

Compare policies on common exogenous paths. GFT proposes a fixed-coverage baseline, continuous-review candidate, periodic-review candidate and seasonal rule only where each fits the same commercial purpose. The inventory lecture material distinguishes periodic and continuous review and different service concepts. The simulation-optimisation manuscript provides a research example of policy comparison, but its case ranking should not be exported to another factory. [S05] [S23]

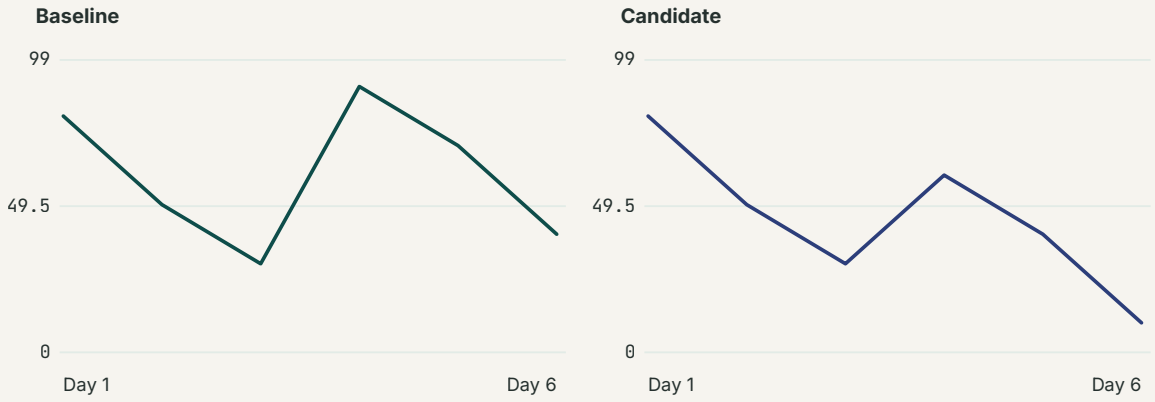
Freeze each rule's review calendar, trigger, order bounds and parameter-tuning procedure. Give all candidates the same cash ceiling, approved material alternatives and supplier constraints. If the current process has an information delay, either retain that delay for every policy or report improved capture as a separate intervention. GFT recommends recording infeasible recommendations as outcomes. Discarding them would reward a rule that appears inexpensive because it never completes difficult orders.

Repeatability requires more than recording a random seed. GFT proposes retaining the generator, input distributions, dependency structure, event ordering, initial pipeline, warm-up treatment and terminal valuation. When a candidate places different orders, define which random inputs remain common and which supplier responses depend on those orders. Use repeated independent paths to examine the distribution of paired differences. Choose replication effort according to the precision needed for the decision; a smooth average alone does not establish dependable performance.

The deliberately short illustrative calculation in **Exhibit 55**, p. 197 isolates arithmetic. It does not estimate a local demand distribution, tune a replenishment rule or claim policy superiority. Both options receive material before demand on the same assumed receipt day. A longer evaluation would need repeated replenishment opportunities and an explicit terminal treatment.

EXHIBIT 55 / USABLE PIECES AT DAY END

Identical assumed requests can produce different closing stock paths



ILLUSTRATIVE Illustrative arithmetic only. Both open with 100 pieces. Requests on days 1 to 6 are 20, 30, 20, 40, 20, 30 pieces. Receipt before day-4 requests is 100 pieces for baseline and 70 for candidate. End stock equals prior stock plus receipt minus served requests. All 160 requested pieces are served. No random generator, factory calibration or full policy simulation is claimed.

WORKED EXAMPLE: LOWER AVERAGE STOCK IS NOT RECURRING PROFIT

In this illustrative case, the sums of day-end usable stock are 360 and 270 pieces. Dividing each by the assumed six equally weighted observations gives average day-end stock of 60 and 45 pieces. At an illustrative material price of XAF 1,000 per piece, the difference in average material value is XAF 15,000. These are sampled closing balances, not a time-weighted average within each day.

The illustrative receipt difference is 30 pieces, costing XAF 30,000 at the same assumed price. With payment on receipt and no other changes, purchases consume XAF 30,000 less cash over this short window. Closing inventory is also 30 pieces lower. That is a balance-sheet difference, not XAF 30,000 of operating profit. Payment later, a deposit, an outstanding order or immediate replenishment after the window would change the cash interpretation. The candidate ends nearer exhaustion and could fail under further demand.

GFT recommends reporting uncertainty in the paired policy difference alongside absolute service and exposure. For repeated generated paths, describe how intervals were calculated and what variation they cover. Simulation intervals reflect the specified model; they do not include every omitted event or calibration error. Preserve an additional sensitivity analysis for uncertain inputs. Keep an independently checked equation, a repeatable model and an externally validated operating result as separate statements.

13.4

Stress tests should include combined rather than isolated failures

Stochastic lead-time research examines how replenishment variation interacts with ordering behaviour. Its assumptions and application context do not calibrate Cameroon transport or supplier performance. GFT proposes using that work as a warning against replacing an uncertain replenishment clock with a deterministic average, while testing local dependencies separately. [S24]

Construct a scenario register around mechanisms. A demand surge can coincide with supplier congestion; a partial fabric receipt can coexist with missing labels; a quality hold can delay material that the ledger already shows as received. A cash freeze may prevent replenishment just as overdue customer receipts increase. GFT recommends recording the combination, affected decision, observable warning, fallback and recovery condition. Avoid attributing a probability to a scenario unless an appropriate local basis has been established.

The heatmap in **Exhibit 56**, p. 198 identifies which synthetic shocks are switched on, rather than assigning measured risk scores. In this illustrative register, an active demand shock multiplies requests by 1.3, the delay shock adds five days, the quality shock releases only 75 percent of a planned receipt in time, and the funding shock funds only 60 percent of planned purchases. Binary cells show scenario composition, not equal severity or joint event frequency.

EXHIBIT 56 / SHOCK ABSENT (0) OR PRESENT (1)

Combined stresses change demand, release and funding together

	Demand surge	Added delay	Quality hold	Funding gap
Reference	0	0	0	0
Demand only	1	0	0	0
Delay and hold	0	1	1	0
Combined	1	1	1	1

Cell scale: 0 to 1

ILLUSTRATIVE Illustrative scenario activation, not outcomes, severity or probabilities. Active demand: requests x1.3; active delay: five extra days; active quality: 75 percent timely release; active funding: 60 percent funded purchases. Zero means the reference assumption. Combined activates all four shocks.

A useful stress run should explain the first binding constraint. GFT would identify whether cash blocks the order, a missing component blocks production release, quality approval blocks usable stock or a transport delay blocks receipt. Record the time

until the breach becomes visible and the action actually available then. A warning issued after the last feasible purchase date is not an effective protection, even if the dashboard eventually diagnoses the shortage correctly.

Test supplier correlation through physical dependencies rather than supplier names alone. OECD describes subcontracting and multi-tiered sourcing in the garment sector. GFT therefore recommends checking whether apparently distinct suppliers rely on the same fabric processor, carrier or constrained route. A scenario may impose a common outage without asserting its probability. Two alternative purchase addresses do not establish independent recovery capacity. [S25]

Connectivity loss needs an operational test as well as a data test. GFT proposes retaining a frozen authorised manual plan during the outage, recording local movements and checking reconciliation before reopening recommendations. The evaluation should count delayed decisions and manual effort. An automatic retry that duplicates a recommendation would be a control failure even if material service happened to remain acceptable. Detailed integration design belongs to the decision-system chapter; this chapter asks whether the fallback preserved the measured boundary.

When no feasible plan meets the customer promise, report that infeasibility. GFT recommends an exception decision on allocation, authorised substitution, negotiated delivery or reduced commitment, with its service consequences retained. Stress testing should not assume unlimited emergency freight, supplier credit or overtime as free recovery. Record the resources and approvals required to recover. The purpose is to expose where the policy stops being acceptable and whether a controllable response exists.

13.5

The field pilot needs a credible comparator

GFT proposes shadow operation before any operating pilot. In shadow mode, the candidate observes current inputs and produces dated recommendations while the incumbent retains decision authority. Compare recommendation timeliness, feasibility, planner interpretation and exception workload. Do not describe shadow recommendations as realised savings. Historical modelling and shadow operation prepare the comparison; they do not reveal what would have happened after different orders were actually placed.

The field design should identify the unit exposed to the policy. An item-location may suit independent repeat-stock materials, but not when competing garments share fabric or a common supplier minimum lot. GFT recommends mapping shared cash, material, capacity and decision makers before assigning comparison groups. If spillovers cannot be contained, consider a larger supplier-material or production cluster, while acknowledging that fewer independent groups make inference more difficult.

Select comparable eligible groups using pre-intervention demand pattern, criticality, supplier exposure, lead-time behaviour and stock-record quality. GFT proposes random allocation within comparable groups when operationally feasible and author-

ised. Where random allocation is impractical, retain a matched comparator and make the selection rationale inspectable. A staged introduction can provide additional comparisons, but the rollout order and concurrent changes must be recorded. Neither matching nor staging automatically removes bias.

Choose the observation window around event opportunities. Rare imported-material receipts, short seasonal selling windows and frequent local call-offs provide different amounts of evidence over the same calendar period. GFT recommends using actual request and replenishment frequencies, baseline variation, acceptable decision error and feasible grouping to plan sample size. Keep dependent observations together in uncertainty calculations. No universal fixed-week duration or convenient transaction count is prescribed here.

> **A COMPARISON SHOULD RETAIN ITS ARITHMETIC AND ITS LIMITS**
>
> An illustrative matched comparison records complete on-time orders of 88 from 100 eligible orders before and 94 from 100 after for the pilot group. The comparator records 90 from 100 before and 92 from 100 after. The raw pilot change is 6 percentage points; the comparator change is 2 percentage points. Their difference is 4 percentage points, calculated as (94 minus 88) minus (92 minus 90). This is illustrative arithmetic, not a causal estimate from a GFT pilot.
>
> To interpret that difference, GFT would examine whether earlier trends and concurrent changes support the comparison. Promotions, order sizes, price changes, different due-date promises and altered suppliers can invalidate an apparently matched contrast. Unequal denominators require calculation from the actual rates, not subtraction of order counts. Shared customers or capacity can transmit effects between groups. The example intentionally supplies no confidence interval because no real sampling design or independence evidence has been established.

GFT would also specify how persistent stock affects the comparison. Material bought under an earlier policy remains available after a switch, and an unusually full opening warehouse can protect an otherwise weak candidate. Record opening pipeline and ageing in each group. Prefer contemporaneous comparable groups where feasible; where a crossover is used, explain the transition and the evidence retained during it. Do not assume that a brief pause removes all carry-over from earlier purchase commitments.

GFT recommends freezing the eligibility rule, primary service measure, cash measure and guardrails before the first treated decision. Record all eligible assigned decisions, including overrides, failed recommendations and missed reviews. Report outcomes by original assignment to answer whether the policy package worked as introduced. An additional analysis of decisions actually followed may explain mechanisms, but planner selection makes it a different question and potentially a biased comparison.

Log concurrent interventions as events rather than explanatory footnotes written after results arrive. A new fabric contract, revised credit terms, extra quality staff or a customer cancellation can change both groups differently. GFT would preserve those events in the result pack and ask whether the conclusion survives their exclusion or

separate analysis. Predefined interim reviews should protect operations without repeatedly searching for a favourable stopping date. A pilot interrupted by a safeguard breach remains an interrupted pilot.

13.6

The scorecard must reveal who paid for the improvement

GFT proposes a joint service and cash gate, with stock truth as an admission condition. Service should improve by the pre-agreed criterion while cash exposure improves on the same operating boundary. A reduction in cash that only maintains service may still justify a separate owner-approved efficiency decision. It should be reported in those terms rather than promoted as evidence that both improved. Keep the acceptance language aligned with the actual objective and observed uncertainty.

Define cycle service, immediate unit fill and complete order delivery separately. The inventory teaching material distinguishes stockout-cycle and fill measures; GFT's proposed order measure additionally fixes the confirmed due date and whole-order boundary. A material may have a high immediate fill rate while its occasional missing units stop complete customer orders. Do not average these different denominators into an invented universal service index. [S05]

Stock accuracy should be reported by usable status and material criticality, alongside blind-count coverage and unresolved discrepancies. The cycle-counting procedure from North Carolina separates recording and approval actions. GFT recommends using that distinction as a control example, not as an apparel performance benchmark. The inventory chapter defines the full ledger method; the pilot scorecard should reveal whether apparent policy improvement coincided with deterioration in the underlying record. [S27]

Separate physical exposure from accounting valuation and cash. The IFRS overview states a lower-of-cost-and-net-realisable-value basis for inventories. That accounting principle does not make a write-down a cash receipt. GFT recommends reconciling opening stock, receipts, consumption, adjustments and closing stock in both quantity and value, then reconciling deposits, supplier payments, payable dates and outstanding commitments separately. The inherited overview has no established publication date and is accounting context, not local professional advice. [S09]

GFT would measure average usable stock, quarantined stock, ageing, material-caused stoppage, expedite spending and peak cash-at-risk. State whether average stock uses time-weighted events or sampled closing balances. Do not omit goods in transit, committed but unpaid orders or customer-owned material when they affect the stated boundary. Age should follow the agreed physical receipt or ownership event rather than reset after a warehouse transfer. Record terminal residuals so a shorter trial cannot export its problem into the following period.

The proposed scorecard in **Exhibit 57**, p. 202 retains measurement families rather than supplying fictitious targets. The owner must approve thresholds and acceptable uncertainty against the actual commercial promise before running a field comparison.

EXHIBIT 57 / PROPOSED METRIC DEFINITIONS, NO OBSERVED VALUES

Service and cash gains remain conditional on supplier and worker safeguards

Measure	Boundary	Review evidence
Immediate fill	Served / eligible requested units	Request and issue events
Order OTIF	Complete on-time / due orders	Original promise and dispatch
Stock truth	Blind counts by usable state	Counts and open discrepancies
Material stops	Hours stopped for missing kits	Cause-coded production log
Stock exposure	Average, aged and held stock	Timed balances and lot age
Cash exposure	Payments and firm commitments	Finance reconciliation
Recovery cost	Expedites and emergency work	Approved costs and invoices
Safeguards	Pay dates and unscheduled hours	Supplier and worker checks

ILLUSTRATIVE [S05] [S09] [S10] [S20] [S27] Illustrative GFT measurement contract. Thresholds require owner approval before a trial. Evidence sources motivate method and safeguards; none supplies local pilot outcomes. Cash includes dated commitments, not inventory valuation alone.

The ILO's global supplier survey provides a reason to examine purchasing practices and working conditions together, with the limits of an older respondent survey retained. OECD's garment guidance places adverse impacts within responsible sourcing. GFT proposes supplier-payment timeliness, disputed deductions, emergency subcontracting and unscheduled overtime as safeguards alongside inventory outcomes. These are proposed checks; the cited publications do not establish the incidence or effect size of such harms in this programme. [S10] [S20]

A lower warehouse balance is not a whole-chain gain if material is merely held unpaid upstream or rushed through unrecorded outsourced work. GFT recommends asking who held the stock, financed it and absorbed the schedule change. Procurement should reconcile supplier terms and invoice settlement; production should reconcile emergency hours and outside conversion; finance should check whether avoided purchases simply become overdue obligations. Confidential worker and supplier feedback should be available without requiring a complaint through the policy's sponsor.

13.7

A failed gate blocks expansion until its cause is resolved

GFT recommends registering a decision gate with explicit outcomes: continue within scope, remediate and repeat, stop the candidate, or declare evidence insufficient. Fix the allowed materials, purchase authority, cash limit, exception route and expiry before starting. Missing inputs are not favourable outcomes. A small positive average with unresolved service uncertainty is not a pass. Expansion should require evidence relevant to the added authority and operating conditions, rather than a general endorsement of the software.

GFT NOTE

Approve the evidence contract before approving the pilot. Give a named operating role authority to pause recommendations when stock truth, customer service, cash limits or supplier and worker safeguards breach their agreed conditions. Preserve the frozen incumbent fallback and an auditable exception record. If the comparison is inconclusive, improve the information or redesign the trial before widening authority. A failed pilot blocks scaling until its cause is resolved and the revised policy is tested.

Distinguish an implementation fault from a policy limitation. GFT would investigate whether the recommendation used the right inputs, arrived before the decision deadline and remained feasible after rounding and approval. Then ask whether its objective and assumptions fitted the actual promise. A parameter adjustment, a repaired capture process and a changed customer commitment are different interventions. Version each change and define which earlier evidence remains applicable. Do not combine pre-repair and post-repair outcomes without showing the break.

A pilot may stop because a serious safeguard condition is triggered even when the cost comparison is favourable. GFT recommends preserving that outcome, the immediate containment action and the investigation. Remediation should demonstrate that the cause was addressed and that stopping and recovery worked. A missing recurrence during a quiet period is weak evidence if the original triggering condition has not reappeared. Retest the mechanism through an appropriate controlled exercise before relying on it again.

Publish claim labels beside findings. GFT proposes distinguishing equation-checked, simulated, historically replayed, shadow-observed and pilot-observed evidence. Add independently checked only when a separate reviewer has actually examined the stated records and method. Identify the reviewer role, scope, version and unresolved findings without implying certification. Keep assumptions and constructed examples visibly illustrative. No stage in this chapter is reported as a completed GFT operating test.

An inconclusive comparison needs its own completion decision. GFT would ask whether uncertainty comes from too few independent events, missing requests, unstable costs or intervention spillovers. More observations may resolve sampling uncertainty, but they will not repair a systematically missing denominator. Retain the bounded incumbent or planner-reviewed rule while the relevant cause is addressed. State the additional evidence needed and a review condition, rather than leaving a trial permanently active without an accountable decision.

GFT also recommends checking transferability before extending a successful bounded result. A locally sourced repeat trim, an imported seasonal fabric and customer-owned material expose different payment, substitution and residual-stock risks. Identify which characteristics the next group shares with the tested group and which require fresh calibration or a narrower approval limit. A pilot can justify continued operation within its tested boundary while remaining insufficient for a different material or commercial promise.

The result pack should contain the protocol, data lineage, exclusions, policy versions, parameter choices, paired comparisons, uncertainty, incident log and finance reconciliation. GFT recommends a short decision statement explaining which conditions passed, which failed and what authority follows. Preserve negative results as part of the same record. Connect the recurring review to the programme's performance-management design and reproducibility controls, while leaving the implementation roadmap to the closing chapter.

The acceptance question remains practical: can the factory make a better commitment without concealing service failures, outstanding cash or transferred harm? This chapter supplies a proposed evaluation method and illustrative arithmetic, not investment, legal or tax advice. Until an operating comparison answers that question within its stated limits, the honest result is a bounded recommendation awaiting proof.

A smaller buffer earns its place when the promise survives and the cost stays visible.

CHAPTER SOURCE NOTES

S22 arXiv, Chen, Jiang, Zhang and Zhou. 2023-10-31.

S23 arXiv, Sarit Maitra. 2024-06-27.

S19 National Institute of Standards and Technology. 2022-01-05.

S05 MIT OpenCourseWare. 2006.

S24 arXiv, Michna, Nielsen and Nielsen. 2015-03-20.

S25 OECD. 2026-02-26.

S27 North Carolina Office of the State Controller. 2023-10-24.

S09 IFRS Foundation. n.d..

S10 International Labour Organization. 2017-06-09.

S20 OECD. 2018-03-07.

GREEN FUTURE TECH

LIMITED / HONG KONG SAR

14

PART IV: PROVE THE POLICY BEFORE WIDENING ITS AUTHORITY / RECOMMENDATIONS

Build stock truth first, prove replenishment second and scale only after the gate holds

GFT would implement a phased inventory decision system whose authority grows only with measured reliability, service and cash evidence.

CHAPTER 14
DECISION TEST

GFT would implement a phased inventory decision system whose authority grows only with measured reliability, service and cash evidence.

An emerging-market apparel manufacturer should never deploy automated ordering or multi-location inventory optimisation on top of an unreconciled stock ledger. [S07] [S1402] Green Future Tech Limited (GFT), Hong Kong SAR, recommends a four-stage gated sequence in which operational authority expands only after the factory proves that its physical records, replenishment rules, and working-capital controls hold under real operating exceptions. [S05] [S27] [S31] Phase 0 locks the item, unit, ownership, and bill-of-material boundary before capital is committed; Phase 1 launches transaction capture, quality-release discipline, and planner-reviewed reorder rules; Phase 2 tests candidate policies in shadow mode and a bounded pilot against pre-registered service and cash gates; and Phase 3 extends the verified baseline to risk pooling, manufacturing postponement, and supplier contract flexibility. [S07] [S08] [S17] [S23]

This staged discipline connects the macro thesis of the paper to the factory floor. Across Cameroon, about 95 percent of cotton fibre is exported raw while about 5 percent is processed locally, according to the inherited register, and the national development strategy restates a policy target of 50 percent local processing by 2030. [R06] [R07] Those benchmarks frame the industrial opportunity, yet they do not tell a cut-make-trim garment plant when to release cash for woven twill, dyed jersey, sewing thread, zips, labels, or cartons. [R06] [S07] Within GFT's three-part Cameroon textile programme, prepared on 31 January 2024 with source-refresh notes retained in 2026, the feasibility design defines the product mix and cash envelope, the performance-management design specifies operating cadences and accountability, and the digital strategy introduces decision support only after transaction truth works. [S01] [S1402] No factory deployment or realised saving is claimed here; instead, this chapter sets out the roadmap, role charter, measurement contract, and stop rules required to make inventory policy verifiable. [S05] [S09]

14.1

Set up the stock boundary before committing capital

Before a garment factory signs its first bulk fabric purchase order or configures a replenishment formula, leadership must lock the boundary of what counts as identifiable, usable stock. [S07] [S27] In GFT's proposed Phase 0 setup stage, engineering, stores, quality, and finance teams establish the master records and physical baselines that every ordering calculation requires. [S07] [S1402] Skipping this preparation mixes incompatible fabric widths, unapproved shade lots, and inconsistent units of measure inside the same item code, stalling production lines even while the balance sheet reports ample material. [S07] [S09]

14.2 ITEM IDENTITY, BILL-OF-MATERIAL VERSIONS AND STOCK STATES FORM THE PREREQUISITE REGISTER

Phase 0 builds an unambiguous item and specification register for every raw material, trim, sub-assembly, and finished stock-keeping unit (SKU). [S01] [S07] Each finished SKU distinguishes style, construction version, colourway, and size rather than pooling demand at the parent style level. [S01] [S06] Fabric and trim definitions specify fibre composition, weight in grams per square metre, cuttable width, shrinkage allowance,

and shade-lot tolerances. [S07] Because cutting rooms cannot mix dye lots within a garment bundle without visible shade defects, the stock ledger tracks lot and roll identity inside each item code. [S07] [S31]

Phase 0 also standardises unit conversions and versioned bills of materials. [S07] Suppliers quote knitted fabric in kilograms, shipping documents record rolls, racks store linear metres, and sewing lines consume cut pieces. [S07] [S29] Every bill of material carries an engineering version specifying gross component consumption per garment, measured or estimated cutting waste from WP06, and permitted substitutions. [S07] [S1402] Before transactions begin, the database enforces mutually exclusive status buckets: `available` (inspected, released, and unreserved), `reserved` (allocated to a confirmed order), `quarantine` (awaiting inspection), `rejected` (failed quality release), and `customer-supplied` (held in bailment without factory title). [S07] [S09] [S27] Only `available` stock may be netted against new demand when calculating available-to-promise quantities or reorder triggers. [S05] [S07]

14.3 BASELINE PHYSICAL WALL-TO-WALL RECONCILIATION AND LOCAL INPUT COLLECTION PRECEDE PARAMETER ENTRY

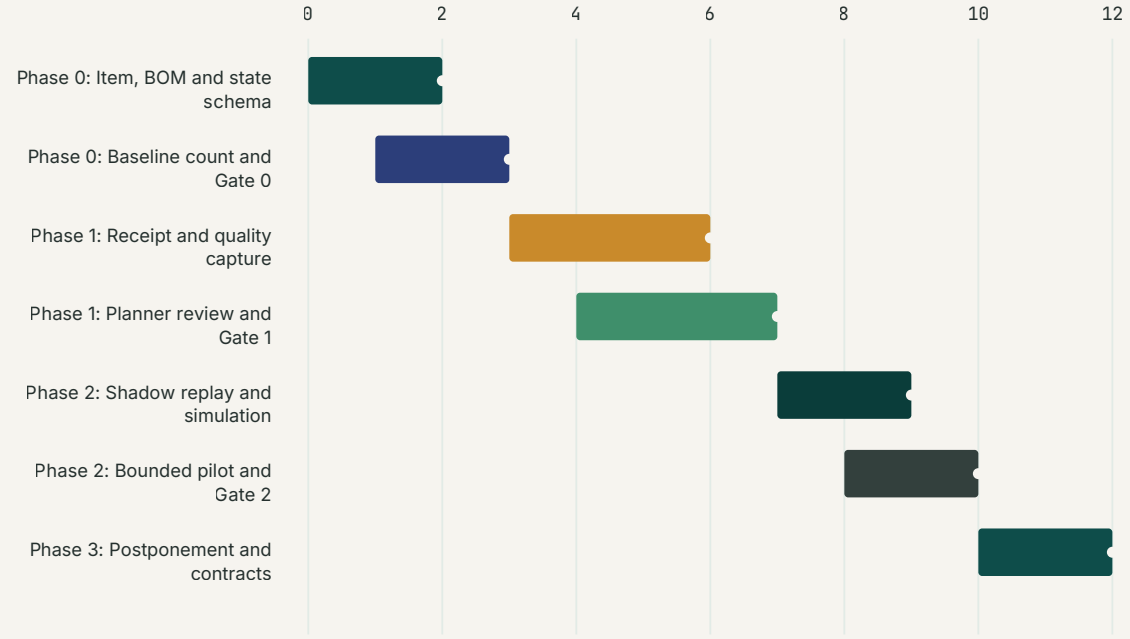
Phase 0 closes with a wall-to-wall physical count and structured local data collection. [S27] [S1402] Following public-sector cycle-counting control principles, physical verification separates the counter from the supervisor who reconciles the ledger and approves adjustments. [S27] Every roll, cone, and carton is tagged with its verified item code, lot identifier, measured quantity, ownership status, and quality state before the opening balance is certified. [S27] [S31]

Concurrently, procurement and finance compile dated local operational parameters. [S12] [S1402] Every supplier is documented with written minimum order quantities, pack multiples, quoted lead times, Incoterms 2020 delivery points, and payment milestones across deposits, shipment, and invoice acceptance. [S03] [S18] [S29] For imported inputs, planners establish a timestamp template informed by World Customs Organization time-release methods so shipments record vessel arrival, declaration, inspection, tax settlement, gate release, inland carriage, and factory quality release as separate intervals. [S12] [S30] Unmeasured route parameters or carrying-cost components remain explicit assumptions with a named owner role and review date. [S09] [S1402]

The progression of this preparatory work and its relationship to later stages is structured in **Exhibit 58**, p. 208. [S05] [S07] [S27] [S31]

EXHIBIT 58 / ILLUSTRATIVE PROGRAMME STAGE INDEX (0 TO 12)

The implementation roadmap advances through evidence gates



ILLUSTRATIVE [S05] [S07] [S27] [S31] Illustrative GFT proposed sequence across four gated phases (Phase 0 setup, Phase 1 capture and review, Phase 2 shadow and pilot, Phase 3 network and contract optimisation). Positions 0 to 12 represent sequential stage gates rather than promised calendar months or a deployed factory schedule.

As **Exhibit 58**, p. 208 shows, phases overlap only where data collection prepares the evidence gate for the next tier, never where authority is widened prematurely. [S07] [S27] The horizontal index from 0 to 12 is an illustrative sequence scale rather than a fixed twelve-month calendar promise: if a plant needs extra time at index 3 to resolve receiving discrepancies or unverified bill-of-material yields, Gate 0 blocks entry into Phase 1 replenishment rules until the boundary is clean. [S07] [S27] [S1402] This roadmap exhibit cannot establish how many calendar weeks a specific Cameroon facility will need, as duration depends on SKU count, supplier responsiveness, and storehouse staffing. [S1402]

14.4

Launch with usable-stock capture and simple review rules

Once Gate 0 confirms verified opening balances, versioned bills of materials, and explicit stock states, the factory enters Phase 1. [S07] [S27] Phase 1 has two objectives: maintaining transaction truth across every physical movement from the receiving bay to the sewing line, and operating simple, planner-reviewed replenishment rules that prevent line-stopping shortages and cash-draining over-orders without complex optimisation software. [S05] [S27] [S31]

14.5 TRANSACTION CAPTURE AT RECEIVING, INSPECTION RELEASE AND LINE ISSUE ESTABLISHES USABLE STOCK

Inventory accuracy degrades whenever physical movements outrun transaction recording. [S15] [S31] Phase 1 enforces a four-event capture chain. [S07] [S31] First, upon gate arrival, the storekeeper logs a physical receipt against the purchase order, recording arrival time, carrier reference, and pack count; following the ship-to-receive reconciliation principles in the GS1 and Auburn University Project Zipper study, any mismatch between shipping documents and physical tallies generates an immediate discrepancy record rather than a silent edit. [S27] [S31] Second, inbound goods enter `quarantine` automatically upon unloading and cannot be allocated until quality assurance completes physical and shade inspection and signs a release voucher moving accepted stock to `available`. [S07] [S31]

Third, movement from the warehouse to the cutting room requires a complete-kit issue check. [S07] [S19] The planner verifies component completeness across the versioned bill of materials before releasing the work order, transitioning issued quantities from `reserved` stock to tracked work in progress (`WIP`). [S07] [S19] Fourth, unused roll remnants and off-shade panels pass through a return-inspection gate so partial rolls are re-measured in linear metres before returning to `available` stock. [S07] [S27] Continuous cycle counting scaled by item criticality verifies this chain: high-value A fabrics and line-stopping C trims are counted on blind sheets omitting book balances, and as retail RFID cycle-count studies confirm, accuracy depends on physical discipline, tag read completeness, and count duration rather than hardware alone. [S01] [S15] [S16] [S27] [S1401] Variances exceeding tolerance trigger a recount and root-cause classification before an authorised supervisor updates the ledger. [S27]

14.6 DAILY, WEEKLY AND MONTHLY OPERATING CADENCES LINK STOCK RECORDS TO HUMAN REVIEW

In Phase 1, replenishment uses auditable rules reviewed by a human planner on a fixed rhythm aligned with GFT's performance-management design. [S05] [S1402] Repeat fabrics and trims follow either a continuous-review reorder-point (`ROP`) rule or a periodic-review order-up-to (`OUT`) rule across A, B, and C segments, while institutional tenders and seasonal capsules follow time-phased material requirements planning (`MRP`) explosions tied to firm cutting dates. [S01] [S05] [S06] [S07] [S1401] Phase 1 embeds these rules in three cadences: [S27] [S1402]

- **Daily shift reconciliation (storekeeper, quality lead, cutting supervisor):** Reconcile receipts, quarantine lots, complete-kit releases for the next two shifts, and line starvation or quality hold events. [S07] [S19] [S31]
- **Weekly replenishment review (planner and procurement lead):** Compute inventory position from usable `available` stock plus firm receipts less reservations, compare against reorder points and order-up-to levels, apply minimum lots and pack multiples, and submit proposals for finance sign-off. [S03] [S05] [S1401]
- **Monthly ledger and ageing audit (finance controller, planner, factory manager):** Review cycle-count accuracy and variance by state, audit overrides, inspect aged lots against IAS 2 net realisable value rules, and confirm exposure stays within the cash ceiling. [S09] [S27] [S1402]

GFT NOTE

Never automate replenishment during Phase 1 to compensate for late transaction entry or uninspected quarantine stock. [S15] [S31] If storekeepers post issues late or quality holds sit inside `available` stock, an automated engine will miss real shortages or order twice against phantom deficits. [S05] [S27] Fix the voucher, barcode, and cycle-count loop first, and keep every order under planner review until Gate 1 holds. [S27] [S1402]

14.7

Scale only after policy comparison survives normal exceptions

Passing Gate 1 requires documented proof that physical-to-book stock accuracy, quarantine segregation, and weekly planner review have operated reliably across consecutive monthly audit cycles. [S27] [S1402] Only then does the factory advance to Phase 2, where statistical safety-stock calibration, demand-forecast integration, and constrained order optimisation are tested against the Phase 1 baseline. [S05] [S22] [S23]

14.8 SHADOW-MODE REPLAY AND STOCHASTIC SIMULATION TEST CANDIDATE RULES WITHOUT RISKING CASH

In early Phase 2, candidate replenishment rules are evaluated offline through point-in-time historical replay and discrete-event stochastic simulation, following Chapter 13. [S19] [S22] [S23] Point-in-time replay requires that the engine see only the forecast vintage, usable stock, open orders, and lead-time sample recorded on each review date. [S22] [S23] As inventory learning research shows, demand and supply observations are frequently censored: stockouts truncate recorded sales below requested demand, and open orders right-censor lead times. [S22] Ignoring censoring overstates the benefit of stock-reduction rules. [S22]

Phase 2 also uses repeated stochastic simulation to stress-test reorder points, periodic-review intervals, and seasonal critical-fractile rules across synthetic demand and lead-time paths. [S05] [S06] [S23] Integrating random and conditional sampling across high-demand and delayed-supply regimes, and representing finite cutting and sewing buffers where breakdowns or quality holds affect flow, reveals upper-tail disruption behaviour. [S19] [S23] [S24] Next, the candidate engine runs in live **shadow mode**, logging weekly recommended quantities alongside the human planner's orders without transmitting supplier instructions. [S23] [S1402]

14.9 BOUNDED FIELD PILOTING INTEGRATES FORECAST DISTRIBUTIONS WHILE PRESERVING MANUAL FALLBACK

Once shadow comparisons confirm that the policy respects pack constraints, avoids duplicate ordering during outages, and improves the shortage-versus-cash trade-off, Phase 2 launches a bounded field pilot. [S03] [S05] [S23] The pilot consumes demand and residual error distributions from WP04 under WP14 reproducibility controls, assigning a matched group of repeat fabrics and trims to the calibrated policy alongside a baseline group. [S01] [S05] [S23] Every recommended lot is checked against the cash ceiling, rack capacity, and supplier minimums; if a WP04 feed fails validation, a cycle count uncovers a discrepancy, or a multi-item allocation is infeasible, the SKU reverts immediately to its Phase 1 manual rule. [S03] [S05] [S09] [S27] [S1402]

SHADOW-TO-PILOT AND GATE PROMOTION VERIFICATION PROTOCOL

To advance an item segment from Phase 1 manual review through Phase 2 shadow mode to Gate 2 approval, GFT requires four sequential checks: [S05] [S22] [S23]

1. **Data-integrity check:** Confirm exact-match cycle-count accuracy meets the pre-registered threshold and zero unapproved quarantine lots enter `available` stock. [S15] [S27]
2. **Censored-data replay check:** Reconstruct inventory positions $IP_t = OH^{\text{avail}}_t + SR^{\text{firm}}_t - BO_t - R^{\text{net}}_t$ at each review date t using only data recorded on or before t, bounding lost-sales demand censoring and right-censored lead times. [S05] [S22]
3. **Stochastic stress check:** Evaluate candidate parameters across simulation replications with correlated demand surges and upper-quantile customs or transport delays, recording cycle service, unit fill rate, and peak cash-at-risk. [S05] [S23] [S24]
4. **Live pilot safeguard check:** Verify across the pilot window that unit fill rate, customer `OTIF`, and complete-kit readiness hold within tolerances while pre-shipment cash-at-risk stays within budget, with zero increase in unscheduled worker overtime or delayed supplier payments. [S05] [S10] [S11] [S20]

14.10

Optimise network and contractual choices after the stock record is dependable

Only after passing Gate 2, proving that item-level replenishment rules work on verified stock data and survive supply exceptions, should leadership enter Phase 3. [S07] [S08] [S17] Phase 3 widens the decision boundary to structural supply-chain choices: risk pooling across shared materials, manufacturing postponement, multi-location distribution requirements planning (`DRP`), and risk-sharing supplier contracts. [S07] [S08] [S17] Attempting these manoeuvres before Phase 1 and Phase 2 hold compounds coordination failures across warehouses and counterparties. [S07] [S1402]

14.11 POSTPONED MANUFACTURING AND SHARED GREIGE OR DYED FABRIC BUFFERS REQUIRE VERIFIED CONVERSION READINESS

In apparel making, holding a pooled buffer of undyed greige fabric, white base fabric, or shared core-dyed twill and postponing final dyeing, printing, cutting, or badging until firm customer orders arrive can reduce finished-goods overage risk. [S06] [S17] Empirical research on Japanese apparel firms shows that postponement is conditional: manufacturing postponement improved inventory performance when aligned with operational control, whereas product-planning postponement could increase inventory when downstream retail control was externalised. [S17]

GFT evaluates any Phase 3 postponement or fabric-pooling proposal as a coupled inventory-and-capacity test rather than a simple square-root calculation. [S05] [S17] [S19] Pooling greige or base-dyed fabric reduces safety-stock variance only if downstream styles share identical technical fabric specifications under versioned bills of materials, and if remaining conversion stages (piece dyeing, printing, quality release, cutting, sewing, and finishing, coordinated with WP15 and WP20) fit inside

the customer delivery window. [S07] [S17] [S19] If piece dyeing or curing adds rework variance or batch queues, upstream fabric pooling can increase work-in-progress congestion and late deliveries. [S17] [S19]

14.12 SUPPLY-CHAIN CONTRACTS AND MULTI-LOCATION DISTRIBUTION RULES MUST PASS WHOLE-CHAIN CASH AND ETHICS CHECKS

Phase 3 evaluates multi-location replenishment and flexible supplier contracts across the whole chain. [S07] [S08] When a manufacturer operates an up-country fabric store and a coastal garment plant, distribution requirements planning (DRP) coordinates echelon stock rather than letting each site pad an independent reorder point. [S07] [S24] Because stochastic replenishment delays interact with local ordering rules to amplify upstream order variance (the bullwhip effect), echelon visibility prevents a coastal cutting room and an up-country storehouse from double-counting the same buffer. [S07] [S24]

When procurement negotiates framework agreements with call-off options, reserved mill capacity, split deliveries, or buyback clauses, Phase 3 models the contract across both balance sheets. [S08] [S18] Buyback and quantity-flexibility terms reallocate overage and underage risk between buyer and supplier, but pushing inventory onto an under-capitalised supplier without fair compensation or liquidity support can trigger insolvency, unannounced subcontracting, or quality shortcuts. [S08] [S10] [S18] [S25] International Labour Organization garment studies show how volatile buyer order changes, compressed lead times, and price pressure transmit into wage pressure and excessive worker overtime. [S10] [S11] Phase 3 therefore audits flexible purchasing terms against OECD garment due-diligence standards and verified subcontracting maps before adoption. [S20] [S25]

14.13

Assign decisions to accountable roles rather than to a dashboard

A recurring failure in industrial digitalisation is treating a software screen or an optimisation script as the decision-maker. [S1402] Dashboards display signals and solvers compute conditional trade-offs, but neither can inspect a damp carton, negotiate credit, or bear accountability for an obsolete stock write-down. [S09] [S27] [S1402] Across all four phases, GFT assigns every parameter, transaction event, override, and stage gate to a named functional role rather than to an automated system or an invented persona. [S27] [S1402]

14.14 SEGREGATION OF DUTIES PREVENTS SELF-APPROVING INVENTORY ADJUSTMENTS AND SILENT TARGET DRIFT

Sound inventory governance separates the role that executes physical custody from the role that reconciles records and approves valuation adjustments. [S09] [S27] GFT defines five functional roles across the inventory cycle: [S27] [S1402]

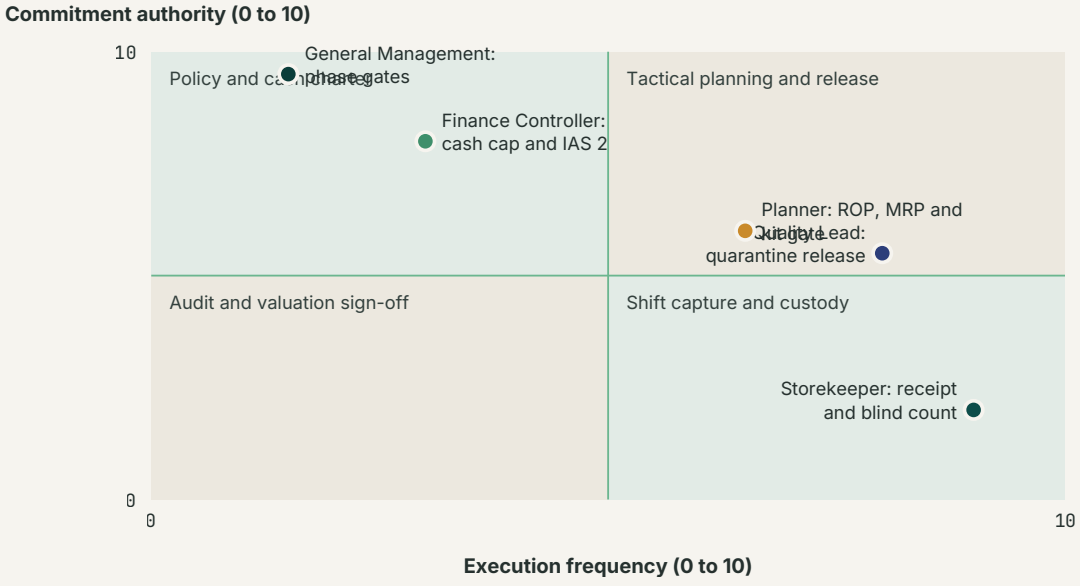
- **Warehouse Storekeeper (custody and capture):** Logs physical receipts, pack checks, bin put-away, blind cycle counts, and complete-kit picks; cannot alter specifications, release quarantine lots, or approve stock adjustments. [S27] [S31]

- **Quality-Assurance Lead (usability release):** Inspects inbound fabrics, trims, and returns, manages `quarantine` and `rejected` states, verifies shade-lot continuity, and signs the release moving material into `available` stock. [S07] [S31]
- **Supply and Production Planner (parameters and orders):** Maintains versioned bills of materials, segments items by purpose and criticality, runs weekly continuous-review, periodic-review, and MRP calculations, manages the complete-kit gate, and proposes orders; cannot breach cash caps or write off stock. [S01] [S05] [S07] [S27]
- **Finance Controller (cash caps and valuation):** Sets pre-shipment cash-at-risk limits, validates holding-cost and currency assumptions, approves discrepancy investigations, and governs IAS 2 lower-of-cost-and-net-realisable-value provisions. [S02] [S09] [S18] [S27]
- **Factory General Management (charter and gates):** Approves segment service targets, authorises progression or rollback across Phase 0 to Phase 3 gates, resolves cash-versus-service escalations, and enforces worker and supplier safeguards. [S05] [S20] [S1402]

The allocation of these recurring responsibilities across operational frequency and commitment authority is mapped in **Exhibit 59**, p. 213. [S27] [S1402]

EXHIBIT 59 / ILLUSTRATIVE GOVERNANCE SCALE (0 TO 10)

Every recurring decision has an accountable role



ILLUSTRATIVE [S27] [S1402] Illustrative GFT governance mapping based on segregation-of-duties principles in S27 and operational practice in S1402. Horizontal axis scores operational execution frequency (0 = quarterly/gate review, 10 = continuous shift capture); vertical axis scores capital and policy commitment authority (0 = execution within fixed rules, 10 = factory-wide cash and gate charter).

As **Exhibit 59**, p. 213 illustrates, high-frequency shift actions (dock receipt logging, blind counting, and quarantine inspection) are separated from high-authority capital and ledger adjustments (cash-ceiling sign-off, discrepancy write-offs, and phase-gate promotion). [S27] [S1402] The coordinates from 0 to 10 are illustrative governance

indices visualising segregation of duties rather than empirical workload measures, and cannot replace formal corporate delegations or statutory audit procedures. [S09] [S27]

14.15 EVERY PARAMETER CHANGE AND QUANTITY OVERRIDE REQUIRES AN IMMUTABLE DECISION LOG

To make role accountability operational, GFT's architecture requires an append-only decision log for every policy change and order override. [S27] [S31] [S1402] Whenever a planner modifies a reorder point, alters a lead-time assumption, or overrides a recommended order quantity, the log records six fields: decision identifier, timestamp, accountable role, frozen input snapshot (usable stock, open receipts, reservations, forecast vintage, and policy version), quantitative change, and structured reason code. [S03] [S05] [S31] Monthly reviews inspect the distribution of override reasons: if planners repeatedly override a rule upward because shipments arrive damaged or late, leadership recalibrates the supplier lead-time and acceptance distribution or re-qualifies the supplier rather than hiding the override. [S27] [S30] [S1402]

14.16

Measure the promise, the buffer and the cash on a common boundary

An inventory policy can look successful on a narrow warehouse metric while quietly damaging delivery, exhausting cash, or transferring pressure onto workers and suppliers. [S05] [S10] [S20] GFT evaluates inventory performance through a balanced measurement contract in which every metric shares an explicit physical boundary, time window, numerator, denominator, and source event. [S05] [S09] [S30] [S31]

14.17 NUMERATORS, DENOMINATORS AND EVENT TIMESTAMPS PREVENT METRIC MANIPULATION

The first pillar is **stock truth by state**: exact-match bin accuracy, absolute discrepancy across `available`, `reserved`, and `quarantine` states, and discrepancy closure time. [S15] [S16] [S27] The second pillar is **usable-receipt supply timing**: following World Customs Organization time-release principles and World Bank logistics tracking, supplier lead time runs from purchase-order approval to post-inspection `available` release, reported via the median, 90th percentile, and right-censored order count. [S12] [S22] [S30]

The third pillar separates **item availability from customer promise fulfilment**. [S05] Cycle service level measures replenishment cycles ending without a stockout, whereas unit fill rate measures requested units satisfied immediately from usable stock divided by total requested units (including documented stockout requests). [S05] [S22] Neither equals customer order on-time-in-full (`OTIF`), which requires complete order delivery across all sizes and trims by the confirmed promise date. [S05] [S07] Alongside external `OTIF`, the plant tracks **complete-kit readiness** and **material-caused line stoppage hours** so lines are not starved by a missing C-item trim. [S07] [S19] [S1401] The fourth pillar tracks **working capital and obsolescence**: pre-shipment **cash-at-risk** (cash in deposits, transit, quarantine, and work in progress not yet backed by an accepted receivable under programmes such as IFC post-shipment supplier finance), **aged stock share**, **IAS 2 write-downs** to net realisable value, and **expedite spend**. [S03] [S09] [S18]

14.18 WORKER OVERTIME AND SUPPLIER PAYMENT GUARDRAILS STOP FALSE INVENTORY SAVINGS

The fifth pillar enforces **social and commercial guardrails**. [S10] [S11] [S20] Cutting fabric and trim buffers without reducing lead-time uncertainty causes late kit releases. [S07] [S24] Recovering late kits by forcing sewing operators into unscheduled overtime or delaying supplier payments beyond agreed terms creates an accounting illusion paid for by workers and suppliers. [S10] [S11] [S20] Grounded in International Labour Organization garment findings and OECD due-diligence guidance, GFT treats material-delay overtime hours and overdue supplier payment days as hard guardrails: breaching either fails the policy gate regardless of stock reductions. [S10] [S11] [S20]

The specification of this balanced scorecard appears in **Exhibit 60**, p. 215. [S05] [S09] [S20] [S27] [S30] [S31]

EXHIBIT 60 / DEFINED OPERATIONAL UNITS PER METRIC

GFT would measure a balanced inventory outcome

Metric and pillar	Numerator and denominator	Boundary and cadence	Accountable owner
Usable stock accuracy (Truth)	Counted item-lot bins within tolerance / Total blind-counted bins	By stock state; weekly cycle count	Storekeeper and Finance
Usable-receipt lead time (Supply)	Timestamp of QA release minus PO approval (median, P90, censored count)	Order-to-usable-release; monthly	Planner and Quality Lead
Unit fill rate (Service)	Requested units filled immediately from available stock / Total requested units	Item-location boundary; weekly	Supply Planner
Complete-kit stoppage (Flow)	Line hours halted by missing approved material or trim / Scheduled line hours	Cutting and sewing release; daily	Planner and Production
Pre-shipment cash-at-risk (Cash)	Deposits + transit + quarantine + unbilled WIP/FG less accepted receivables	XAF working-capital ledger; monthly	Finance Controller
Aged and NRV exposure (Cash)	Cost of stock exceeding age limit or above IAS 2 NRV / Total book stock	Lot-level valuation; monthly	Finance Controller
Worker and supplier guardrails (Ethics)	Material-delay overtime hours and overdue supplier invoices / Agreed limits	Plant payroll and AP ledger; monthly	General Management

ILLUSTRATIVE [S05] [S09] [S20] [S27] [S30] [S31] Illustrative GFT proposed measurement contract synthesising foundational inventory definitions (S05), IAS 2 valuation boundaries (S09), OECD due-diligence guardrails (S20), cycle-count controls (S27), WCO time-release intervals (S30), and ship-to-receive reconciliation (S31). No realised factory values are claimed.

As **Exhibit 60**, p. 215 demonstrates, every metric pairs a numerator with an explicit denominator, event boundary, cadence, and owner role. [S05] [S09] [S27] Evaluating all seven rows together at each gate prevents inflating fill rate via aged stock or squeez-

ing cash-at-risk via cutting stoppages and worker overtime. [S05] [S09] [S20] Because this table defines a measurement contract rather than reporting empirical results, numerical thresholds must be calibrated locally during Phase 0 and Phase 1 before Gate 2 is judged. [S1402]

WORKED EXAMPLE

Illustrative reconciliation of a gated policy transition from Phase 0 to Phase 2

Consider an illustrative case for a Cameroon uniform maker ordering dyed poly-cotton twill (`FAB-TWL-240-NVY`, 2,500 XAF/m) and sewing thread (`TRM-THD-402-NVY`, 1,500 XAF/cone). [S02] [S05] [S07] [R32] All figures are teaching assumptions; the CFA franc is pegged at 655.957 XAF per euro in the inherited register, and no GFT operating result is claimed. [R32] [S05]

Step 1: Phase 0 stock reconciliation. Unadjusted book stock reports 6,000 metres of navy twill ($15,000,000\text{ XAF}$). [S09] Phase 0 audit identifies three exclusions: 400 metres missing from unrecorded scrap ($1,000,000\text{ XAF}$), 600 metres in `quarantine` ($1,500,000\text{ XAF}$), and 1,800 metres `reserved` for a confirmed cut ($4,500,000\text{ XAF}$). [S07] [S27] [S31] Deducting these leaves $OH^{\text{avail}} = 3,200\text{ metres}$ ($8,000,000\text{ XAF}$), preventing a 2,800-metre ($7,000,000\text{ XAF}$) overstatement. [S07] [S27]

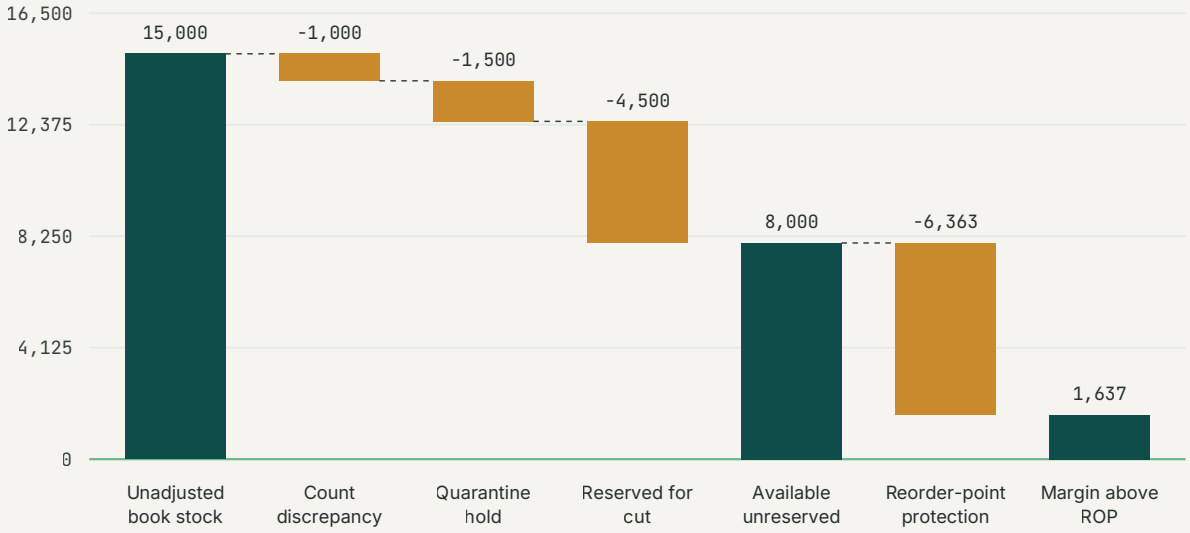
Step 2: Phase 1 and Phase 2 reorder calibration. Repeat twill demand has illustrative mean $\mu_D = 120\text{ m/day}$ and residual standard deviation $\sigma_D = 30\text{ m/day}$; usable-receipt lead time has mean $\mu_L = 16\text{ working days}$ and standard deviation $\sigma_L = 3\text{ working days}$. [S05] [S24] [S30] Lead-time demand variance is $\text{Var}(D_L) = 16(30)^2 + (120)^2(3)^2 = 14,400 + 129,600 = 144,000\text{ m}^2$, giving $\sigma_{D_L} \approx 379.47\text{ m}$ (90 percent from lead-time variance). [S05] [S24] At 95 percent cycle service ($z = 1.645$), safety stock is $SS \approx 625\text{ m}$ ($1,562,500\text{ XAF}$), and $ROP = 1,920 + 625 = 2,545\text{ m}$ ($6,362,500\text{ XAF}$). [S05]

Step 3: Cash and complete-kit gate. With $OH^{\text{avail}} = 3,200\text{ m}$, fabric sits 655 metres ($1,637,500\text{ XAF}$) above ROP. [S05] Yet the complete-kit check for the reserved 1,800-metre cut (requiring 360 thread cones worth $540,000\text{ XAF}$) finds only 200 `available` cones and 200 in `quarantine`. [S07] [S1401] Holding 200 extra cones at an illustrative 20 percent annual carrying rate costs $60,000\text{ XAF}$/year while protecting $4,500,000\text{ XAF}$ of cut fabric. [S02] [S07] [S1401]

The financial reconciliation from unadjusted book stock to usable unreserved fabric and safety protection in this illustrative case is shown in **Exhibit 61**, p. 217. [S02] [S05] [S09] [S31]

EXHIBIT 61 / THOUSAND XAF (ILLUSTRATIVE WORKED CASE)

State exclusions and safety buffers reconcile book value to free stock



ILLUSTRATIVE [S02] [S05] [S09] [S31] Illustrative arithmetic waterfall from the worked case for navy twill fabric (2,500 XAF/metre). Opening book value (15,000 thousand XAF for 6,000 m) less count discrepancy (-1,000), quarantine hold (-1,500), and confirmed reservation (-4,500) yields 8,000 thousand XAF (3,200 m) of available unreserved stock; deducting the reorder-point buffer (-6,363 thousand XAF for 2,545 m) leaves 1,637 thousand XAF (655 m) above the trigger.

As **Exhibit 61**, p. 217 shows, nearly half of the 15,000 thousand XAF opening book stock is unavailable once the 1,000 thousand XAF count error, 1,500 thousand XAF quarantine hold, and 4,500 thousand XAF reservation are deducted. [S07] [S09] [S31] Of the remaining 8,000 thousand XAF in `available` fabric, 6,363 thousand XAF covers expected demand and stochastic lead-time protection over the 16-day horizon, leaving 1,637 thousand XAF above the trigger. [S05] [S24] This waterfall illustrates the reconciliation method on constructed inputs rather than an observed Cameroon factory balance sheet. [S09]

14.19

Close with explicit conditions for changing the recommendation

A credible engineering recommendation must state in advance what evidence would cause GFT to halt expansion, revert to a simpler method, or recommend changing the commercial promise. [S05] [S23] [S1402] Widening automation when a gate fails accelerates the rate at which bad data or underfunded commitments generate losses. [S27] [S1402]

14.20 FAILED GATES TRIGGER SPECIFIC STRUCTURAL REMEDIES RATHER THAN WIDER SOFTWARE DEPLOYMENT

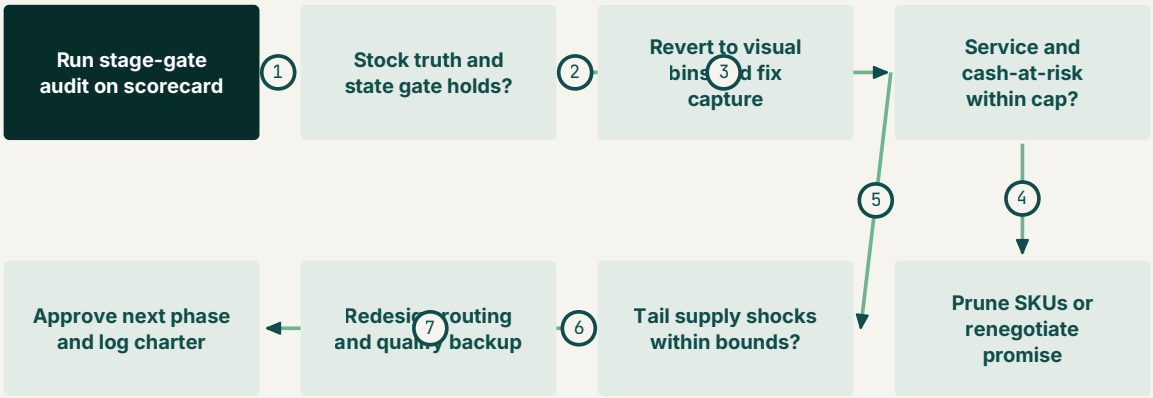
GFT defines three conditional fallback branches whenever a stage gate fails: [S05] [S27] [S1402]

1. **If physical-to-book capture cannot be trusted, revert to simpler manual policies:** If cycle counts in Phase 1 or Phase 2 show persistent discrepancies from unrecorded issues, power dropouts without voucher queues, or unsegregated quarantine stock, halt statistical reorder and solver deployment. [S15] [S27] [S31] Two-bin kanban racks for C trims and locked complete-kit cages for A fabrics with weekly manual counts outperform formulas that trust corrupted balances. [S27] [S1401] [S1402]
2. **If working capital cannot fund the buffer, lower commitment or renegotiate the promise:** If cash-constrained allocation in Chapter 11 and Phase 2 shows that pre-shipment cash cannot finance a 95 percent cycle service target across the active SKU range, no algorithm can invent liquidity. [S03] [S05] [S18] Leadership must prune the SKU assortment, negotiate smaller minimums or staged call-offs, shift seasonal lines to deposit-backed make-to-order terms, or lengthen quoted delivery lead times. [S01] [S03] [S08] [S18]
3. **If tail supply shocks dominate statistical buffers, redesign sourcing and routing:** If World Customs Organization time-release sampling and corridor tracking show that multi-week port, border, or supplier disruptions overwhelm feasible safety stock, holding larger piles is the wrong remedy. [S12] [S24] [S30] Instead, pre-qualify secondary regional or domestic textile sources as capacity develops, establish pre-cleared customs procedures, or structure contracts around contingent delivery windows tied to verified inbound release. [R06] [S13] [S25] [S30]

The decision logic linking stage-gate audit outcomes to these remedies is summarised in **Exhibit 62**, p. 219. [S05] [S22] [S23] [S27]

EXHIBIT 62 /

Evidence can require a simpler policy or a changed promise



1. Evaluate
2. Fail
3. Pass
4. Fail
5. Pass
6. Fail
7. Pass

ILLUSTRATIVE [S05] [S22] [S23] [S27] Illustrative GFT conditional decision tree governing stage-gate reviews across Phases 0 to 3. Every failed gate routes to a specific operational, commercial, or sourcing remedy rather than permitting unverified expansion.

As **Exhibit 62**, p. 219 shows, progression to `scale` requires passing three sequential verification gates: stock truth, cash-feasible service, and bounded supply-shock exposure. [S05] [S23] [S27] A failure at any node routes management to a targeted structural intervention (`manual`, `promise`, or `sourcing`) and requires a fresh gate audit before authority widens. [S27] [S1402] As a logical decision diagram, this exhibit defines the governance protocol rather than predicting which branch a pilot will trigger. [S23]

14.21 UNRESOLVED LOCAL INPUTS AND ADVISORY BOUNDARIES GOVERN ANY REAL PILOT COMMITMENT

To translate this framework into an operational Phase 0 launch in Cameroon or a comparable apparel corridor, the sponsoring enterprise must assemble six local evidence packages outside the scope of this paper: [S12] [S1402]

- **Verified item and bill-of-material files:** Technical specifications, fabric widths, weights, shrinkage tolerances, and cutting waste baselines coordinated with WP06 and WP20. [S07]
- **Dated supplier and logistics quotes:** Written price schedules, minimum order quantities, pack multiples, Incoterms 2020 delivery places, and empirical port and inland transit timestamps across Douala, Kribi, or northern corridors, coordinated with WP08. [S03] [S12] [S29] [S30]

- **Calibrated demand distributions:** Point-in-time demand forecasts, residual variance estimates, and stockout-censoring records from WP04. [S05] [S22]
- **Working-capital and financing terms:** Confirmed credit lines, deposit schedules, XAF conversion terms, and pre-shipment or post-shipment trade finance rules coordinated with WP13. [R32] [S18]
- **Shop-floor capacity and cadences:** Finite cutting, sewing, and finishing capacities, maintenance schedules, and shift routines linked to WP02 and WP15. [S19] [S28]
- **Resilient capture and governance:** Offline-capable barcode or RFID event logging, idempotent sync queues, and model audit trails aligned with WP07, WP14, and WP18. [S15] [S31]

SCOPE, VERIFICATION LIMITS AND PROFESSIONAL ADVICE NOTICE

This paper presents an engineering and operating-method framework by Green Future Tech Limited, Hong Kong SAR, drawing on its three-part Cameroon textile programme (feasibility study, performance-management system design, and digital strategy, prepared 31 January 2024 with source-refresh notes retained in 2026) and public academic, institutional, and standard-setting sources. [R06] [S05] [S09] No factory inventory ledger, shipment sample, or pilot outcome is claimed as an achieved GFT operating result, and third-party references (SODECOTON, CICAM, Camtext SA, McKinsey) remain statements of their respective sources rather than GFT engagements. [R03] [R09] [R12] [R27] Nothing here constitutes investment, legal, accounting, or tax advice; enterprises must obtain professional review before executing contracts, write-downs, customs declarations, or financing agreements. [S09] [S18] [S29]

Count the roll before trusting the ledger, prove the rule before trusting the model, and keep every order inside the cash and promise the factory can honour.

CHAPTER SOURCE NOTES

S07 MIT OpenCourseWare. 2006.

S1402 MIT OpenCourseWare. 2006.

S05 MIT OpenCourseWare. 2006.

S27 North Carolina Office of the State Controller. 2023-10-24.

S31 GS1 US and Auburn University RFID Lab. 2018.

S08 MIT OpenCourseWare. 2006.

S17 Japan Marketing Academy, Quarterly Journal of Marketing. 2020-03-31.

S23 arXiv, Sarit Maitra. 2024-06-27.

R06 Cameroon Trade Hub. Publication date unrecorded; inherited source-text check.

R07 Ecofin Agency. Publication date unrecorded; inherited source-text check.

S01 MIT OpenCourseWare. 2006.

S09 IFRS Foundation. n.d..

S06 MIT OpenCourseWare. 2006.

S29 International Chamber of Commerce. 2020.

S12 World Bank. 2025.

S03 MIT OpenCourseWare. 2006.

S18 International Finance Corporation. 2026.

S30 World Customs Organization. 2025.

S15 Association for Information Systems Electronic Library. 2013.

S19 National Institute of Standards and Technology. 2022-01-05.

S16 University of Modena and Reggio Emilia institutional repository. 2015.

S1401 MIT OpenCourseWare. 2006.

S22 arXiv, Chen, Jiang, Zhang and Zhou. 2023-10-31.

S24 arXiv, Michna, Nielsen and Nielsen. 2015-03-20.

S10 International Labour Organization. 2017-06-09.

S11 International Labour Organization. 2017-05-15.

S20 OECD. 2018-03-07.

S25 OECD. 2026-02-26.

S02 MIT OpenCourseWare. 2006.

R32 Wise; XE. Publication date unrecorded; inherited source-text check.

S13 World Trade Organization. n.d..

S28 MIT OpenCourseWare. 2012.

R03 SODECOTON. Publication date unrecorded; inherited source-text check.

R09 Ecofin Agency. Publication date unrecorded; inherited source-text check.

R12 Business in Cameroon. Publication date unrecorded; inherited source-text check.

R27 McKinsey, The State of Fashion 2025. Publication date unrecorded; inherited source-text check.

Provenance travels with the claim.

Available inherited records and all chapter-local source records are graded below. Identical source identities are consolidated without losing citation aliases. Missing annex references are excluded. Unrecorded dates are disclosed rather than invented.

R01 Inherited evidence R01 GRADE C

Business in Cameroon. Publication date unrecorded. businessincameroon.com

Cameroon seed-cotton production 2022-2024: 330,000 t (2022), 314,455 t (2023), about 340,000 t (2024).

Inherited source-text check; original source not independently re-opened for this edition.

R02 Inherited evidence R02 GRADE C

Fibre2Fashion; Business in Cameroon. Publication date unrecorded. fibre2fashion.com; businessincameroon.com

2025 seed-cotton projection above 350,000 t.

Inherited source-text check; original source not independently re-opened for this edition.

R03 Inherited evidence R03 GRADE C

SODECOTON. Publication date unrecorded. sodecoton.cm

Company profile: more than 220,000 cotton producers supported in northern Cameroon; nine ginning plants; two oil mills.

Inherited source-text check; original source not independently re-opened for this edition.

R04 Inherited evidence R04 GRADE C

Financial Fortune Media. Publication date unrecorded. financialfortunemedia.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R05 Inherited evidence R05 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R06 Inherited evidence R06 GRADE C

Cameroon Trade Hub. Publication date unrecorded. cameroontradehub.cm

About 95 percent of cotton fibre exported raw against 5 percent processed locally; industrialisation target of 50 percent local processing.

Inherited source-text check; original source not independently re-opened for this edition.

R07 Inherited evidence R07 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

SND30 restates the 50 percent local-processing target by 2030.

Inherited source-text check; original source not independently re-opened for this edition.

R08 Inherited evidence R08 GRADE C

Business in Cameroon. Publication date unrecorded. businessincameroon.com

Announced SODECOTON and CICAM merger to create a cotton-textile industrial pole targeted by 2030.

Inherited source-text check; original source not independently re-opened for this edition.

R09 Inherited evidence R09 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

Camtext SA integrated cotton-to-textile project, US$300 million (180 billion XAF); Panafritex and Arise with SODECOTON, CNPS and Marlo; 12 November 2025.

Inherited source-text check; original source not independently re-opened for this edition.

R10 Inherited evidence R10 GRADE C

Ecofin Agency. Publication date unrecorded. ecofinagency.com

Camtext SA: 12,000 t a year processed; about 12,000 jobs at the Dibamba hub and 3,000 at Garoua; return on investment near 11 years; optimal capacity around 2033.

Inherited source-text check; original source not independently re-opened for this edition.

R11 Inherited evidence R11 GRADE C

Business in Cameroon; MMI News. Publication date unrecorded. businessincameroon.com; mimimefoinfos.com

CICAM market share from about 85 percent three decades ago to about 5 percent today.

Inherited source-text check; original source not independently re-opened for this edition.

R12 Inherited evidence R12 GRADE C

Business in Cameroon. Publication date unrecorded. businessincameroon.com

CICAM capacity utilisation, July 2023 report: Garoua 33.7 percent, CICAM 1 48.9 percent, CICAM 2 10.83 percent; underutilisation cost about CFAF 2 billion in 2020.

Inherited source-text check; original source not independently re-opened for this edition.

R13 Inherited evidence R13 GRADE C

Business in Cameroon; MMI News. Publication date unrecorded. businessincameroon.com; mimimefoinfos.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R14 Inherited evidence R14 GRADE C

WageIndicator. Publication date unrecorded. wageindicator.org

Cameroon minimum wage (SMIG) XAF 41,875 a month in 2023, raised to XAF 43,969 a month from 23 February 2024, about US$76, a rise of 5 percent.

Inherited source-text check; original source not independently re-opened for this edition.

R15 Inherited evidence R15 GRADE C

wage.is. Publication date unrecorded. wage.is/cameroon

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R16 Inherited evidence R16 GRADE C

WageIndicator. Publication date unrecorded. wageindicator.org

Informal share of employment about 90 percent; standard 40-hour week, 48-hour maximum, overtime at 1.2 times.

Inherited source-text check; original source not independently re-opened for this edition.

R17 Inherited evidence R17 GRADE C

just-style. Publication date unrecorded. just-style.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R18 Inherited evidence R18 GRADE B

Congressional Research Service IF10149; AGOA (Wikipedia). Publication date unrecorded. congress.gov; en.wikipedia.org

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R23 Inherited evidence R23 GRADE C

MacMap; AfCFTA. Publication date unrecorded. macmap.org

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R24 Inherited evidence R24 GRADE C

Mordor Intelligence. Publication date unrecorded. mordorintelligence.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R25 Inherited evidence R25 GRADE C

Market.us. Publication date unrecorded. market.us

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R26 Inherited evidence R26 GRADE C

Statista. Publication date unrecorded. statista.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R27 Inherited evidence R27 GRADE B

McKinsey, The State of Fashion 2025. Publication date unrecorded. mckinsey.com

AI at scale in fashion: inventory cost down 35 to 40 percent; forecast accuracy up 20 to 30 percent; revenue per customer up 15 to 25 percent; forecast errors cut 20 to 50 percent; lost sales down up to 65 percent; 75 percent of executives prioritising; sector operating-profit potential US$150 to 275 billion over three to five years.

Inherited source-text check; original source not independently re-opened for this edition.

R28 Inherited evidence R28 GRADE C

DataIntelo. Publication date unrecorded. dataintelo.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R29 Inherited evidence R29 GRADE C

MDPI; industry inspection literature. Publication date unrecorded. mdpi.com

Manual defect inspection 60 to 70 percent accurate, missing 20 to 30 percent of defects; computer vision above 90 percent across more than 40 defect types; line-scan inspection up to 60 metres a minute.

Inherited source-text check; original source not independently re-opened for this edition.

R30 Inherited evidence R30 GRADE C

Design Journal (Taylor and Francis). Publication date unrecorded. tandfonline.com

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R31 Inherited evidence R31 GRADE C

arXiv 2509.04705. Publication date unrecorded. arxiv.org/abs/2509.04705

Source metadata retained for traceability. This entry is not cited in the chapter text.

Inherited source-text check; original source not independently re-opened for this edition.

R32 Inherited evidence R32 GRADE C

Wise; XE. Publication date unrecorded. wise.com; xe.com

CFA franc to US dollar reference rates; franc pegged to the euro at 655.957. The GFT study's method note places the franc near 590 to 610 per US dollar across the 2023 study window and near 563 to 577 more recently, and treats currency conversions as estimates.

Inherited source-text check; original source not independently re-opened for this edition.

c01-S01 Supply chain fundamentals and segmentation analysis GRADE A

MIT OpenCourseWare. 2006. https://ocw.mit.edu/courses/esd-260j-logistics-systems-fall-2006/resources/lect2/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S02 Inventory management I: Level demand, EOQ, sensitivity GRADE A

MIT OpenCourseWare. 2006. https://ocw.mit.edu/courses/esd-260j-logistics-systems-fall-2006/resources/lect7/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S03 Inventory management II: EOQ extensions and discounts GRADE A

MIT OpenCourseWare. 2006. https://ocw.mit.edu/courses/esd-260j-logistics-systems-fall-2006/resources/lect8/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S05 Inventory management IV: Probabilistic demand and safety stock GRADE A

MIT OpenCourseWare. 2006. https://ocw.mit.edu/courses/esd-260j-logistics-systems-fall-2006/resources/lect11/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S06 Inventory management VI: Fashion, style and perishable goods GRADE A

MIT OpenCourseWare. 2006. https://ocw.mit.edu/courses/esd-260j-logistics-systems-fall-2006/755a26f6aae-bae8928294146a2484cd4_lect14.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S07 Inventory management VII: MRP and DRP systems GRADE A

MIT OpenCourseWare. 2006. https://ocw.mit.edu/courses/esd-260j-logistics-systems-fall-2006/resources/lect15/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S08 Supply chain contracts GRADE A

MIT OpenCourseWare. 2006. https://ocw.mit.edu/courses/esd-260j-logistics-systems-fall-2006/resources/lect18/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S09 IAS 2 Inventories GRADE A

IFRS Foundation. n.d.. https://www.ifrs.org/issued-standards/list-of-standards/ias-2-inventories/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S10 Purchasing practices and working conditions in global supply chains: Global Survey results GRADE A

International Labour Organization. 2017-06-09. https://www.ilo.org/publications/purchasing-practices-and-working-conditions-global-supply-chains-global

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S14 Revisiting the Sources of Growth: Enhancing the Efficiency of the Port of Douala GRADE A

World Bank. 2015-02-12. https://www.worldbank.org/en/country/cameroon/publication/cameroon-economic-update-revisiting-the-sources-of-growth-enhancing-the-efficiency-of-the-port-of-douala

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S15 The Performance and Reliability of a RFID Cycle-Count: A Quantitative Approach from Fashion Retail GRADE B

Association for Information Systems Electronic Library. 2013. https://aisel.aisnet.org/icmb2013/25/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S18 Global Trade Supplier Finance GRADE A

International Finance Corporation. 2026. https://www.ifc.org/en/what-we-do/sector-expertise/trade-and-supply-chain-finance/global-trade-supplier-finance

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S20 OECD Due Diligence Guidance for Responsible Supply Chains in the Garment and Footwear Sector GRADE A

OECD. 2018-03-07. https://www.oecd.org/en/publications/oecd-due-diligence-guidance-for-responsible-supply-chains-in-the-garment-and-footwear-sector_9789264290587-en.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S22 Learning to Order for Inventory Systems with Lost Sales and Uncertain Supplies GRADE B

arXiv, Chen, Jiang, Zhang and Zhou. 2023-10-31. https://arxiv.org/abs/2207.04550

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S23 Inventory Management Under Stochastic Demand: A Simulation-Optimization Approach GRADE B

arXiv, Sarit Maitra. 2024-06-27. https://arxiv.org/abs/2406.19425

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S27 INV-44 Cycle Counting Process GRADE A

North Carolina Office of the State Controller. 2023-10-24. https://www.ncosc.gov/training/ncfs-system-training/ncfs-help-documents/inv-44-cycle-counting-process

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c01-S30 Time Release Study Guide: Version 4 GRADE A

World Customs Organization. 2025. https://www.wcoomd.org/en/topics/facilitation/instrument-and-tools/tools/time-release-study.aspx?p=1

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S01 Same source record as c01-S01.

c02-S05 Same source record as c01-S05.

c02-S07 Same source record as c01-S07.

c02-S09 Same source record as c01-S09.

c02-S13 Trade facilitation GRADE A

World Trade Organization. n.d.. https://www.wto.org/english/tratop_e/tradfa_e/tradfa_e.htm

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S15 Same source record as c01-S15.

---

c02-S16 The impact of RFID technologies on inventory accuracy in the apparel retailing: Evidence from the field GRADE B

University of Modena and Reggio Emilia institutional repository. 2015. https://iris.unimo.it/handle/11380/1188598

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S20 Same source record as c01-S20.

---

c02-S22 Same source record as c01-S22.

---

c02-S26 The destruction of returned and unsold textiles in Europe’s circular economy GRADE A

European Environment Agency. 2024. https://www.eea.europa.eu/en/analysis/publications/the-destruction-of-returned-and-unsold-textiles-in-europes-circular-economy

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S27 Same source record as c01-S27.

---

c02-S29 Incoterms 2020 GRADE A

International Chamber of Commerce. 2020. https://iccwbo.org/business-solutions/incoterms-rules/incoterms-2020/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S30 Same source record as c01-S30.

---

c02-S31 EPC/RFID Data Exchange Study: Project Zipper Executive Summary GRADE B

GS1 US and Auburn University RFID Lab. 2018. https://documents.gs1us.org/adobe/assets/deliver/urn:aaid:aem:a8d86484-ffce-498d-8c59-fc1223437d58/Case-Study-Executive-Summary-EPC-RFID-Retail-Supply-Chain-Data-Exchange-Study.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c02-S0201 INV-15 Enter Cycle Count Results GRADE A

North Carolina Office of the State Controller. 2023-07-26. https://www.ncosc.gov/training/ncfs-system-training/ncfs-help-documents/inv-15-enter-cycle-count-results

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S0202 INV-28 Approve Cycle Counts GRADE A

North Carolina Office of the State Controller. 2023-07-27. https://www.ncosc.gov/training/ncfs-system-training/ncfs-help-documents/inv-28-approve-cycle-counts

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c02-S0203 Inventory management VIII: Inventory management and optimization in practice GRADE A

MIT OpenCourseWare. 2006. https://ocw.mit.edu/courses/esd-260j-logistics-systems-fall-2006/resources/lect16/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S01 Same source record as c01-S01.

c03-S04 Inventory management III: Time-varying demand GRADE B

MIT OpenCourseWare. 2006. https://ocw.mit.edu/courses/esd-260j-logistics-systems-fall-2006/resources/lect9/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S05 Same source record as c01-S05.

c03-S06 Same source record as c01-S06.

c03-S07 Same source record as c01-S07.

c03-S08 Same source record as c01-S08.

c03-S09 Same source record as c01-S09.

c03-S17 The Effects of Postponed Manufacturing and Product Planning on Inventory Performance: An Empirical Analysis of Japanese Apparel Firms GRADE B

Japan Marketing Academy, Quarterly Journal of Marketing. 2020-03-31. https://www.jstage.jst.go.jp/article/marketing/39/4/39_2020.015/_article/-char/en

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S20 Same source record as c01-S20.

c03-S22 Same source record as c01-S22.

c03-S25 Due diligence essentials for responsible garment and footwear GRADE A

OECD. 2026-02-26. https://www.oecd.org/en/publications/responsible-business-conduct-spotlights_03a75bf9-en/due-diligence-essentials-for-responsible-garment-and-footwear_c7ae4e4a-en.html

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S27 Same source record as c01-S27.

c03-S0301 Why do zeroes happen? A model-based approach for demand classification GRADE B

arXiv, Svetunkov and Sroginis. 2025-11-13. https://arxiv.org/abs/2504.05894

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c03-S0302 A Simulation-Optimization Technique for Service Level Analysis in Conjunction with Reorder Point Estimation and Lead-Time consideration: A Case Study in Sea Port GRADE B

arXiv, Arani, Abdolmaleki, Maleki, Momenitabar and Liu. 2021-05-28. https://arxiv.org/abs/2106.00767

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S04 Same source record as c03-S04.

c04-S05 Same source record as c01-S05.

c04-S06 Same source record as c01-S06.

c04-S07 Same source record as c01-S07.

c04-S08 Same source record as c01-S08.

c04-S09 Same source record as c01-S09.

c04-S12 Logistics Performance Indicators: LPI Report 2025 GRADE B

World Bank. 2025. https://lpi.worldbank.org/en/home

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c04-S13 Same source record as c02-S13.

c04-S14 Same source record as c01-S14.

c04-S15 Same source record as c01-S15.

c04-S18 Same source record as c01-S18.

c04-S22 Same source record as c01-S22.

c04-S23 Same source record as c01-S23.

c04-S24 The impact of stochastic lead times on the bullwhip effect GRADE B

arXiv, Michna, Nielsen and Nielsen. 2015-03-20. https://arxiv.org/pdf/1411.4289

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S25 Same source record as c03-S25.

---

c04-S27 Same source record as c01-S27.

---

c04-S29 Same source record as c02-S29.

---

c04-S30 Same source record as c01-S30.

---

c04-S31 Same source record as c02-S31.

---

c04-S0401 Agreement on Trade Facilitation GRADE A

World Trade Organization. 2017-02-22. https://www.wto.org/english/docs_e/legal_e/tfa_e.htm

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c04-S0402 NIST/SEMATECH e-Handbook of Statistical Methods: Section 8.2.1.5 Empirical model fitting, distribution free (Kaplan-Meier) approach GRADE A

National Institute of Standards and Technology. 2012-04. https://www.itl.nist.gov/div898/handbook/apr/section2/apr215.htm

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

---

c05-S03 Same source record as c01-S03.

---

c05-S05 Same source record as c01-S05.

---

c05-S07 Same source record as c01-S07.

---

c05-S22 Same source record as c01-S22.

---

c05-S23 Same source record as c01-S23.

---

c05-S24 Same source record as c04-S24.

---

c05-S30 Same source record as c01-S30.

c05-S0501 Property Safety Stock Policy for Correlated Commodities Based on Probability Inequality GRADE B

arXiv, Takashi Shinzato. 2017-01-09. https://arxiv.org/abs/1701.02245

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c05-S0502 Extensions to the Guaranteed Service Model for Industrial Applications of Multi-Echelon Inventory Optimization GRADE B

arXiv, Achkar, Brunaud, Perez, Musa, Mendez and Grossmann. 2023-06-19. https://arxiv.org/abs/2306.10961

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S02 Same source record as c01-S02.

c06-S03 Same source record as c01-S03.

c06-S04 Same source record as c03-S04.

c06-S05 Same source record as c01-S05.

c06-S06 Same source record as c01-S06.

c06-S07 Same source record as c01-S07.

c06-S08 Same source record as c01-S08.

c06-S09 Same source record as c01-S09.

c06-S10 Same source record as c01-S10.

c06-S11 Purchasing practices and low wages in global supply chains: Empirical cases from the garment industry GRADE A

International Labour Organization. 2017-05-15. https://www.ilo.org/publications/purchasing-practices-and-low-wages-global-supply-chains-empirical-cases

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c06-S18 Same source record as c01-S18.

c06-S20 Same source record as c01-S20.

c06-S22 Same source record as c01-S22.

c06-S29 Same source record as c02-S29.

c07-S01 Same source record as c01-S01.

c07-S04 Same source record as c03-S04.

c07-S05 Same source record as c01-S05.

c07-S06 Same source record as c01-S06.

c07-S07 Same source record as c01-S07.

c07-S08 Same source record as c01-S08.

c07-S09 Same source record as c01-S09.

c07-S17 Same source record as c03-S17.

c07-S21 Ban on destruction of unsold clothes and shoes enters into application GRADE A

European Commission, Directorate-General for Environment. 2026-07-17. https://environment.ec.europa.eu/news/ban-destruction-unsold-clothes-and-shoes-enters-application-2026-07-17_en

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c07-S22 Same source record as c01-S22.

c07-S24 Same source record as c04-S24.

c07-S26 Same source record as c02-S26.

c08-S03 Same source record as c01-S03.

c08-S04 Same source record as c03-S04.

c08-S07 Same source record as c01-S07.

c08-S19 Simantha: Simulation for Manufacturing GRADE A

National Institute of Standards and Technology. 2022-01-05. https://www.nist.gov/services-resources/software/simantha-simulation-manufacturing

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c08-S28 **Lean Thinking Part II** GRADE A

MIT OpenCourseWare. 2012. https://ocw.mit.edu/courses/16-660j-introduction-to-lean-six-sigma-methods-january-iap-2012/f044431285fe865dd395b27f5d3c83ce_MIT16_660JIAP12_1-3part2.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c08-S31 Same source record as c02-S31.

c09-S01 Same source record as c01-S01.

c09-S03 Same source record as c01-S03.

c09-S05 Same source record as c01-S05.

c09-S07 Same source record as c01-S07.

c09-S08 Same source record as c01-S08.

c09-S09 Same source record as c01-S09.

c09-S10 Same source record as c01-S10.

c09-S12 Same source record as c04-S12.

c09-S17 Same source record as c03-S17.

c09-S19 Same source record as c08-S19.

c09-S24 Same source record as c04-S24.

c09-S25 Same source record as c03-S25.

c09-S29 Same source record as c02-S29.

c09-S30 Same source record as c01-S30.

c10-S03 Same source record as c01-S03.

c10-S05 Same source record as c01-S05.

c10-S07 Same source record as c01-S07.

c10-S08 Same source record as c01-S08.

c10-S10 Same source record as c01-S10.

c10-S11 Same source record as c06-S11.

c10-S18 Same source record as c01-S18.

c10-S20 Same source record as c01-S20.

c10-S22 Same source record as c01-S22.

c10-S24 Same source record as c04-S24.

c10-S25 Same source record as c03-S25.

c10-S30 Same source record as c01-S30.

c10-S31 Same source record as c02-S31.

c10-S1001 ESD.260J Logistics Systems, Lecture 18: Supply Contracts GRADE A

MIT OpenCourseWare. 2006-11-09. https://ocw.mit.edu/courses/esd-260j-logistics-systems-fall-2006/9382bc01d4a703b-d28bec0f6e2d3e536_lect18.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c10-S1002 INWORK Policy Brief No. 10: Purchasing practices and working conditions in global supply chains: Global Survey results GRADE A

International Labour Organization. 2017-06-09. https://www.ilo.org/sites/default/files/wcmsp5/groups/public/%40ed_protect/%40protrav/%40travail/documents/publication/wcms_556336.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c10-S1003 Conditions of Work and Employment Series No. 86: Purchasing practices and low wages in global supply chains: Empirical cases from the garment industry GRADE A

International Labour Organization. 2017-05-15. https://www.ilo.org/sites/default/files/wcmsp5/groups/public/%40ed_protect/%40protrav/%40travail/documents/publication/wcms_561141.pdf

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c11-S02 Same source record as c01-S02.

c11-S03 Same source record as c01-S03.

c11-S05 Same source record as c01-S05.

c11-S08 Same source record as c01-S08.

c11-S09 Same source record as c01-S09.

c11-S18 Same source record as c01-S18.

c11-S21 Same source record as c07-S21.

c11-S26 Same source record as c02-S26.

c11-S31 Same source record as c02-S31.

c11-S1101 Incoterms 2020 C rules vs D rules: Risk transfer explained GRADE B

ICC Academy. 2026-01-20. https://academy.iccwbo.org/incoterms/article/incoterms-2020-c-or-d-rules/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c12-S03 Same source record as c01-S03.

c12-S05 Same source record as c01-S05.

c12-S07 Same source record as c01-S07.

c12-S22 Same source record as c01-S22.

c12-S23 Same source record as c01-S23.

c12-S24 Same source record as c04-S24.

c12-S27 Same source record as c01-S27.

c12-S31 Same source record as c02-S31.

c13-S05 Same source record as c01-S05.

c13-S09 Same source record as c01-S09.

c13-S10 Same source record as c01-S10.

c13-S19 Same source record as c08-S19.

c13-S20 Same source record as c01-S20.

c13-S22 Same source record as c01-S22.

c13-S23 Same source record as c01-S23.

c13-S24 Same source record as c04-S24.

c13-S25 Same source record as c03-S25.

c13-S27 Same source record as c01-S27.

c14-S01 Same source record as c01-S01.

c14-S02 Same source record as c01-S02.

c14-S03 Same source record as c01-S03.

c14-S05 Same source record as c01-S05.

c14-S06 Same source record as c01-S06.

c14-S07 Same source record as c01-S07.

c14-S08 Same source record as c01-S08.

c14-S09 Same source record as c01-S09.

c14-S10 Same source record as c01-S10.

c14-S11 Same source record as c06-S11.

c14-S12 Same source record as c04-S12.

c14-S13 Same source record as c02-S13.

c14-S15 Same source record as c01-S15.

c14-S16 Same source record as c02-S16.

c14-S17 Same source record as c03-S17.

c14-S18 Same source record as c01-S18.

c14-S19 Same source record as c08-S19.

c14-S20 Same source record as c01-S20.

c14-S22 Same source record as c01-S22.

c14-S23 Same source record as c01-S23.

c14-S24 Same source record as c04-S24.

c14-S25 Same source record as c03-S25.

c14-S27 Same source record as c01-S27.

c14-S28 Same source record as c08-S28.

c14-S29 Same source record as c02-S29.

c14-S30 Same source record as c01-S30.

c14-S31 Same source record as c02-S31.

c14-S1401 Inventory management V: Backorder costs, implied costs, extensions, A and C items GRADE B

MIT OpenCourseWare. 2006. https://ocw.mit.edu/courses/esd-260j-logistics-systems-fall-2006/resources/lect12/

Opened source record.

Accessed 2026-10-30. Quote retained in the machine-readable source record.

c14-S1402 Same source record as c02-S0203.

Stable definitions make comparisons useful.

ABC analysis

Classification by cumulative annual usage value, with locally chosen breakpoints and explicit criticality overrides.

Available-to-promise

Quantity that can be committed to a dated customer requirement under defined stock, receipt, allocation and capacity rules.

Backorder

Unfilled demand retained as an outstanding obligation to deliver later.

Bill of materials

Versioned component quantities and specifications required to produce an item.

Bullwhip effect

Amplification of order variability upstream relative to downstream demand under a consistently defined comparison.

Call-off

Binding release of quantity or timing against a framework agreement, according to its stated terms.

Carrying cost

Decision-model cost of holding an additional unit for a stated time, distinct from accounting stock value.

Censored demand

Demand only partly observable because availability or another constraint limits recorded sales.

Complete-kit gate

Proposed production-release check that all required accepted components and specifications are ready.

Consignment stock

Stock physically held by a party while ownership remains with another under agreed contractual terms.

Critical fractile

Ratio of underage cost to the sum of underage and overage cost in the basic single-period inventory model.

Cycle count

A periodic physical verification of selected stock records with investigation and approval of discrepancies.

Cycle service level

Probability, or observed share, of replenishment cycles with no stockout under a defined convention.

Cycle stock

Inventory associated with replenishment lot sizes rather than uncertainty protection.

Decoupling point

Stage separating forecast-driven commitment from activities triggered by a confirmed customer order.

Dependent demand

Component demand derived from requirements for parent items through their bills of materials.

Distribution requirements planning

Time-phased calculation of replenishment requirements across distribution locations.

Economic order quantity

Lot size minimising ordering and holding costs under the basic model's restrictive assumptions.

Echelon stock

Stock measured across a defined downstream network boundary, with consistent treatment of pipeline and backorders.

Fill rate

Proportion of requested unit demand filled immediately from available stock under a stated boundary.

Firm receipt

Expected inbound quantity whose commitment status and timing meet a defined planning acceptance rule.

Forecast residual

Observed demand minus the forecast available at the decision date, measured on a consistent time basis.

Goods in transit

Materials moving between physical locations, with ownership and cost treatment determined separately.

Inventory position

Defined planning quantity combining on-hand stock, outstanding supply and unfilled commitments without double counting.

Item-location policy

Replenishment rule, parameter set, constraints and approval authority for one item at one location.

Landed cost

Cost assigned to bringing an input to a stated location and condition, with dated currency and contractual assumptions.

Lead-time demand

Demand occurring during the interval from an effective replenishment decision to usable receipt.

Little's Law

Stable-system relationship linking average inventory to average throughput multiplied by average time within a consistent boundary.

Lost sale

Demand that is not fulfilled and does not remain as a backorder.

Material requirements planning

Time-phased netting and bill-of-material explosion of dependent requirements and available supply.

Minimum order quantity

Smallest order quantity a supplier will accept under specified terms.

Net realisable value

Estimated selling price in ordinary business less estimated completion and necessary selling costs, as described by IAS 2.

Newsvendor model

Single-period inventory model balancing the incremental costs of ordering too little and too much.

On-time in full

Order-level delivery metric requiring the agreed complete quantity by the agreed date under a defined acceptance rule.

Order-up-to level

Target inventory position used to calculate replenishment at a review event.

Overage cost

Incremental economic loss from a unit ordered beyond demand in a specified inventory model.

Periodic review

Inventory policy that assesses replenishment at scheduled intervals.

Postponement

Delay of product differentiation or another irreversible commitment until additional demand information is available.

Protection period

Horizon whose uncertain demand a stock policy must cover, including a review gap where applicable.

Quarantine stock

Stock unavailable for use pending inspection, investigation or authorised release.

Reorder point

Inventory-position threshold that triggers replenishment under a continuous-review policy.

Right censoring

Observation whose eventual completion time exceeds the end of the measurement window.

Risk pooling

Combining compatible requirements or resources to reduce uncertainty exposure, conditional on dependence and service constraints.

Safety stock

Buffer intended to protect a defined service outcome against uncertainty beyond expected demand.

Shadow mode

Operation of a candidate recommendation process without executing its proposed orders.

Stock-keeping unit

Distinct inventory item tracked under an agreed specification, such as a style-colour-size variant.

Underage cost

Incremental economic loss from failing to supply a unit of demand in a specified inventory model.

Usable receipt

Arrival that has met the physical, specification, documentation and quality-release conditions required for consumption.

Work in progress

Items that have entered but not completed the defined production process.

Yield

Accepted output divided by stated input on a consistent physical and process boundary.

Follow the terms back to their context.

Page references point to occurrences in chapter prose. When a term does not occur in a chapter, its glossary definition is indexed.

ABC analysis 49

Available-to-promise 45 46 207

Backorder 81 90 108

Bill of materials 15 16 17 29 39 50 100 122 162 209

Bullwhip effect 154 212

Call-off 18 59 101 104 105 112 113 117 156 157 158 159 162 212

Carrying cost 103 153

Censored demand 190

Complete-kit gate 122 132 133 213 216

Consignment stock 104

Critical fractile 58 61 108 109 110 112 159

Cycle count 23 39 40 42 43 46 210

Cycle service level 56 214

Cycle stock 86 87 94 99 102 138 142 144

Decoupling point 49 53 54 55 59 117

Dependent demand 17 60 122 124

Distribution requirements planning 136 211 212

Echelon stock 143 148 212

Economic order quantity 94

Fill rate 19 56 57 78 80 82 87 89 90 111 116 138 189 201 211 214 215

Firm receipt 81 124

Forecast residual 81 138

Goods in transit 18 35 75 96 166 201

Inventory position 80 81 83 85 130 143 166 179 190 209

Item-location policy 22 25 62 179

Landed cost 61

Lead-time demand 56 74 81 83 84 87 89 153 189 190 216

Little's Law 240

Lost sale 22 111

Material requirements planning 15 30 36 53 80 99 102 122 124 209

Minimum order quantity 101

Net realisable value 19 35 50 58 96 97 110 118 147 168 175 209 214

Newsvendor model 58

On-time in full 56

Order-up-to level 80 85 86 87

Overage cost 58 108 110 159

Periodic review 60 85 86 90 179

Postponement 22 54 116 117 136 141 206 211

Protection period 25 137 138 140

Quarantine stock 210 218

Reorder point 20 29 52 56 58 61 66 80 81 83 84 89 90 95 212 214

Right censoring 69 151

Risk pooling 206 211

Safety stock 62 74 77 80 82 83 84 86 87 96 131 140 142 150 153 163 190 216 218

Shadow mode 199 206 210 211

Stock-keeping unit 29 206

Underage cost 58 108 159

Usable receipt 81 84 137 142 144 146 150 151 157 163 167 180 184

Work in progress 18 21 96 125 129 130 131 209 214

Yield 22 32 103 124 151 152 153 162 163

A reproducible edition keeps its limits visible.

Production is the warp. Measurement is the weft.

Green Future Tech prepared this white paper from published sources opened during research, each recorded with its URL, publisher, date and a quoted sentence in the evidence register, and from GFT's Cameroon textile programme evidence register (R entries, transcribed from the programme's data annex and flagged where refreshed in 2026). Derived figures show their inputs; illustrative models are labelled with their basis. No figure is presented as audited, and nothing here is legal, tax or investment advice.

Source treatment preserves uncertainty.

Inherited records reproduce the supplied finding and attribution only where cited. Their source-text status is preserved. Original publication dates and exact source URLs are incomplete in that register. These records require independent validation before external reliance. New chapter-local sources carry a publication date, complete URL, publisher, exact supporting quotation and access date.

The exports share a content model.

The PDF and web edition use the same parsed chapter blocks and SVG exhibits. The build records input hashes, artifact hashes, layout measurements and elapsed time. Cross-references, source notes and the term index are generated from identifiers rather than manually typed page numbers.

IMPRINT

Issued by Green Future Tech Limited, Hong Kong SAR.
Engineering the sustainable production of tomorrow.

GFT Research. Edition 2026-10-30. Document wp12. Locally generated proof. No named author, office, certification or client endorsement is implied.

GREEN FUTURE TECH

LIMITED / HONG KONG SAR

The factory that makes the garment should make the evidence.

Engineering the sustainable production of tomorrow.