An Approach to Effective Supply Chain: A Conceptual Framework and Research Agenda

A Digital, IoT and AI-Driven Supply Chain Framework for FMCG, Food and Beverage Industries under Industry 5.0

Santanu Mukherjee

B.Tech (Mechanical), MBA (Operations Management)

How to Cite this Article:

Mukherjee, S. (2026). An Approach to Effective Supply Chain: A Conceptual Framework and Research Agenda. International Journal of Creative and Open Research in Engineering and Management, <i>02</i>(8), 1-9. https://doi.org/10.55041/ijcope.v2i8.012

https://doi.org/10.55041/ijcope.v2i8.012

Abstract

The accuracy and the limits of ERP, WMS, and TMS architectures consistently fail when confronted with the physical findings of a modern supply chain system, especially in food, FMCG & FMCD sectors, where complexity of the Supply chain system is high and product shelf presence is the top sought factor. Latency sits at the centre of the difficulty. Digital system signals and actual shop-floor movement drift apart, whole segments of the chain go unobserved for hours at a stretch, and the systems often fail to track the mapping of a semi-finished product movement through different stages. The framework addresses the operational reality of high-velocity manufacturing: OEE stabilization, OTIF performance and shelf availability are inseparable. When one breaks, all three fail. This framework treats these elements as an interdependent system rather than isolated KPIs. Drawing on practical floor-level experience, it introduces the IPC Network Model—incorporating Industry 4.0 principles (IPC-N4.0)—to construct a highly stable, visible, and mappable supply chain management system.

Unlike traditional software-first approaches, the IPC framework is grounded in physical execution. It addresses the systemic lag in traditional software (though their importance in Industry is quite high) by integrating physical Lean principles, such as dynamic Kanban implementation, with decentralized data structures like the Vendor Intelligence Index (VII) and the Product Digital Passport (PDP). The result is a pragmatic, Industry 4.0-aligned framework designed to bridge the gap between static enterprise systems and the complex, often chaotic reality of modern logistics.

The framework introduces a set of new operating instruments — Inventory Health Score, Inventory Behaviour Index, Product Digital Passport, Product Journey Map, Vendor Intelligence Index, Customer Fulfilment Index, Retail Shelf Availability Index and others — together with a low-cost, blockchain-free chain-of-custody architecture suitable for FMCG cost structures. The paper illustrates the framework with publicly documented industry disruptions (KFC UK 2018, Abbott 2022, Maggi India 2015, Huy Fong Sriracha, COVID-19 FMCG disruption and others), and provides simulated but transparently constructed financial impact estimates. The paper is offered as a conceptual contribution; pilot validation in FMCG and F&B environments is explicitly recommended as the next step.

Keywords: *Supply chain management; IPC Network Model; FMCG; Food and Beverage; Industry 5.0; digital twin; IoT; RFID; QR code; vendor intelligence; inventory health; product traceability; human-centred decision making.*

1. Introduction

1.1 Overview of the Argument

I have spent a little over two decades inside FMCG and food manufacturing supply chains, first on the plant floor and later across procurement, warehousing and distribution. In that time the tooling moved from paper ledgers and spreadsheets to ERP suites and multi-tier global networks. The transactional backbone, in most large enterprises, now works well enough. However, on the shop floor and in the boardroom, the fundamental operational question remains unanswered: do we know, right now, exactly where our stock is, the condition of our products, and whether our customers are actually getting what they need before a failure hits the P&L?

The IPC (Inventory–Product–Customer) Network Model is built to answer this by design, bypassing the lag of standard exception reporting. This framework is anchored in operational reality:

- **Three anchors carry the argument. Supply chain performance, stripped of its vocabulary, comes down to inventory (what is held), product (what is moving, and in what condition) and the customer (whose expectations are met or missed). Vendors, warehouses and fleets are infrastructure in service of those three; they are not ends in themselves. The model I set out here is built on that ordering, and I return to it whenever a design choice has to be settled.**

- **Practical Observability: The tools required for continuous visibility—RFID, IoT gateways, predictive analytics, and human-in-the-loop AI—are finally mature and cost-effective enough for margin-tight FMCG environments. The framework does not require heavy, over-engineered solutions like blockchain to achieve this.**

- **Human-Centric Execution:** Resilient, Industry 5.0 operations require an intelligence layer that surfaces actionable decisions to supply chain leaders, rather than burying them under automated black-box processes.

This document outlines the reference architecture for the IPC Network Model, introducing practical operating instruments for inventory and vendor management. It details a low-cost, event-based chain of custody and introduces the Vendor Intelligence Index—a scoring framework that tracks performance far beyond baseline price and OTIF metrics. It also provides reference screens for a digital IPC platform and a structured comparison with traditional SCM using simulated data.

The framework has been stress-tested against existing FMCG operational constraints and addressed the gap between shop floor reality and a software capability, 2it is an architectural blueprint designed specifically for live pilot testing.

The paper contributes: (i) a formal definition and reference architecture for the IPC Network Model;

(ii) a set of new operating instruments across inventory, product, customer and vendor dimensions;

(iii) a low-cost, event-based chain-of-custody design that does not require blockchain; (iv) a scoring framework — the Vendor Intelligence Index — that goes beyond price and OTIF; (v) reference screens for a digital IPC platform; and (vi) a structured comparison with traditional SCM supported by reference case simulated data. The paper does not claim empirical validation. It is written to be pilot-tested.

Expected benefits, based on scenario analysis presented later in the paper, include material reductions in inventory holding, waste and stock-outs, meaningful improvements in vendor governance and customer fulfilment, and a shortened cash-to-cash cycle. These are presented as plausible directional outcomes rather than guaranteed results, and every assumption is stated in the open.

1.2 From Traditional Supply Chain Management to the IPC Layer

1.2.1 A Case that Continues to alert the industry

In February 2018, KFC UK shuttered roughly two-thirds of its outlets. The surface-level cause is well known: transferring logistics to a new provider whose single Midlands distribution centre simply collapsed under the volume of daily chilled deliveries. Within days, the core raw material—chicken—vanished from stores. The result was fullpage newspaper apologies and massive revenue haemorrhage.

The apologies in the national press were the visible part. What sat underneath was an execution failure in a stretch of the chain nobody was watching. Visibility effectively died at the ERP purchase order. The system recorded the order but had zero real-time grasp on transit status, failing routes, or which specific restaurants would hit zero inventory first. Well-configured traditional SCM systems are excellent at logging history. They are practically useless at telling an operations director where the network will break in the next four to twenty-four hours.

The pattern recurs across FMCG. During the 2015 Maggi episode in India, the regulatory action itself concerned alleged lead content and labelling (FSSAI, 2015); what the public record also shows is that the withdrawal was executed nationally rather than by batch or region. Whether granular batch isolation would have narrowed the recall is, strictly, an open question, and I treat it as a design requirement to be tested rather than as an established finding. The 2022 Abbott infant formula recall highlighted the fragility of over-relying on a single facility without plant-to-shelf visibility for rapid, targeted containment. Huy Fong's Sriracha shortage proved how single-source raw material bottlenecks, invisible until it is too late, can paralyze a global brand. COVID-19 then did to demand planning what no consultant's slide deck had managed: it showed, across every FMCG operator at once, that historical velocity is a poor guide to tomorrow's demand.

These incidents share a common structure. The transactional system worked. The operating system — the layer that should have translated data into anticipation — did not exist. That absence, rather than any deficiency in the transactional record, is what this study treats as the structural flaw.

1.2.2 The Operational Shift to Industry 5.0

Supply chain architecture has continuously chased complexity. The field spent decades implementing MRPs for basic inventory math, integrating siloed departments into monolithic ERP suites, and eventually establishing S&OP disciplines. Recently, Industry 4.0 brought IoT, cloud computing, and virtual replica of the nodes to the shop floor. Yet, despite this massive tech stack, the operational latency remains.

Industry 5.0 is not about rejecting Industry 4.0; it is about course-correcting it. It shifts the focus from blind automation to resilient, human-centered execution. The IPC framework aligns directly with this operational mandate. It assumes the presence of rich telemetry but channels that data into actionable insights for the supply chain leaders who actually have to make the decisions.

Industry 5.0, as articulated by the European Commission and increasingly by industrial practitioners, is not a rejection of Industry 4.0 but a re-centring of it: technology in service of people, sustainability and resilience, rather than automation for its own sake. IPC is proposed as a framework that fits inside this era. It presumes rich telemetry, but it exists to present that telemetry to human decision-makers in a form they can act on.

1.2.3 Positioning of the IPC Network

The IPC is not a replacement for the ERP, WMS, TMS, or MES. It acts as an overarching decision layer designed to continuously answer three critical operational questions:

- **Inventory:** Is this stock healthy, aging, at risk, or already dead capital—and what is the immediate corrective action?
- **Product:** Where is it physically, what is its actual condition in transit, and does the digital record match the physical reality on the floor?
- **Customer:** Are we hitting our delivery promise, and if we are failing the shelf or the distributor, where is the bottleneck and how do we bypass it fastest?

The remainder of this described model details this framework. This is not a software specification document. It is an operational blueprint for running a supply chain where transactional systems are treated as baseline infrastructure, and the true competitive advantage is built in the diagnostic layer operating actively above them.

1.3 Why Traditional Supply Chain Management Needs Another Layer

Traditional SCM, as implemented in most large FMCG and F&B organisations, is built around a small number of powerful systems: an ERP as the transactional spine, a WMS at each large distribution centre, a TMS for outbound logistics, an MES on the shop floor, and a demand planning or S&OP tool. Each of these systems, taken individually, is mature. The gaps discussed below are not defects of any one system; they are properties of the way the systems are stitched together and of what none of them, by design, does.

1.3.1 Inventory Visibility

ERP stock is what the ledger believes to be true. Physical stock is what the shelf actually holds. In practice these two numbers diverge continuously — because of scan errors, put-away mistakes, damage, silent returns, expiry, sampling, in-transit ambiguity and simple human error. Organisations manage the divergence with cycle counts and periodic wall-to-wall inventories. Between counts, decisions are made on assumed truth. IPC treats inventory visibility as a continuous signal rather than a periodic reconciliation.

1.3.2 Demand Distortion

The bullwhip effect is a textbook phenomenon; it is also, thirty years after Lee, Padmanabhan and Whang formalised it, a daily reality on every FMCG demand review. Distributors buffer against the manufacturer, retailers buffer against the distributor, and the manufacturer plans against a signal that is already three amplifications away from real consumer demand. IPC does not eliminate distortion, but it insists that primary sales, secondary sales and retail off-take are held together in one view.

1.3.3 Vendor Evaluation

Most vendor scorecards, in my observation, still weight price and on-time delivery disproportionately. Quality, statutory compliance, ESG performance, innovation capability and behavioural attributes such as responsiveness are often assessed qualitatively, if at all. The result is that vendors who look cheap on the purchase order are frequently expensive on the P&L. Section 4.5 introduces the Vendor Intelligence Index as an explicit remedy.

1.3.4 Warehouse Blind Spots

Even the best WMS treats a warehouse largely as a set of locations with quantities. It rarely knows, without additional instrumentation, that a particular slot is congested, that a particular SKU family is being repeatedly touched to reach another, that certain dispatch windows are systematically breached, or that a certain aisle experiences temperature excursions in summer. IPC fills these blind spots with IoT signals and analytics.

1.3.5 Product Movement Opacity

From the moment a case leaves the mother warehouse, most FMCG companies enter a period of near-blindness that ends only when a distributor uploads a secondary sales file or a retailer scans at point of sale. In the intervening days, product may be at a transhipment hub, on a truck, at a distributor godown, or informally rerouted. IPC introduces an event-based chain of custody, described in Section 4.6, that closes this gap without requiring blockchain.

1.3.6 Customer Disconnect

Most supply chain organisations meet the customer only through the complaint log and the order fill rate. Neither is sufficient. Complaints are lagging indicators of failures; fill rate is an internal metric that does not tell whether the consumer found the right product on the right shelf. IPC treats the customer, the retail shelf and the consumer as first-class entities within the operating model.

1.3.7 Information Latency and Cross-Functional Silos

Even where data exists, it is often locked behind report refresh cycles and departmental boundaries. Sales sees one dashboard, operations sees another, quality sees a third, and the composite picture is assembled only when there is already a problem. IPC assumes a single integrated view built on top of the existing systems.

1.3.8 The Case for an Operating layer

The proposition of this framework is that these gaps are not solved by replacing the existing systems, which represent enormous embedded investment and generally work. They are solved by adding a layer above them whose sole purpose is to produce decisions. That layer is the IPC Network Model.

2. Literature Review

2.1 Scope and approach of the review

The review that follows is selective rather than exhaustive. Four bodies of work bear directly on the argument: the visibility and traceability literature, the Industry 4.0 to Industry 5.0 debate, the work on distributed ledgers in food and consumer-goods chains, and the supplier-evaluation literature. Where a source is practitioner rather than peer-reviewed, I say so, and I treat its claims as evidence of industry belief rather than of measured effect. The paragraphs below take a position on these works rather than summarising them.

2.2 Visibility, traceability and the persistence of latency

Supply chain visibility has been studied for at least three decades, and yet the definitions remain loose. Barratt and Oke (2007) treat visibility as the extent to which actors have access to timely, accurate information they consider useful; Swift, Guide and Muthulingam (2019) show empirically that firms with better internal information linkage carry less inventory. Both streams share an assumption that I question here, namely that information which exists somewhere in the enterprise is information the operator can act on. In FMCG practice the two are far apart. A batch record sitting in a quality module and a temperature excursion recorded by a third-party carrier are, in principle, available; in the four to twenty-four hour window in which a decision must be taken they are effectively unavailable.

Traceability research is more advanced in food than in most sectors, largely because regulation forced it. Aung and Chang (2014) set out the technical requirements for food traceability, and Olsen and Borit (2013) show how inconsistently the term is defined across studies. What the literature has not resolved is the granularity problem: firms achieve traceability at the level of the production lot but lose it at the level of the case once the case enters a distributor's godown. Recall studies in the food safety literature repeatedly describe this loss of granularity as the reason withdrawals become national rather than local. I take that finding as the empirical anchor for the product layer of the framework proposed here, and I have avoided restating the popular reading of individual recalls as settled fact where the published record does not support it.

2.3 From Industry 4.0 to Industry 5.0

The Industry 4.0 literature is now large and, in places, uncritical. Kagermann and colleagues framed the original agenda around cyber-physical systems; subsequent reviews, including Frank, Dalenogare and Ayala (2019), report that adoption in process and consumer-goods manufacturing is patchy and concentrated in sensing rather than in decision-making. The pattern is consistent: firms instrument the plant, accumulate telemetry, and then discover that no one owns the decision the telemetry implies.

Industry 5.0, as set out in the European Commission's 2021 formulation and elaborated by Xu et al. (2021), reacts to precisely this. Its three commitments are human-centricity, resilience and sustainability. The criticism levelled at it, fairly in my view, is that it is a policy statement rather than an operating theory: it says what industry should value without saying how a plant manager should behave differently on a Tuesday morning. Breque, De Nul and Petridis (2021) do not offer an operating model, and the subsequent literature has largely stayed at the level of principle. That gap is where this paper positions itself. The framework developed in Section 4 is an attempt to translate the human-centric commitment into named instruments, defined owners and a decision cadence.

2.4 Distributed ledgers: promise, cost and the question of the first mile

Blockchain attracted disproportionate attention in food supply chains between roughly 2017 and 2021. Saberi et al. (2019) provide the most cited treatment, and Kamble, Gunasekaran and Sharma (2020) survey adoption barriers. The reported pilots — Walmart's leafy greens work with IBM Food Trust being the best known — demonstrate that trace time can fall dramatically once data is on a shared ledger.

The unresolved objection is the oracle problem, and it is fatal for the way the technology is usually sold. A ledger guarantees that a record, once written, is not altered without consensus. It guarantees nothing about whether the record was true when it was written. In an FMCG chain the weak link is almost never the ledger; it is the loading-bay scan that did not happen, the pallet re-labelled in a hurry, the temperature logger left in the cab rather than the trailer. Wang, Han and Beynon-Davies (2019) note this in passing; most of the adoption literature does not. There is also a plain cost argument. Consortium governance, node operation and partner onboarding impose a running cost that a category earning single-digit margins on a case of biscuits will not carry. My reading of the evidence is that the marginal benefit of immutability over a well-governed, cryptographically hashed and independently audited event store is small, while the marginal cost is large. Section 4.6 develops that position into a design choice rather than leaving it as an opinion.

2.5 Supplier evaluation beyond price and OTIF

The supplier-selection literature is methodologically rich and practically under-used. Dickson's early enumeration of criteria, Weber, Current and Benton's review, and the large body of AHP, TOPSIS and DEA applications summarised by Chai, Liu and Ngai (2013) establish that multi-criteria evaluation is well understood. Sustainability criteria have been added steadily since Seuring and Müller (2008). What the field has not produced is an instrument that a category buyer will actually maintain every month. The methods that are analytically strongest tend to demand pairwise comparisons and expert panels that no procurement team sustains past the first quarter.

This is a design failure rather than a theoretical one, and it explains why practitioner scorecards regress to price and on-time-in-full. The Vendor Intelligence Index described in Section 4.5 is therefore built for maintainability first: additive weighting, transparent sub-scores, monthly refresh from data that procurement, quality and logistics already generate. I accept that an additive model is less defensible than an outranking method when criteria interact. I argue that a mediocre instrument that is updated is worth more than an elegant one that is abandoned, and I return to this trade-off in Section 7.

2.6 The gap this study addresses

Three conclusions follow from the review. First, visibility research has documented the value of information but has under-theorised the latency between information and action. Second, Industry 5.0 supplies a normative direction without an operating model. Third, the technology debate in food chains has been distracted by immutability at the expense of first-mile data capture and affordability. Taken together, these leave a gap for a middle-range framework: not another enterprise system, and not a statement of values, but a defined operating layer with named instruments, explicit formulae, and an implementation and validation path. That is the contribution this paper attempts.

3. Methodology

3.1 Research design

This is a conceptual, theory-building study. It does not test a hypothesis against a dataset; it constructs a framework and specifies how that framework could be tested. Meredith's (1993) defence of conceptual modelling in operations management and Jaakkola's (2020) typology of conceptual papers describe the genre precisely: the contribution is a model, and the standard of evaluation is internal coherence, grounding in prior work, and testability rather than statistical significance. I state this plainly because the alternative — presenting a framework as though it had been validated — is the more common and the less honest option.

3.2 Position of the researcher and the observational base

The framework is best described as a conceptual framework derived from participant observation and industry case studies. Over roughly two decades in mechanical engineering and operations management in FMCG and food manufacturing, I have worked on plant operations, inventory control, vendor development and distribution in environments running mature ERP and WMS installations. The recurring failures described in Section 1 are ones I have watched, argued about in review meetings and occasionally caused. That proximity is the study's principal source of insight and, simultaneously, its principal threat to validity. Participant observation carries a known risk that the observer's habits of mind are reproduced as findings. I have tried to control for this in three ways: by grounding each proposed instrument in published literature or a publicly documented incident rather than in recollection alone; by stating the assumptions behind every quantitative illustration; and by writing the validation plan in Section 7.3 so that a sceptical reader can attempt to falsify the claims.

No proprietary data from any employer has been used. Where operating figures appear, they are either drawn from published industry benchmarks cited in Appendix D.2 or constructed as reference cases, and they are labelled as such.

3.3 Case selection

Five publicly documented disruptions anchor the argument: the KFC UK distribution failure of 2018, the Maggi withdrawal in India in 2015, the Abbott infant-formula recall of 2022, the Huy Fong Sriracha raw-material shortage, and the COVID-19 demand shock across FMCG. They were selected on three criteria: the events are documented in the public record rather than in internal reports; they span different failure modes (logistics transition, regulatory and traceability, single-site concentration, upstream agricultural supply, systemic demand shock); and each occurred in a firm with a functioning enterprise system, which is what makes them relevant. These are illustrative rather than analytically sampled cases; I use them to motivate design requirements, not to establish frequency or causation, and I have been careful to describe what the public record actually establishes about each.

3.4 Framework construction

Construction proceeded in four passes. The first derived design requirements from the review in Section 2 and from the case set. The second organised the requirements around the three entities that appeared in every failure — inventory, product and customer — and treated vendors, warehouses, fleets and distributors as infrastructure serving those entities rather than as peers of them. The third defined instruments for each entity, with the constraint that any score must be computable from data that a mid-sized FMCG business already holds or can capture with commodity sensing, and must be explicable to a shift manager in one sentence. The fourth pass added the governance, planning and validation apparatus: the MAPE planning loop, the extended Kanban arrangement, the roadmap and the pilot design.

Weight sets, thresholds and coefficients in the instruments are proposals, not estimates. They are starting values intended to be recalibrated against local data during the pilot described in Section 7.3.

3.5 Basis and limits of the quantitative illustrations

All numerical material in Section 6 and Appendix A is scenario analysis built on stated assumptions and published benchmark ranges. It is not measured outcome data, and no claim of empirical validation is made anywhere in this paper. The purpose of the scenarios is to establish whether the framework is worth testing — whether the order of magnitude of plausible benefit justifies the implementation cost — and to expose the assumptions to challenge. Section 6.5 converts the scenarios into formal hypotheses so that a pilot can confirm or refute them rather than simply illustrate them.

4. The IPC Conceptual Framework

4.1 Definition and Reference Architecture

4.1.1 Definition

The IPC Network Model serves as an operating framework and reference architecture where Inventory, Product, and Customer are continuously observable, digitally represented, and analytically linked. Instead of ripping and replacing, it acts as an intelligence layer that extracts data from existing ERP, WMS, TMS, and MES systems, translating that telemetry into immediate, decision-ready insights for the operators actually running the floor.

4.1.2 Guiding Principles

- **Integrate instead of ripping and replacing: IPC bolts onto your existing transactional stack; it is not designed to trigger a high-risk system migration or compete with the ERP.**
- **Real-time signals over batch updates: Relying on periodic reconciliation leaves operations blind. The design prioritises continuous, live signals wherever physically possible.**
- **Establish physical identity first: Analytics are useless if they do not match the reality on the floor. This study argues that analytics cannot be trusted until the physical identity of the item, batch, pallet, or vehicle is absolutely unambiguous.**
- **Decisions stay with people: Aligning with Industry 5.0, the IPC surfaces actionable decisions to people. Unattended automation is reserved for low-risk, reversible tasks.**
- **Deploy practical, affordable technology: The framework relies on proven, cost-effective tools like QR, RFID, barcodes, IoT gateways, and cloud databases. Exotic, expensive technology is unnecessary for margin-tight FMCG environments.**
- **Keep the math explainable:** If a shift manager or non-technical operator cannot understand the scoring formulas or prediction models, the tools will be ignored. Transparency is mandatory.

4.1.3 Reference Architecture

1. **Layer 1 — Physical:** The actual items, pallets, cases, vehicles, warehouses, and retail shelves. Each carries a clear identity via QR, RFID, or barcode.
2. **Layer 2 — Sensing:** The hardware capturing the floor reality—handheld scanners, fixed RFID readers, IoT gateways, temperature and humidity sensors, vehicle GPS units, and mobile devices used by distributor sales representatives.
3. **Layer 3 — Systems:** The legacy backbone. IPC integrates with existing ERP, WMS, TMS, MES, and CRM systems through standard APIs, explicitly avoiding costly replacements.
4. **Layer 4 — IPC Intelligence:** The analytical engine handling identity resolution, event stream processing, and scoring (health, behaviour, fulfilment, vendor). This layer houses the predictive models (expiry, complaint, delay), digital twins of critical nodes, and the rule engine generating active alerts.
5. **Layer 5 — Decision:** The interface layer pushing dashboards, alerts, mobile screens, and structured decision briefs directly to operations, sales, quality, procurement, and executive leadership.

Figure1— Five-layer reference architecture of the IPC Network Model (conceptual).

4.1.4 Data Model

The IPC data model revolves around three master entities: the Inventory Position, the Product Instance, and the Customer Node. These are tied together by an Event Stream capturing every observable physical transition. Table 1 outlines this simplified structure

Entity	Key attributes	Sources
Inventory Position	SKU, batch, location, quantity, condition, age, expiry, health score	ERP, WMS, IoT sensors
Product Instance	Serial or batch ID, digital passport ID, current location, temperature history, ownership	QR/RFID scans, IoT gateways
Customer Node	Distributor, retailer, shelf, consumer segment, fulfilment index, complaint history	TMS, CRM, secondary sales feeds
Event	Timestamp, location, actor, action, entity references	Scanners, mobile apps, integrations
Vendor Node	Vendor ID, category, VII score, contract terms, ESG profile	Procurement, quality, ESG audits

Table 1 — Simplified entity model for the IPC Network. Source: constructed for this study from floor-level operating practice.

4.1.5 What IPC Is Not

To be completely clear about the framework's scope: IPC is not an ERP replacement, nor is it a generic control tower product. It does not operate on the assumption that AI can substitute for hard-earned supply chain judgement on the floor, and the design deliberately sets aside heavy, blockchain-based provenance tracking. Finally, the financial analysis presented in Section 6 represents directional scenario modelling—not a guaranteed outcome or a promise of specific percentage benefits. And it is not a pitch for idealized software performance. The baseline numbers outlined in Section 6 reflect the actual operational grind of pushing OTIF past legacy plateaus in high-complexity food manufacturing, serving as practical targets rather than guaranteed Operational forecasts.

4.2 Inventory Intelligence: New Instruments

Traditional inventory management, as taught in operations textbooks, offers classical instruments: ABC and XYZ classification, EOQ, safety stock formulae, cycle counting policies and ageing reports. These remain useful. They are not, however, sufficient in an environment where SKU counts routinely exceed one thousand, where shelf life may be measured in weeks, and where a single expiry event can consume the margin of a profitable quarter. What the classical set lacks is not rigour but refresh rate: it describes a position monthly, when the position changes hourly. The instruments defined below are computed continuously from transactional and sensor data already present in most FMCG estates.

4.2.1 Inventory Health Score (IHS)

The Inventory Health Score is a composite indicator, computed per SKU per location per day, that summarises the risk profile of a given inventory position. It is defined conceptually as:

$$IHS = w1 \cdot Freshness + w2 \cdot Coverage + w3 \cdot Turnover + w4 \cdot Damage\text{-}Free + w5 \cdot Location\ Fit$$

Freshness expresses remaining shelf life as a proportion of total shelf life. Coverage expresses days-of-cover relative to the acceptable band for the SKU. Turnover expresses trailing velocity against the SKU's norm. Damage-Free expresses the proportion of stock in saleable condition. Location Fit expresses whether the stock sits in the correct temperature and slot category. Each term is normalised to a zero-to-one scale before weighting, so that the composite stays interpretable and the weights sum to unity.

4.2.1.1 Setting the weights

The weights are the part of this instrument that a reviewer should press hardest, because they carry the judgement. I do not treat w_1 to w_5 as universal constants. They are category parameters, set by the operations leadership of the business unit and re-examined at least annually, and the framework is transparent about them precisely so that they can be argued over in a room rather than buried in a model.

Three rules govern how they are set. First, weight follows the dominant failure mode of the category: whichever variable most often destroys value in that category takes the largest share. Second, weights sum to one, so that any increase is an explicit trade-off rather than an addition. Third, no single weight is allowed to exceed 0.50, because a composite dominated by one term is not a composite; if freshness is genuinely the only thing that matters, the business should manage freshness directly and drop the index.

The contrast between two products makes the logic concrete. For pasteurised milk at seven days of shelf life, remaining life is the failure mode: a plausible setting is $w_1 = 0.45$ on Freshness, 0.20 on Coverage, 0.15 on Turnover, 0.10 on Damage-Free and 0.10 on Location Fit, with Location Fit carrying real weight because a cold-chain excursion destroys the batch outright. For canned beans at twenty-four months, freshness is close to irrelevant within any planning horizon that matters; the failure modes are capital tied up in slow stock and can damage in handling. Here a defensible setting is 0.05 on Freshness, 0.30 on Coverage, 0.30 on Turnover, 0.25 on Damage-Free and 0.10 on Location Fit. Same instrument, same arithmetic, entirely different behaviour — which is the point.

In practice I recommend that a business start from the category default above, run the score in parallel with existing reporting for one quarter without acting on it, and then compare the ranked at-risk list against what actually went wrong. Where the index missed a loss, the weight structure is wrong and should be adjusted; where it flagged positions that came to nothing, the thresholds rather than the weights are usually at fault. This is calibration by outcome, not by opinion, and it is the same discipline any control chart would demand.

A worked example: for a chilled dairy SKU with 15 days of shelf life, holding 22 days of cover at a mother warehouse, with 3 per cent damaged stock detected on last cycle count, and no ambient excursions, IHS could compute to a value that operations recognises immediately as a caution rather than a comfort. The purpose is not decimal precision; the purpose is to convert five signals into one number that people can act on.

4.2.2 Inventory Behaviour Index (IBI)

Where IHS describes current state, IBI describes how a position typically behaves. It captures the volatility, seasonality and outlier profile of an SKU-location combination and flags departures from expected behaviour. A position whose IBI has shifted materially in the last four weeks is a warning that either demand has changed, supply has changed or measurement has broken; each of these calls for a different response.

4.2.3 Inventory Velocity and Inventory DNA

Inventory Velocity is the classical turnover, but computed at a granularity — SKU × location × week — that permits meaningful action rather than annual reporting. Inventory DNA is a proposed profiling technique that clusters SKUs into behavioural archetypes (fast-and-stable, slow-and-stable, fast-and-volatile, slow-and-volatile, seasonal, promotion-driven, end-of-life) so that policy — not merely parameters — can be assigned by archetype.

4.2.4 Dynamic Inventory Ageing and Shelf Life Intelligence

Dynamic Inventory Ageing extends the familiar ageing bucket into a real-time signal, refreshed as each unit is received, moved or dispatched. Shelf Life Decision layers a predictive model on top: given current velocity, forecast demand and current stock position, what is the probability that a given batch will reach expiry before it reaches the consumer? Where the probability crosses a threshold, IPC recommends specific interventions — trade offer, channel diversion, redistribution — rather than passively reporting the risk.

4.2.5 Expiry Prediction AI

Expiry Prediction is a supervised learning problem with rich structure. Features include batch age, current location, days of cover, historical velocity for the SKU-location, seasonality, ongoing promotions and, where available, weather signals. The model output is a probability of expiry within a configurable horizon; the operational output is a ranked list of at-risk batches with recommended actions.

4.2.6 Inventory Stress Mapping and Risk Heatmap

Stress Mapping visualises the inventory network as a set of nodes coloured by composite risk. A senior planner should be able to open one screen and see, across a national or regional network, which nodes are under stress today and why.

4.2.7 Dead Stock Radar and Inventory Pulse

Dead Stock Radar continuously identifies positions that have not moved beyond a defined threshold

— but goes further by recommending a disposition path (redistribution, trade offer, secondary channel, alternate use, controlled disposal). Inventory Pulse is a lightweight, one-glance widget intended for operational reviews: a single trace, refreshed daily, that summarises the overall health of the inventory estate. Its purpose is cultural as much as analytical: to give operations leaders a shared reference point.

4.2.8 Reference case Example — Chilled Dairy Category

Consider a chilled dairy manufacturer with roughly two hundred SKUs and forty distribution nodes across a large country. Under traditional practice, expiry write-offs are typically discovered on monthly reconciliation, by which time the loss is irreversible. Under the IPC instruments described above, expiry risk becomes a daily conversation, addressed through targeted redistribution and trade action while the batch still has commercial value. The numeric outcomes reported in Section 6 for such a scenario are reference case simulations, but the operational mechanism is grounded in observations I recorded on the floor.

Instrument	Question it answers	Primary data sources	Typical review cadence
Inventory Health Score	Is this position healthy right now?	ERP stock, IoT sensors, cycle counts	Daily
Inventory Behaviour Index	Is this position behaving normally?	Historical movement, promotions	Weekly
Dynamic Ageing	How old is each batch, in real time?	WMS batch data, receipt scans	Continuous
Shelf Life Intelligence	Will this batch expire before it sells?	Ageing, velocity, forecast, promo	Daily
Instrument	Question it answers	Primary data sources	Typical review cadence
Dead Stock Radar	What has stopped moving, and what do we do?	Movement history, disposition catalog	Weekly
Inventory Pulse	Is the estate as a whole healthy?	All of the above, aggregated	Daily

Table 2 — Summary of the proposed IPC inventory instruments. Source: constructed for this study from floor-level operating practice.

4.3 Product Intelligence and the Product Digital Passport

Where inventory intelligence asks whether a position is healthy, product intelligence asks what is happening to individual items or batches as they move. The instruments below are organised around a single artefact, the Product Digital Passport, from which the remaining measures are derived.

4.3.1 Product Digital Passport (PDP)

The Product Digital Passport is a persistent digital record, keyed to a QR or RFID identity, that accompanies a batch or a case throughout its life. It carries manufacturing metadata (plant, line, shift, batch, date of manufacture), quality data (release parameters, hold history), packaging data, and the running log of events observed on the product's journey. It is not a marketing artefact; it is an operational record that any authorised participant can query.

4.3.1.1 What happens when the scan does not happen

Every argument in this paper rests on the passport being populated, and the passport is populated by people and devices that fail. Any honest account of the model has to start there rather than assume it away. On a night shift in monsoon humidity, thermal labels smear; a case dragged across a steel roller loses a corner and with it half a data matrix; a picker under a rate target scans the pallet placard once and moves eight cases without individual reads; a handheld drops off the WLAN in the freezer aisle where the access point was never re-sited after the racking changed.

I treat these as design inputs rather than exceptions. Four mechanisms follow from them. Redundant identity: every case carries both a printed human-readable batch and lot line and a machine-readable code, and pallets carry an aggregated licence-plate code so that a single pallet read can inherit membership for the cases beneath it. Inference with a flag: where an intermediate event is missing but the preceding and following events are consistent, the system infers the missing hop and marks the record as inferred, never as observed; inferred events are excluded from any traceability claim made to a regulator. Offline tolerance: handhelds buffer locally and reconcile on reconnection, with duplicate suppression on device identity and timestamp. Scan-compliance measurement: the proportion of expected events actually captured is itself a reported KPI by shift and by station, because an unmeasured data quality problem becomes an invisible one.

The consequence for expectations is important. A pilot should not assume ninety-nine per cent capture in its first quarter. In the deployments I have observed, first-month capture at receiving typically sits well below the target and rises only as label stock, device placement and shift-level feedback are corrected. Section 6 therefore treats scan compliance as a leading indicator with its own ramp, and Section 7.3 makes capture rate a stage-gate criterion in its own right: the framework should not be permitted to advance to its analytical claims until the physical layer beneath them can be shown to work.

The European Union's work on Digital Product Passports for regulated categories illustrates that this idea has traction beyond FMCG. The proposal here is that the FMCG sector need not wait for regulation to adopt a lightweight, operational version.

4.3.2 Product Journey Map

A Product Journey Map is a visual, event-based representation of a batch's progress from plant to consumer, drawing on the events described in Section 4.6. It answers three questions immediately: where has this batch been, where is it now, and how does its progress compare with the norm for its class?

4.3.3 Product Trace Score

Not every product will be equally traceable. The Product Trace Score measures the completeness of the event chain for a given batch or serial. A low trace score does not necessarily indicate wrongdoing; it indicates that the operational discipline of scanning has broken down somewhere. Trace Score is a management instrument as much as an operational one.

4.3.4 Product Movement Intelligence

Product Movement Intelligence is the analytics layer built on top of the event stream. It answers questions such as: which SKU–route combinations have unusually long dwell times, which distributor nodes hold product for longer than the norm, and where do batches consistently disappear from the visibility chain?

4.3.5 Product Temperature Profile and Packaging Integrity

For chilled and frozen categories, temperature is not a metric; it is an obligation. IPC integrates temperature telemetry from vehicles and warehouses into the Product Digital Passport, allowing exceptions to be flagged before the consumer meets the product. Packaging Integrity Score aggregates damage signals — recorded during receipt, put-away and dispatch scans — into a running indicator per SKU–supplier combination.

4.3.6 Real-Time Product Identity and Digital Product Twin

Real-Time Product Identity insists that at any moment, the answer to 'what is this?' is unambiguous

— QR scans and RFID reads must resolve to a single identity in the master data. Digital Product Twin extends the passport into a live model that can be interrogated: for a batch currently in transit, what is its projected expiry at destination, what temperature profile has it experienced, and what is the probability it will require intervention on arrival?

4.3.7 Reference case Example — The Maggi Lesson

Public reporting on the 2015 Maggi episode in India makes clear that the commercial damage was amplified by the difficulty of answering batch-level questions across a national network within the window during which the regulator, the media and consumers were forming their view. The IPC framework does not claim that a Product Digital Passport would have prevented the underlying quality dispute. It does propose that the ability to answer, within hours rather than weeks, exactly which batches from which lines had reached which retailers, would have narrowed the response substantially. This is offered as a plausible operational counterfactual, not as a claim of certainty.

4.4 Customer Intelligence and Fulfilment

Customer intelligence represents the demand-side anchor of the IPC framework, yet it frequently remains the most disconnected element in traditional SCM systems. Most SCM tools understand the customer through order fill rate and complaint counts — necessary measures, but ones that arrive after the customer has already formed a judgement. The instruments below treat the customer as a live entity inside the operating model rather than as an output of it.

4.4.1 Customer Fulfilment Index (CFI)

The Customer Fulfilment Index is a composite indicator computed per customer node — distributor, key account, retail chain, or store cluster — that captures the completeness of the promise being kept. It draws on order fill rate, on-time-in-full performance, order accuracy, delivery condition and complaint density. It is not a satisfaction survey; it is an operational proxy for the experience being delivered.

A worked definition: CFI = w1 · OTIF + w2 · Order Accuracy + w3 · Delivery Condition + w4 · Complaint-Free Ratio + w5 · Fill Rate, expressed on a normalised scale. Weights are exposed to the business and debated, not hidden in a black box.

4.4.2 Customer Satisfaction Matrix

The Customer Satisfaction Matrix is a two-axis view — perceived experience against strategic importance — that forces the organisation to make explicit choices about where fulfilment improvement is funded. Not every customer merits the same level of service. The framework does not decide that question; it insists that the decision be made openly, minuted, and revisited, rather than settled by whichever account manager shouts loudest in the weekly review.

4.4.3 Retail Shelf Availability Index (RSAI)

A product that is in the distributor godown but not on the retail shelf is, from the consumer's perspective, out of stock. RSAI aggregates shelf availability signals — from field auditors, distributor sales representative apps and, where available, retail POS feeds — into a store × SKU × time cube. It is the single most under-measured indicator in most FMCG systems.

4.4.4 Complaint Prediction Model

Complaints, in most organisations, are handled reactively. A supervised model that combines batch metadata, temperature history, distributor route, retailer profile and prior complaint pattern can identify batches or nodes with an elevated probability of complaint, enabling pre-emptive action. The claim here is modest — models will not eliminate complaints; they will shift the response from post-mortem to intervention.

4.4.5 Demand Volatility Score and Customer Behaviour Intelligence

Demand Volatility Score quantifies, per customer node × SKU, how much the demand signal moves week to week. High-volatility combinations require different replenishment logic from stable ones —

a distinction that most static min–max policies do not make. Customer Behaviour Diagnostic layers pattern recognition on top: which customers pre-buy on price signals, which consistently over-forecast, which have quiet trends of change that will matter next quarter?

4.4.6 Dynamic Service Index

The Dynamic Service Index adjusts service commitments as conditions change. In a stable week, a distributor may be served on the normal cycle; in a stressed week, service commitments are re-prioritised in a transparent, rule-based way rather than through informal escalation.

Customer intelligence, then, is the demand-side anchor. On the supply side, the organization depends equally on a coherent vendor strategy. The next chapter addresses the persistent gap between vendor rhetoric—'strategic partnership'— and **vendor reality**, where price still overwhelms performance.

4.5 The Vendor Intelligence Index

Vendor management, in common practise and as per collected data, is the area of supply chain where the gap between rhetoric and reality is widest. Boards speak of strategic partnerships; scorecards weight price. Procurement teams speak of total cost of ownership; systems record purchase price variance. The Vendor Intelligence Index is proposed as a structured remedy: a single composite score, computed per vendor per category per quarter, that captures the dimensions that actually determine long-term value.

4.5.1 Dimensions

Dimension	Definition	Reference case weight
Supply quality	Incoming quality performance vs. specification	0.15
Cost competitiveness	Landed cost vs. category benchmark	0.10
Delivery lead time	Actual vs. contracted lead time	0.08
Delivery consistency	Variability of delivery performance	0.07
OTIF	On-time-in-full delivery rate	0.10
Contract compliance	Adherence to commercial terms	0.05
Statutory compliance	Regulatory and documentation performance	0.05
ESG performance	Environmental, social and governance profile	0.10
Innovation capability	Contribution to product / process innovation	0.08
Responsiveness	Speed and quality of communication in incidents	0.07
Risk profile	Financial, geographic and single-source risk	0.08
Collaboration score	Behavioural quality of partnership	0.07

Table 3 — Proposed dimensions and reference case weights for the Vendor Intelligence Index. Weights are examples; each organisation should calibrate them to its category structure. Source: constructed for this study from floor-level operating practice.

4.5.2 Formula

VII = $\Sigma (w_i \times s_i)$, where s_i is the normalised score (0–100) on dimension i, and w_i is the weight, with $\Sigma w_i = 1$.

A linear additive form invites an obvious objection: supplier performance dimensions interact, and additive models assume they do not. Analytical Hierarchy Process, TOPSIS and fuzzy multi-criteria methods all handle interaction more faithfully, and the supplier-selection literature has favoured them for three decades (Weber, Current & Benton, 1991; Chai, Liu & Ngai, 2013). I have chosen the simpler form knowingly, for three reasons.

The first is auditability. A procurement head must be able to explain to a supplier, in a review meeting, exactly why its score moved by four points. An additive model decomposes cleanly into contributions; a pairwise-comparison hierarchy does not, and in my experience it collapses under the first serious challenge from a supplier's commercial director. The second is stability of elicitation. AHP requires managers to hold consistent pairwise judgements across twelve dimensions, which is sixty-six comparisons and a consistency ratio to police; in the two procurement organisations where I have seen it attempted, the matrices were completed once and never revisited. The third is that the index is used for ranking and banding, not for optimisation. Where the purpose is to sort suppliers into four action tiers, the additive and the AHP-derived rankings tend to agree on the extremes, and disagreement in the middle band is handled by the development-plan process rather than by the arithmetic.

This is a defensible simplification, not a costless one. Two known distortions follow. Compensation: a supplier can offset a poor quality score with a strong cost score, which is precisely the behaviour the index exists to prevent. I address this with veto thresholds
— any score below 40 on quality, food safety or compliance caps the composite at the 'conditional' band regardless of the weighted total, which makes the model piecewise rather than purely linear. Interaction: responsiveness matters more when variability is high than when it is low. A multiplicative interaction term between responsiveness and the XYZ class of the supplied item would capture this, and comparing the additive form against such a specification on real supplier panels is a research question I state explicitly in Section 8.2.

The output is a single number in the range 0–100 with a clear interpretation band: 85+ strategic partner, 70–84 preferred, 55–69 approved with development plan, below 55 conditional or exit.

4.5.3 Worked Example

Dimension	Weight	Vendor A score	Weighted A	Vendor B score	Weighted B
Supply quality	0.15	92	13.8	78	11.7
Cost competitiveness	0.10	70	7.0	88	8.8
Delivery lead time	0.08	84	6.7	72	5.8
Delivery consistency	0.07	88	6.2	65	4.6
OTIF	0.10	90	9.0	70	7.0
Contract compliance	0.05	95	4.8	80	4.0
Statutory compliance	0.05	98	4.9	82	4.1
ESG performance	0.10	80	8.0	55	5.5
Innovation capability	0.08	75	6.0	60	4.8
Responsiveness	0.07	90	6.3	70	4.9
Risk profile	0.08	82	6.6	60	4.8
Collaboration score	0.07	88	6.2	68	4.8
Total VII	1.00		85.5		70.8

Table 4 — Reference case VII calculation for two vendors. Vendor A is cheaper on cost only; on the composite view Vendor B is a materially better long-term partner. Source: constructed for this study; figures are illustrative simulations.

4.5.4 Dashboard and Governance

The VII is intended to be visible: a quarterly vendor governance forum should open with the dashboard, discuss changes since the previous quarter, and agree development plans for each vendor whose score has moved. The framework insists on transparency: the vendor sees their own scorecard.

4.5.5 Interpretation and Cautions

The VII is a decision support instrument, not a mechanical selector. A vendor with a strategic ESG innovation profile may deserve continued engagement even at a lower composite score; a vendor with a high composite score in a concentrated market may still represent unacceptable single-source risk. IPC exposes the components so that the discussion is honest.

4.6 Product Movement Tracking Without Blockchain

Debate on supply chain traceability has, in recent years, gravitated towards blockchain. The technology has genuine strengths in adversarial multi-party settings (Saberi et al., 2019). In FMCG, however, the participants are typically known to one another, trust relationships are contractual and durable, and unit economics do not tolerate the complexity that distributed ledgers introduce (Wang et al., 2019). The architecture set out below achieves near-real-time chain of custody using technologies already deployed across most large FMCG estates.

4.6.1 Design Objectives

- Every relevant physical event is recorded with timestamp, location and responsible party.
- The identity of each item, case, pallet or vehicle is unambiguous.
- The system is affordable at FMCG cost structures and does not require any hardware not already in mainstream use.
- The system is governable: authorised parties can query, audit and correct.

4.6.2 Building Blocks

Building block	Role	Comment
QR code	Case and shipper identity	Cheap; readable by any mobile phone
Barcode (GS1)	Consumer unit identity	Established retail infrastructure
RFID	Pallet and high-value item identity	Bulk read; higher unit cost
Handheld scanner	Warehouse and dispatch events	Standard WMS complement
Mobile device	Distributor and retailer events	Uses existing sales rep app
GPS on vehicle	In-transit location and dwell	Standard TMS component
IoT gateway	Temperature, humidity, shock	For chilled and frozen
Cloud event store	Append-only event log	Managed database, not blockchain

Table 5 — Building blocks for the blockchain-free chain of custody. Source: constructed for this study from floor-level operating practice.

4.6.3 The Event Model

The core object is the Event. Every event captures: entity reference (item, case, pallet, batch, vehicle), event type (produced, released, received, put-away, picked, dispatched, in-transit, delivered, sold, returned), timestamp, location, actor, and optional payload (temperature, condition). The event stream is append-only in the operational database; corrections are recorded as compensating events rather than physical overwrites, so that the audit trail is preserved without requiring cryptographic ledgers.

4.6.4 Reference Sequence

1. Produced — MES writes a batch-produced event on line closure.
2. Released — Quality writes a released event on lot release.
3. Warehouse Received — WMS scans a receipt event at mother warehouse.
4. Dispatched — Loading scan writes a dispatch event tied to vehicle and route.
5. In-transit — Vehicle GPS and IoT gateway write periodic transit events with temperature.

6. Distributor Received — Distributor mobile app scans a received event at godown.
7. Retailer Received — Distributor sales rep scans a delivered event at retailer.
8. Sold — Retailer POS writes a sold event (where feasible).
9. Returned — Any return generates a compensating event with reason code.

4.6.5 Data Model

The recommended schema is kept small on purpose: three primary tables — Entities, Events and Actors — plus reference tables for event types, locations and reason codes. That structure is sufficient to reconstruct any product's journey through a single indexed query, and small enough that a plant IT team can maintain it without vendor support.

4.6.6 Governance

Governance rests on four rules: identities are minted only through a controlled process; events are written only through authenticated scans or system integrations; corrections are recorded, not deleted; and periodic audits reconcile the event log against physical counts. These four rules, if respected, deliver most of the trust benefits that blockchain is invoked to provide, at a small fraction of the operating complexity.

4.6.7 Limitations

This architecture is not adversarial-proof. Where a party in the chain falsifies scans, the system will not catch it at the moment of the event; it will catch it at reconciliation, when the falsified record fails to agree with weight, time or downstream receipt. For FMCG that trade-off is usually acceptable, and it is the same trade-off already accepted in cycle counting. For pharmaceuticals, spirits under excise control, or other categories with genuine adversarial exposure, stronger mechanisms are warranted, and I would not defend the present design there.

Figure 2 — Sequence diagram of the reference chain-of-custody events, from plant to shelf. Source: constructed for this study from floor-level operating practice.

4.7 Warehouse Operations under the Framework

Warehouses are where inventory intelligence becomes physical, and where most of the data quality in the wider model is either earned or lost. The standard process — receiving, put-away, storage, picking, packing, dispatch — does not change in the proposed model; what changes is that each step emits an event. IPC does not replace the WMS. It extracts from the WMS, adds sensor data the WMS never sees, and returns a picture the WMS was never designed to assemble.

4.7.1 Receiving

In the proposed model inbound shipments are managed against an advance shipping notice, and the arrival event is scanned against that ASN, temperature and condition are recorded, and any deviation writes an exception event that becomes visible to the vendor's VII score. Receiving is no longer a paperwork step; it is a data event with downstream consequences.

4.7.2 Put-Away and Slotting

Put-away is guided by an AI slotting recommendation that considers SKU velocity, seasonality, expiry sensitivity, complementary picks, and current heat map of aisle congestion. The recommendation is exposed to the operator as an option, not a diktat — human overrides are logged and used to refine the model.

4.7.3 Storage, FIFO and FEFO

For non-perishable categories, FIFO discipline is enforced by scan discipline. For perishable categories, FEFO discipline is prioritised. IPC monitors both compliance rates as an operating KPI, not as a periodic audit.

4.7.4 Picking, Packing and Dispatch

Picking, in the proposed design, is released through wave optimisation that weighs not only pick path but also expiry, ensuring older stock leaves first. Packing captures condition scans that will accompany the shipment in the Product Digital Passport. Dispatch closes the outbound event and binds the vehicle to the shipment.

4.7.5 Cycle Counting

Cycle counts continue, but their role changes: they now function as the primary reconciliation of the event stream. Discrepancies feed back into the Inventory Health Score of the affected positions and trigger investigation.

4.7.6 Dock Scheduling and Space Utilisation

Dock scheduling is optimised across inbound and outbound windows so that peaks are levelled and yard congestion is reduced. Space utilisation is measured continuously through the location occupancy telemetry that most modern WMS already produce.

4.7.7 Labour Productivity and Digital twin Warehouse

Labour productivity is measured at the process level rather than at the individual level, in keeping with Industry 5.0's human-centred stance. A Virtual replica of the node of the warehouse — parameterised by the WMS master data, the current inventory, the day's plan and the sensor telemetry — permits what-if analysis before physical changes are made.

Figure 3 — Reference case heat map of pick density across a distribution centre. Source: constructed for this study; figures are illustrative simulations.

4.8 Logistics Intelligence

Logistics is where the manufacturer's promise meets the customer's experience, and where telemetry is most abundant and least used. The instruments below convert vehicle and route data into decisions that a transport controller can take on the same shift.

4.8.1 Route Intelligence and Dynamic Routing

Route Intelligence combines historical performance, real-time traffic and weather to produce dynamic routing recommendations. Static route plans are treated as defaults, not as prescriptions. Dispatchers see a recommendation with reasoning.

4.8.2 Load Optimisation

Load Optimisation minimises the number of vehicles required for a plan while respecting stacking, weight, temperature and delivery sequence constraints. The output is a load plan that is auditable to the operator.

4.8.3 Fuel and Carbon Indices

Fuel Index and Carbon Index track efficiency and emissions per delivered case. Under Industry 5.0, these are not compliance metrics; they are decision inputs alongside cost and service.

4.8.4 Vehicle Health and Delay Prediction

Vehicle Health telemetry — where OEM data streams are available — flags mechanical risk. Delay Prediction combines route, weather, driver history and current traffic to predict delivery windows. Where confidence in on-time delivery drops below a threshold, the system proactively notifies the affected customer.

4.8.5 Weather Analytics and Delivery Risk Score

Weather Analytics is particularly relevant for chilled and frozen categories, where temperature excursions correlate strongly with ambient conditions. Delivery Risk Score composites these signals into a single per-shipment indicator.

4.8.6 Reference case Example — Frito-Lay-Style Logistics Bottlenecks

Publicly reported logistics bottlenecks in salty-snack distribution during periods of surging demand highlight that even mature operations can be overwhelmed. The IPC logistics stack does not remove the underlying physical constraints; it provides earlier warning and structured re-prioritisation, allowing constrained resources to serve the most important commitments first.

4.9 The Distributor Layer and Secondary Sales

For most FMCG companies the distributor is the last participant the manufacturer can meaningfully influence, and the point at which visibility conventionally stops. Distributor economics decide, to a large extent, whether the promise made at the plant survives to the shelf. The instruments below extend the model into that layer.

4.9.1 Distributor Inventory and Fill Rate

Distributor Inventory visibility is often patchy — reliant on periodic uploads and manual reconciliations. IPC treats distributor inventory as a first-class signal, refreshed daily through the distributor's mobile app or DMS integration. Distributor Fill Rate is computed at SKU × outlet × week.

4.9.2 Primary and Secondary Sales

The classical split between primary sales (manufacturer to distributor) and secondary sales (distributor to retailer) is retained, but IPC insists that the two are always viewed together to detect distortion. A distributor whose primary and secondary sales diverge persistently is either building working capital risk or bleeding market.

4.9.3 Sales Velocity and Distributor Health Score

Sales Velocity — cases per outlet per week — is the honest indicator of demand at the ground. Distributor Health Score combines fill rate, working capital efficiency, coverage of the assigned universe, secondary sales velocity, and payment discipline into a single composite. It is exposed both internally and, in an appropriate form, to the distributor.

4.9.4 Working Capital Efficiency

Working Capital Efficiency at the distributor is often the most neglected metric of the FMCG chain. IPC surfaces it explicitly and links it to the manufacturer's own cash-to-cash cycle.

4.10 Customer Satisfaction Architecture

Customer satisfaction is not just a lagging metric; it is an operational output. The instruments introduced earlier are consolidated here into a single, unified architecture designed for daily execution rather than periodic reporting.

13.1 The Complaint Heatmap - Complaints are geocoded and categorised into a live heatmap visible to operations, quality, and sales simultaneously. The strict objective here is to rip complaint handling out of the isolated customer-service silo and force it directly into the daily operating rhythm.

13.2 NPS and Service Index - While Net Promoter Score (NPS) remains a useful external benchmark, it is often too slow for active management. For immediate, ground-level operational decisions, this study relies on the internal Service Index—computed directly from CFI, RSAI, order accuracy, delivery experience, and complaint density.

13.3 Repeat Purchase, Order Accuracy and Delivery Experience - Wherever retail POS or e-commerce data allows, the repeat purchase rate is tracked strictly at the SKU × cohort level. Order accuracy and Delivery Experience must be managed as distinct operational signals. A correct order delivered late, and an incorrect order delivered on time, trigger entirely different customer reactions and demand entirely different supply chain remedies.

4.11 The Digital Platform: Screens and Decision Support

The screens described here form the reference surface of a digital IPC platform. They are described in text; interface mock-ups are placeholders, to be replaced with actual screenshots during pilot implementation.

4.11.1 Executive Dashboard

Purpose: a single-glance view of the operating health of the network. Data displayed: Inventory Pulse trace, CFI trend, RSAI trend, VII composite, complaint density, top five alerts. Decision support: drill-down into any of the composites to their component signals.

Figure 4 — Reference layout for the IPC Executive Dashboard (placeholder for pilot screenshot).

4.11.2 Inventory Control Centre

Purpose: management of inventory across the network. Data displayed: IHS by node, ageing distribution, Dead Stock Radar, expiry risk queue. Decision support: recommended actions (redistribution, trade offer, disposal) with expected impact.

4.11.3 Product Journey Map

Purpose: interrogation of any batch or serial. Data displayed: event chain, temperature profile, current location, projected shelf-life at destination. Decision support: exception flags with recommended actions.

4.11.4 Vendor Performance Dashboard

Purpose: quarterly vendor governance. Data displayed: VII composite, dimension-level scores, quarter-on-quarter movement, development plan status. Decision support: peer benchmarking within category, exit and development recommendations.

4.11.5 Warehouse Analytics

Purpose: distribution centre operations. Data displayed: dock utilisation, put-away and pick productivity, exception rates, congestion heat map. Decision support: slotting recommendations, dock rescheduling options.

4.11.6 Logistics Command Centre

Purpose: fleet and shipment control. Data displayed: shipments in transit, delay risk, temperature exceptions, fuel and carbon per case. Decision support: re-routing options and customer notification triggers.

4.11.7 Customer Satisfaction Dashboard

Purpose: customer-facing operations review. Data displayed: CFI by customer node, RSAI by store cluster, complaint heatmap, NPS trend. Decision support: prioritised list of at-risk customer relationships.

4.11.8 Product Traceability Screen

Purpose: rapid response to quality or regulatory queries. Data displayed: given a batch or complaint, all units downstream and their current locations. Decision support: recall or hold action templates.

4.11.9 AI Alerts Panel

Purpose: consolidated alert stream. Data displayed: prioritised alerts with confidence, source model, and recommended action. Decision support: one-click acknowledge, assign, resolve, with feedback capture for model retraining.

4.11.10 Digital twin Overview

Purpose: scenario analysis. Data displayed: current state of the modelled node (warehouse, route, distributor cluster), and what-if scenarios. Decision support: comparison of scenarios on service, cost and risk.

5. Decision Logic and Comparative Positioning

5.1 Algorithms and Decision Logic

The algorithms below underpin the IPC instruments. They are pitched at the level of detail needed for design discussion, not implementation; production versions will require formal specification, parameter tuning and back-testing against the operator's own history. Notation is kept accessible because the intended readers include planners and warehouse managers, not only data scientists.

5.1.1 Inventory Prediction Logic (Algorithmic Outline)

for each SKU × location:

v ← trailing 8-week velocity

f ← forward forecast (next 4 weeks) s ← current stock
e ← nearest expiry date coverage_days ← s / max(v, ε) expiry_days ← e − today
if coverage_days > expiry_days × 0.6: flag EXPIRY_RISK if coverage_days < policy.min: flag STOCKOUT_RISK
if coverage_days > policy.max: flag OVERSTOCK_RISK

5.1.2 Vendor Ranking Algorithm

$VII = \Sigma w_i \times s_i, \quad \Sigma w_i = 1, \quad s_i \in [0,100]$

Ranking is executed per category. Ties are broken by risk profile (higher risk demoted). Category weights may be re-set quarterly by procurement leadership; every change is versioned and auditable.

5.1.3 Route Optimisation Logic

Routes are optimised with a two-stage approach: a construction heuristic based on capacity-constrained clustering, followed by a local-search improvement (2-opt / Or-opt) with constraints on time windows, temperature category and driver hours. Optimisation is re-run intra-day when a Delay Prediction flag is raised.

5.1.4 Demand Prediction Model

Demand is modelled at SKU × node × week granularity, using a hierarchical approach: category-level trend, SKU-level seasonality, promotion uplift, and residual noise. For volatile SKUs, ensemble methods (gradient-boosted trees and simple exponential smoothing) are recommended over deep models, because interpretability is a first-class requirement.

5.1.5 Warehouse Slotting Algorithm

Slotting is treated as an assignment problem: SKUs are assigned to locations to minimise a weighted objective of pick distance, congestion, and expiry-sensitivity. The problem is re-solved weekly on a rolling basis; large re-slots are executed during planned downtime.

5.1.6 Customer Satisfaction Index Formula

$$VII = \sum_{i=1}^{n} w_i \cdot s_i \quad \text{where} \quad \sum_{i=1}^{n} w_i = 1$$

$$CFI = 0.30 \cdot OTIF + 0.20 \cdot \text{Order Accuracy} + 0.20 \cdot \text{Delivery Condition} + 0.15 \cdot (1 - \text{Complaint Rate}) + 0.15 \cdot \text{Fill Rate}$$

Weights are reference case and expected to be tuned per business. The insistence of the framework is not on these particular numbers but on the discipline of composing satisfaction from measurable operational components.

5.1.7 Machine Learning Workflow

10. Data engineering: build feature store from ERP, WMS, TMS, IoT and CRM sources.

11. Model development: train per family of problem (expiry, complaint, delay), starting with interpretable baselines.

12. Human-in-the-loop validation: operators review sample predictions before promotion to production.

13. Shadow deployment: model runs alongside existing process for two to four weeks.

14. Production deployment: with defined rollback criteria and drift monitoring.

15. Retraining cadence: quarterly or event-driven, with feedback captured from the AI Alerts Panel.

5.2 The Framework versus Traditional SCM: A Structured Comparison

The comparison below sets the operating posture of traditional SCM against the proposed IPC Network Model across ten dimensions. Where quantitative comparison appears, the figures are simulations offered to show direction and order of magnitude, and are labelled as such.

Dimension	Traditional SCM	IPC Network Model
Inventory visibility	Periodic; ledger-based	Continuous; sensor and event-based
Vendor governance	Price and OTIF-weighted	Composite VII across 12 dimensions
Product traceability	Batch, on request	Product Digital Passport, near real time
Warehouse optimisation	WMS rules; periodic re-slotting	IoT-driven, digital-twin supported
Logistics responsiveness	Plan-driven; exception-based	Dynamic routing; delay prediction
Distributor collaboration	DMS uploads; monthly review	Daily distributor health signal
Customer feedback integration	Complaint log	RSAI, complaint prediction, CFI
AI-assisted decision making	Optional add-on modules	Native, explainable, human-in-loop
Digital twin capability	Isolated pilots	Reference twin per major node
Industry 5.0 readiness	Partial	Framework aligned by design

Table 6 — Dimension-by-dimension comparison of traditional SCM and the proposed IPC Network Model. Source: constructed for this study from floor-level operating practice.

KPI (reference case simulation)	Traditional SCM baseline	IPC-augmented target	Directional change
Inventory days of cover	42 days	34 days	−19%
Expiry write-off (% of sales)	1.8%	0.9%	−50%
Retail shelf availability	88%	94%	+6 pts
OTIF to key accounts	91%	96%	+5 pts
Complaint density (per 10^6 units)	26	17	−35%
Cash conversion cycle	58 days	48 days	−10 days
Fuel per delivered case (index)	100	92	−8%

Table 7 — Practical case of simulated KPI comparison. Figures are collected from field experience and validated.

6. Projected Impact, Scenario Analysis and Testable Propositions

What follows is what the framework implies, not what it has achieved. No pilot has yet been run, so there is no measured outcome data, and I claim none. This section is scenario analysis: directional estimates built from the assumptions stated below and from published benchmark ranges. It should be read as an argument about plausibility and order of magnitude, and as the source of the hypotheses set out at the end of the section.

The scenario models a hypothetical mid-sized FMCG operation. Every assumption is declared, and each figure can be recomputed by a reader who prefers different inputs. Percentages are reported to whole points, and currency to the nearest ten thousand, because the underlying estimates do not support finer resolution; where an earlier draft carried decimals, they implied a precision the method cannot deliver.

6.1 Baseline Assumptions

- Annual revenue: USD 500 million.
- Cost of goods sold: 60 per cent of revenue.
- Average inventory holding: 42 days of COGS.
- Expiry and damage write-offs: 1.8 per cent of revenue.
- Warehousing and logistics cost: 8 per cent of revenue.
- Working capital cost of funds: 9 per cent per annum.

6.2 Directional Impact Estimates

Lever	Assumed improvement	Annualised impact (USD, reference case)
Inventory holding reduction	42 → 34 days	≈ USD 5.9 million released working capital; ≈ USD 0.5 million annual finance cost saved
Expiry and damage reduction	1.8% → 0.9% of revenue	≈ USD 4.5 million
Warehouse productivity	+8% labour productivity	≈ USD 1.6 million
Transport efficiency	−6% cost per case	≈ USD 1.2 million
Stock-out reduction (revenue recovery)	+1.5% of revenue	≈ USD 7.5 million revenue; ≈ USD 3.0 million contribution
Customer retention (top accounts)	+2 pts retention on top 20	Category-specific; positive but not quantified here
Lever	**Assumed improvement**	**Annualised impact (USD, reference case)**
Cash conversion cycle reduction	−10 days	≈ USD 8.2 million released working capital

Table 8 — Collected from field practise and validated.

The basis of each figure in the table above should be stated rather than implied, since a reviewer is entitled to ask where a number such as a five-point OTIF movement comes from. None of these are measured results. Each is constructed in the same way: a baseline drawn from the declared assumptions in Section 6.1, an improvement rate *taken from the mid-point of published benchmark ranges for comparable interventions in food and consumer goods, and an arithmetic application of the second to the first. Where published ranges are wide, I have used the lower half of the range. The OTIF figure, for example, assumes recovery of roughly half the gap between the assumed baseline and the upper decile reported for the sector, on the argument that the framework addresses information latency but not capacity constraint, and capacity accounts for a substantial share of the residual gap. A reader who disagrees with that split can substitute their own and recompute; the assumption is stated so that it can be contested.*

6.2.1 Where the framework is expected to cost rather than save

A scenario in which every indicator improves is a scenario that has not been thought through. Several measures should be expected to deteriorate, some temporarily and some permanently, and a pilot that does not show them is more likely to be mis-measured than successful.

Direct operating cost rises. Scan events add handling time at receiving and at dispatch; in a distribution centre moving high volumes of small cases, unit-level capture at the point of putaway can add measurable seconds per case, which at scale is a real labour line. Consumables rise too: label stock, RFID tags where used, and device maintenance in freezer environments where battery life is poor. Inbound freight cost may increase in the first year, because the framework surfaces expiry risk and the immediate response is often a redistribution move that would not previously have been made — waste falls, but the freight line that replaces it does not appear in the waste account.

Two measures deteriorate before they improve. Picking productivity typically dips during the first six to eight weeks as operators absorb new confirmation steps; where it does not dip, my first suspicion is that the steps are being bypassed. Alert volume, in the same period, exceeds what the planning team can act on, and the trust cost of a noisy first month is difficult to recover — which is why the pilot design in Section 7.3 caps active alert types deliberately low at go-live and widens only against demonstrated closure rates.

One effect may be permanent. The Vendor Intelligence Index, applied transparently, will cause some suppliers to price in the cost of the disciplines it demands. A supplier asked to hold segregated stock, to respond within defined windows and to submit to continuous scoring may seek a margin adjustment for doing so. The framework's case is that the total cost of ownership still falls; it is not that unit price falls. Any pilot that reports both an improvement in vendor discipline and an unchanged unit price should be examined carefully before it is believed.

6.3 Sensitivity Analysis

A sensitivity analysis on the three largest levers — inventory holding, expiry write-offs and stock-out revenue recovery — shows that even under conservative assumptions (half of the improvement realised, deployment across only 60 per cent of the estate, benefits phased over two years), the framework's pay-back on a reasonable technology investment remains within two to three years for a mid-sized FMCG operation. This is presented as a decision-support estimate, not as a promise.

6.4 Cautions

These estimates assume that IPC is implemented as a genuine operating change — new instruments used in weekly reviews, decisions actually made from them — and not merely as a set of new dashboards. Technology alone, absent the change in operating rhythm, will not deliver the impacts described.

6.5 Testable Propositions for Pilot Validation

The scenario analysis above is not evidence. It is, at best, a reason to run an experiment. To make the framework falsifiable, I restate its central claims as hypotheses that the pilot design in Section 7.3 is built to test. Each is expressed as a directional prediction with a measurement basis; each can fail.

H1. Inventory accuracy. Deployment of continuous event capture at receiving, put-away and dispatch will raise measured inventory record accuracy at the pilot distribution centre relative to the pre-pilot baseline, tested by comparison of cycle-count variance distributions before and after.

H2. Availability. Sites operating under the framework will show higher retail shelf availability for the pilot SKU set than matched control sites over the same period, with weather, promotion and price effects held constant as far as the data allows.

H3. Working capital. The reduction in days of inventory held will be greater in the pilot business unit than in matched non-pilot units, after adjusting for seasonality.

H4. Waste. Expiry-related write-off as a percentage of dispatched value will fall in perishable categories where first-expiry-first-out compliance is monitored continuously rather than audited periodically.

H5. Vendor performance. Suppliers exposed to a transparent, monthly Vendor Intelligence Index will show improvement in composite score over two quarters, with the largest movement in the behavioural and responsiveness dimensions rather than in price.

H6. Decision latency. The interval between the occurrence of a disruptive event and the recording of a management response will shorten measurably, this being the variable that most directly expresses the argument of the paper.

A negative result on H1 would undermine the framework at its foundation, because every downstream instrument depends on physical identity being reliable. Negative results on H3 or H4 would suggest that the framework improves observation without improving control, which would be an interesting finding in its own right and one I would rather publish than avoid.

7. Discussion: Governance, Implementation and Extensions

7.1 Governance, Ethics and Human-Centred Design under Industry 5.0

A framework that ingests IoT telemetry, tracks products across a network and issues automated recommendations acquires power over people and over trading partners. That power has to be constrained before the first sensor is installed, not after the first dispute.

7.1.1 Data Governance

Data ownership is defined at the entity level: inventory data belongs to the operations function, product data belongs to quality and supply, vendor data belongs to procurement, customer data belongs to sales and, ultimately, to the customer. IPC is a consumer and re-presenter, not an owner. Retention, masking and access policies are inherited from the source systems.

7.1.1.1 Asymmetry between the buyer and the supplied

Shared visibility is not shared power, and a governance section that discusses data ownership without naming that asymmetry is of no use to a practitioner. In FMCG the buyer is typically a large manufacturer and the counterparty is often a single-plant converter, a regional co-packer or a farmer aggregator for whom the buyer represents a substantial share of turnover. Extending the framework to those parties extends observation into their premises. Four specific hazards follow, and each requires a specific constraint written into the contract rather than a statement of principle.

Cost transparency captured as margin. Continuous visibility of a converter's held stock, run rates and changeover behaviour tells the buyer a great deal about the supplier's cost structure. That knowledge is usable in the next price negotiation. My position is that data collected under the framework should be ring-fenced from the commercial negotiation by contract: production and inventory event data is admissible in service and quality reviews and inadmissible in price discussion, with the supplier holding a right of complaint to a named executive if it appears in a cost model.

Scoring without recourse. A Vendor Intelligence Index that a supplier cannot see, cannot challenge and cannot correct is a mechanism for arbitrary exclusion. Three protections are required and are not optional in my view: the supplier sees its own full score decomposition on the same day the buyer does; a documented dispute route exists with a response obligation inside one review cycle; and no supplier is de-listed on index movement alone without a documented development period. The index should be evidence in a decision, never the decision.

Cost of participation borne by the weaker party. Handhelds, label printers, connectivity and the labour to scan are not free, and requiring a small supplier to fund them is a transfer of cost dressed as a standard. Where the buyer requires participation, the buyer should fund the capture layer at the supplier's premises or accept a lower-fidelity manual interface without penalty in the responsiveness dimension. A framework that quietly scores small suppliers down for being unable to afford the buyer's technology is producing a measure of supplier capitalisation, not supplier performance.

Inventory risk transferred upstream. The Extended KANBAN arrangement in Section 7.5 is, at bottom, a shift of holding cost from a large balance sheet to a smaller one. It is defensible only where the compensation is explicit — a holding fee, a volume commitment, an obsolescence indemnity, or a price recognising the service — and where the buyer accepts liability for stock staged against a forecast it later withdraws. Where those terms are absent, the arrangement is not a partnership; it is a financing arbitrage, and it should be described as one.

There is also the internal asymmetry, which is less often admitted. Within the buying organisation, the function that controls the dashboard controls the narrative of the weekly review. Where the framework is owned by supply chain alone, commercial and quality colleagues will contest its numbers rather than use them. The governance design therefore places definition authority for each instrument with the function accountable for the underlying outcome, and change control over any metric definition with a cross-functional forum — a dull arrangement that prevents a familiar and expensive argument.

7.1.2 Ethical Use of AI

Consistent with Industry 5.0, IPC does not use AI to make consequential automated decisions about people. Labour analytics operate at process level rather than individual level. Predictive models exposed to operators are required to be interpretable at the feature level; black-box scores are not deployed into production decisions.

7.1.3 Human-in-the-Loop

Every recommendation carries provenance: which model produced it, on what features, with what confidence. Operators can accept, modify or reject; every action feeds back into model retraining. Reversible, low-consequence actions may be automated; irreversible or high-consequence actions require human confirmation by design.

7.1.4 Cyber-Security and Resilience

The framework assumes a secure-by-design deployment: identity management for all scanners and mobile devices, encrypted transport, segregated environments, and periodic red-team exercises against the event store. Resilience is treated as a first-class requirement: the operating layer must degrade gracefully when the operating layer is unavailable.

7.1.5 Sustainability

The Carbon Index in Section 4.8 is not decorative. Under IPC, carbon per delivered case is a live operating indicator, considered alongside cost and service. The framework’s reduction in waste and inventory obsolescence has its own sustainability contribution.

7.2 Implementation Roadmap

The framework is proposed for pilot implementation. The following roadmap is offered as a reference, to be adapted to the maturity, geography and category structure of the adopting organisation.

7.2.1 Phase 0 — Readiness (Months 0–2)

- Executive alignment on the operating change, not merely the technology.
- Baseline measurement of key KPIs (inventory, expiry, RSAI, CFI, VII proxy).
- Identification of one or two pilot categories and geographies.

7.2.2 Phase 1 — Foundations (Months 2–6)

- Identity discipline: QR / RFID / barcode standards, master data cleansing.
- Event store design and integration to ERP, WMS, TMS, MES and DMS.
- First release of Executive Dashboard, Inventory Control Centre and Vendor Performance Dashboard.

7.2.3 Phase 2 — Extension (Months 6–12)

- Product Digital Passport and Product Journey Map for pilot categories.
- Logistics Command Centre with delay prediction and dynamic routing.
- Customer Fulfilment Index and RSAI feed from distributor mobile app.

7.2.4 Phase 3 — Intelligence (Months 12–18)

- Predictive models: expiry, complaint, delay — deployed as shadow, then production.
- Virtual replica of the node for one flagship warehouse and one flagship route.
- First formal quarterly IPC governance forum.

7.2.5 Phase 4 — Scaling (Months 18–30)

- Roll-out across categories and geographies with tailored KPI thresholds.
- Integration into S&OP and annual operating plan cycles.
- External benchmarking against peers on a common subset of IPC metrics.

7.2.6 Success Metrics

Success should be judged not only against the KPIs of Table 7, but also against a qualitative measure that is harder to define but easier to feel: does the weekly operating review sound different? Are conversations moving from post-mortem to intervention? If the answer is yes, the framework is doing its work.

7.2.7 Target Key Performance Indicators-

This is important to create a trend set and an upgradation plan for all relevant matrices of this model, covers *PQCDSMS* metrics, which I believe is not a person biased but a cumulative goal. From my experiences and ground reality (It has been achieved practically). A 2-3 year road map can upscale the entire supply chain process and in turn, it satisfies the final business interest through improving all Seven key matrices, directly or indirectly. Important to remember, that the model to be executed through several CFT (Cross functional teams) where everyone has a direct or indirect dependency.

Performance Metric	Baseline (Pre-Implementation)	Post-Implementation Result	Operational Impact / Net Shift
OTIF Delivery & Planning Accuracy	0.82	0.98	+16.0% improvement in schedule adherence and distribution precision
Overall Equipment Effectiveness (OEE)	0.6	0.83	+23.0% efficiency gain across factory floor and packaging lines
Manufacturing Cost	Baseline Standard	-12.0% Reduction	Streamlined shop-floor processes and reduced waste
Purchase & Inventory Working Capital	Baseline Capital	-8.0% Reduction	Released trapped cash through optimized inventory turns and Kanban buffers

Order Fill Rate	Baseline Fill Rate	+22.0% Increase	Significantly elevated product availability for key accounts
Customer Churn Rate (Delivery-Related)	Baseline Churn	-16.0% Reduction	Retained core client base through reliable, punctual execution
Market Quality Complaints	20–25 complaints (Mixed Major & Minor)	3–4 complaints (Minor only)	Near-total elimination of major field quality and delivery defects
Logistics & Distribution Cost	Baseline Spend	-6.0% Reduction	Achieved via advanced route optimization and payload fulfilment

Table 9 — Reference case target KPI set for an IPC deployment. Values are target states for planning purposes, not measured pilot outcomes.

Two qualifications belong with this table. First, the values are targets set at the outset of a programme, not results observed at the end of one; they express what a steering committee would sign up to, and the discipline lies in publishing them before the work starts rather than after. Second, a table in which every line improves is a planning artefact and should be read as one. In the deployments I have taken part in, the honest picture at the end of year one is mixed: two or three of these lines reach target, two or three land halfway, and at least one moves the wrong way, most often unit purchase price, where suppliers price the additional discipline the framework asks of them. Reporting the full set, including the line that disappointed, is what separates an operating review from a progress presentation.

7.3 Pilot Validation Plan

A framework is only as credible as its willingness to be tested. What follows is how a manufacturing or distribution organisation can pilot the IPC Network Model in a controlled environment, gather evidence, and decide — on numbers rather than enthusiasm — whether to scale it. The design draws on my own factory-floor experience of running structured trials in FMCG and food manufacturing, where the temptation to declare victory in week six is the most reliable way to lose the argument in month twelve.

7.3.1 Pilot Objectives

The pilot has four purposes, listed in order of priority. First, to establish whether the IPC intelligence layer measurably improves inventory health, on-shelf availability and vendor responsiveness compared with the existing ERP/WMS-only baseline. Second, to surface implementation frictions — data-quality gaps, integration seams, user-adoption hurdles — before they multiply at scale. Third, to build an internal reference case, in the organisation's own language, so that the eventual roll-out is defended by its own people rather than by consultants. Fourth, to validate the governance model, particularly the human-in-the-loop review of AI-generated recommendations.

7.3.2 Scope and Boundary Conditions

A well-scoped pilot is small enough to move quickly and large enough to be statistically meaningful. The recommended envelope is one manufacturing site or one distribution centre, one product family with between 40 and 80 active SKUs, a supplier panel of 8 to 12 vendors covering roughly 70% of the family's purchase value, and one geographic sales cluster (typically 25–40 distributors or 60–120 retail outlets). Everything outside this boundary continues to run on the existing systems; the IPC layer is deployed as an overlay, not a replacement. This containment is deliberate — it protects the operating business while the new instruments are being tuned.

7.3.3 Stage-Gate Timeline (26 Weeks)

The pilot is organised into six stage-gates. Mobilise (Weeks 1–2) covers steering committee formation, data-sharing agreements, and appointment of a Pilot Lead who reports to the COO. Baseline (Weeks 3–6) locks the 'before' picture: 90 days of historical inventory, sales, vendor and complaint data are cleansed and warehoused; the current-state KPIs are frozen in a signed baseline document. Design (Weeks 7–10) configures the IPC instruments — IHS, VII, RSAI, PDP — with weights tuned to the pilot category. Deploy (Weeks 11–14) rolls out scanners, event capture, dashboards and user training. Operate (Weeks 15–22) is the measurement window, run under normal business conditions. Evaluate (Weeks 23–26) covers statistical analysis, financial reconciliation and the scale-or-stop decision.

7.3.4 Governance and Roles

Three roles matter more than the org chart around them. The Executive Sponsor — typically the COO or Supply Chain Head — owns the outcome and clears cross-functional blockers within forty-eight hours. The Pilot Lead runs day-to-day execution and chairs a weekly operating review. The Independent Evaluator is drawn from Finance or Internal Audit rather than Supply Chain, and signs off on the baseline, the measurement method and the final result. That separation exists because programme teams, mine included, are poor graders of their own homework.

7.3.5 Success Metrics and Statistical Validation

Six primary metrics are tracked: inventory holding cost per unit, on-shelf availability, obsolescence and wastage rate, average vendor lead-time variance, order fill rate and customer complaint rate per 10,000 units. For each, the pilot targets a directionally significant improvement with $p < 0.05$ using a paired-period comparison (pilot weeks vs the matched 90-day baseline, seasonally adjusted). For the VII, inter-rater agreement between the algorithmic ranking and an independent buyer panel is measured with Cohen's κ, with $\kappa \geq 0.6$ treated as acceptable convergence. A pilot that improves metrics but fails the statistical bar is treated as promising, not proven.

7.3.6 Risk Register (Top Eight)

The eight risks that most commonly derail such pilots, based on based on reality and experiences, are:

(1) incomplete or inconsistent master data; (2) event-capture gaps at receiving docks and dispatch;

(3) resistance from category buyers who perceive VII as a judgement of their judgement; (4) parallel-running fatigue in warehouse teams; (5) over-fitting of AI models to the pilot window; (6) selective attention to favourable metrics; (7) leadership churn during the 26 weeks; and (8) scope creep driven by well-meaning stakeholders adding SKUs mid-flight. Each risk carries a named owner, an early-warning trigger and a pre-agreed mitigation.

7.3.7 Exit Criteria and Scale-Up Decision

The pilot concludes with a binary recommendation to the Board: scale, or stop and restructure. 'Scale' requires that at least four of the six primary metrics move in the intended direction with statistical significance, that the reference case payback period stays within 18 months on realistic assumptions, and that the Independent Evaluator signs the results without material qualification. 'Stop and restructure' is not a failure verdict — it is an honest acknowledgement that the model, the data foundation or the organisation is not yet ready, and it prevents the far larger cost of a premature enterprise roll-out.

7.4 The S&OP Backbone: MAPE-Based Planning Integrated with the Framework

Sales and Operations Planning has been, for four decades, the connective tissue between commercial ambition and operational reality. In most organisations it remains a monthly ritual: a deck reconciled in a conference room, already drifting out of date by the time it is signed. The planning architecture set out below — Monitor, Analyse, Plan, Execute — treats forecast error not as an embarrassment to be buried in a variance report but as the primary signal on which the plan is corrected between cycles.

7.4.1 Why MAPE, and Why Now

Mean Absolute Percentage Error (MAPE) is a well-worn statistic. Its power in the IPC context is not the number itself but the discipline it enforces. When MAPE is measured at SKU-week-channel granularity and published on the same dashboard as inventory health and shelf availability, every planning conversation is anchored in a shared truth: how wrong was the last plan, where, and why. The IPC layer supplies that truth in near-real time; the S&OP process converts it into decisions.

7.4.2 The MAPE-IPC Planning Loop

The loop has four phases. Monitor: the IPC event stream feeds a rolling MAPE calculation at SKU × region × week granularity, alongside sell-in, sell-out, secondary sales velocity and RSAI. Analyse: the system decomposes MAPE into bias (systematic over- or under-forecasting), noise (random error) and shock (single-period breaks). Plan: three planning postures — aggressive, resistive and balanced — are proposed to the S&OP committee, each with its inventory, capacity and working-capital implications quantified. Execute: the approved plan is pushed to ERP and

WMS as production, procurement and deployment orders, and the loop restarts. The cadence is weekly at operational level, monthly at tactical level, and event-driven at exception level.

7.4.3 Three Planning Postures

The IPC-enabled S&OP recognises that there is no single 'right' plan; there is a right posture for the moment. An aggressive plan is warranted when the organisation is preparing for a demand surge that is knowable in advance — a festive season, a new product launch tied to a sporting event, a movie release, a political conclave, a regional festival. Here the plan deliberately over-builds inventory in a controlled corridor, accepts a higher holding cost, and prioritises availability over efficiency. A resistive plan is invoked when the environment turns adverse: a financial slowdown, a natural calamity, an epidemic on the pattern of Covid-19, a sudden change in government rule such as a packaging ban, or an unexpected shift in levy or tax structure. Here the plan pulls back on procurement, protects cash, thins the pipeline and prioritises the fastest-moving 20% of SKUs. A balanced plan is the default — invoked when the sales domain shows no clear directional signal but is disturbed by localised, unforeseen issues (a regional strike, a competitor stock-out, a distributor insolvency). Here the plan holds central inventory steady but redistributes it geographically to absorb the disturbance.

7.4.4 How IPC Senses the Shift

The critical question is how the IPC layer decides which posture to recommend without waiting for a human to notice. Six sensing mechanisms operate in parallel. First, a historic pattern library holds signatures of past aggressive, resistive and balanced episodes — festive lifts, Covid-era demand collapses, GST-transition disruptions, flood-affected quarters — indexed by SKU family, region and calendar week. Second, a present-momentum tracker computes short-window velocity (7-day and 21-day) against the trailing 90-day baseline, flagging deviations beyond two standard deviations. Third, an external-signal ingestor scrapes structured public sources — weather bureaus, commodity indices, official epidemic bulletins, gazette notifications on tax and rule changes, sports and entertainment calendars — and maps them to affected SKU families. Fourth, a sell-in versus sell-out gap monitor detects channel-stuffing or channel-starvation early. Fifth, a complaint and return velocity monitor catches quality or acceptance shocks. Sixth, a distributor-level cash-flow proxy (derived from payment behaviour and order frequency) flags financial stress in the network before it becomes an order collapse.

7.4.5 Signal Fusion and Posture Recommendation

The six sensing streams feed a signal-fusion module that produces a Posture Recommendation Score for each SKU-region cell. The score is a weighted composite whose weights are themselves learned from the historic pattern library — festive signals dominate in August–November for confectionery in India, weather signals dominate for beverages in the pre-monsoon window, regulatory signals dominate for tobacco and packaged food across the year. When the score crosses a configurable threshold, the IPC platform does not auto-execute; it drafts a Posture Change Note for the S&OP committee, showing the triggering signals, the proposed shift (say, from balanced to aggressive on beverage SKUs in three southern states for the next four weeks), the inventory and cash implications, and the reversal criteria. Human judgement remains sovereign; the machine supplies the evidence and the option set.

7.4.6 Worked Illustration — A Festive Aggressive Plan

Consider a snack manufacturer preparing for a regional festival that historically lifts category demand by 35–45% over a four-week window. The IPC layer, six weeks in advance, detects the calendar trigger, retrieves the past three years' festive signatures for the affected states, adjusts for the current baseline momentum (which is already 8% above last year) and recommends an aggressive posture: build finished-goods inventory to 22 days of forward cover (versus the standing 14 days), pre-position 60% of that inventory in the two nearest distribution centres to the festive belt, and place advance procurement orders for high-velocity flavours with a 15% upside option clause. The S&OP committee reviews, adjusts the flavour mix based on qualitative trade-partner feedback, approves. Four weeks later, actual demand comes in at 38% above the base — well inside the planned corridor — and the RSAI holds above 95% through the festive peak while a competitor, planning on last year's number, stocks out by day nine.

7.4.7 Worked Illustration — A Resistive Plan under Sudden Regulation

In the second illustration, a sudden government notification changes the permissible sugar content in a beverage category, with a 90-day compliance window. The IPC's regulatory-signal ingestor flags the notification within hours of gazette publication. The pattern library retrieves the closest analogue — a comparable reformulation event from a previous year — and models the likely three-phase demand response: initial hoarding of the old formulation, a mid-window dip during transition, a slow recovery on the new formulation. The recommended resistive posture stops all speculative procurement of the old formulation, accelerates depletion of existing stock through targeted trade schemes, aligns new-formulation launch to day 75, and preserves approximately 12% of working capital that would otherwise have been trapped in obsolete inventory. Again, the committee decides; the IPC merely makes the invisible visible on the day it becomes visible.

7.4.8 The Balanced Posture — the Quiet Majority of Weeks

Plainly put, most weeks are neither festive nor catastrophic. The balanced posture is where the IPC-enabled S&OP earns its keep unspectacularly — smoothing a distributor's short-term liquidity issue by re-routing rather than re-producing, absorbing a competitor's regional stock-out by nudging shelf presence up two percentage points, letting a slow-moving SKU age out gracefully rather than forcing a promotional dump. Good S&OP, like good governance, is measured largely by the crises that never happened.

7.4.9 Integration with Existing S&OP Cadence

The MAPE-IPC loop does not replace the traditional monthly S&OP cycle; it feeds it. The Pre-S&OP demand review is presented with rolling MAPE, decomposed error, active posture recommendations and the pattern-library evidence behind each. The Supply Review sees the inventory and capacity implications of each posture side by side. The Executive S&OP meeting decides on posture changes explicitly, with reversal criteria written into the minutes. Between cycles, the weekly operating review handles exception-driven posture adjustments under delegated authority.

7.4.10 Guardrails

Three guardrails prevent the MAPE-IPC loop from becoming a runaway optimiser. First, no posture change beyond a defined magnitude (say, a 15% swing in forward cover) may be executed without S&OP committee approval, even if the signal score exceeds the threshold. Second, every posture change carries an automatic 30-day review, with a default reversal to balanced if the triggering signal has decayed. Third, model drift is monitored monthly — if the pattern library's predictive accuracy on the last 90 days falls below a set floor, the library is retrained and the sensitivity of the signal-fusion module is dampened until confidence is restored.

7.5 Extended KANBAN and the IPC Network

The KANBAN system, born on the Toyota shop floor, has spent seventy years teaching manufacturers to pull rather than push. Its extension beyond the factory gate — into the premises of the vendor — is not a new idea in the literature, but it has remained curiously rare in practice. Two things have held it back: the discomfort of trusting a supplier to hold inventory on the buyer's behalf, and the absence of a shared visibility layer that would make such trust verifiable. The IPC Network resolves the second problem, and in doing so it clears the path for the first. This chapter introduces an Extended KANBAN model, describes how it is anchored to the ABC and XYZ classifications and to vendor ranking, and works through examples drawn from FMCG and packaged-food practice.

7.5.0 Position of this section within the literature

Extending pull signals to the supplier's premises is not, in itself, novel, and the section that follows should not be read as claiming otherwise. Vendor-managed inventory, consignment stock and supplier-hub arrangements have all been studied and practised for decades, and the general finding is consistent: the arrangements work where information is shared and demand is reasonably stable, and they transfer cost without transferring benefit where either condition fails. The specific contribution I claim here is narrower and, I hope, more testable than a restatement of that literature.

It is this. Existing vendor-side holding arrangements decide what the supplier holds through negotiation — typically a flat weeks-of-cover figure applied across a contract, adjusted when someone complains. What I propose instead is that the held quantity be a function of three variables that are already computed within the framework: the value class of the item (ABC), the demand-variability class of the item (XYZ), and the supplier's own composite reliability score. The holding percentage is then a rule rather than a bargain, it recalculates as those inputs move, and it can be audited by both parties against the same data. Whether a rule-based holding policy of this kind outperforms a negotiated flat policy on total landed cost and service is a question that admits an experiment, and I set it out as one in Section 8.2 rather than asserting the answer here.

7.5.1 The Idea in Plain Terms

In classical KANBAN, the buyer's own warehouse or line-side rack signals replenishment to the buyer's production planning, which in turn issues a call-off to the supplier. Extended KANBAN pushes the signal one step upstream. The supplier holds a defined quantum of finished, packed, buyer-specified inventory at its own premises, segregated and labelled, but not yet dispatched and not yet invoiced. When the buyer's KANBAN signal fires — triggered automatically by the IPC layer's inventory-consumption event stream — the supplier releases the pre-staged quantum, invoices at that moment, and simultaneously starts replenishing its own held pool. Ownership transfers on release; the buyer's cash moves only when goods move.

7.5.2 Why the IPC Layer Makes It Workable

The historical objection to holding vendor-side buyer inventory has been visibility. Buyers feared they were paying twice — once for the inventory in transit, and again for the risk of stock they could not see. The IPC layer removes that fear. The vendor's held pool is a live entity on the buyer's Inventory Control Centre, with its own Inventory Health Score, expiry clock, batch traceability and photographic proof of segregation. The vendor sees the same screen, filtered to its own scope. Auditability is continuous, not periodic. The commercial contract simply codifies what the shared system already makes transparent.

7.5.3 The Positive Case

Five benefits accrue to the buyer, and — critically — three of them accrue to the vendor as well, which is what makes the arrangement durable. For the buyer: working capital released from the pipeline (typically 15–25% of the affected SKUs' inventory value), warehouse space freed at the buyer's site (often 8–14% of storage footprint for the affected categories), shortened effective lead-time (the supplier's manufacturing time is decoupled from the buyer's replenishment time), reduced obsolescence exposure (the buyer owns less aged stock at any moment), and sharper responsiveness to demand shocks (pre-staged inventory absorbs the first wave without factory disruption). For the vendor: a demand signal of higher quality (real consumption, not forecast), preferential status in the buyer's category (Extended KANBAN partnerships are rarely dual-sourced casually), and improved production smoothing (the vendor plans against a known holding contract rather than volatile call-offs). The negotiation, done properly, is not a zero-sum transfer of pain but a joint reduction of it.

7.5.4 The Honest Downsides

Three cautions deserve mention, though experience suggests each is manageable. First, the vendor's own working capital is temporarily engaged in the held pool; this must be compensated, typically through a modest holding-cost adjustment in the unit price or a shorter payment cycle on release. Second, if the buyer's demand for a covered SKU collapses (through discontinuation, regulatory change or catastrophic loss of market), the vendor is left with pre-staged inventory whose disposition must be pre-agreed in the contract — usually through a buy-out clause with a defined ceiling. Third, quality escapes in a held pool can multiply if not caught early, since larger quantities may be affected before rotation. The IPC layer's batch traceability and expiry prediction significantly narrow this window, and the practical incidence in well-run programmes is modest. On balance, the arithmetic favours the arrangement in most fast-moving categories.

7.5.5 Negotiating and Initiating an Extended KANBAN Agreement

Extended KANBAN cannot be imposed; it must be negotiated. The productive sequence begins with a joint value-mapping workshop between the buyer's supply chain leadership and the vendor's commercial and operations leadership. The workshop quantifies, in shared spreadsheets, the working capital and space savings on the buyer's side and the demand-visibility and revenue-stability gains on the vendor's side. A pilot SKU set is chosen — typically five to ten high-velocity items where both sides can absorb learning cost. The commercial contract addresses seven clauses: the covered SKU list and volumes; the holding quantum and its refresh trigger; segregation and identification standards; release, invoicing and payment mechanics; quality and recall responsibilities; a demand-collapse buy-out ceiling; and an exit protocol with a 90-day notice window. The IPC layer's shared dashboard is written into the contract as the system of record. Once signed, a 90-day soft-launch is run in parallel with conventional replenishment before the switch-over is completed.

7.5.6 ABC and XYZ — The Basis for Deciding What to Hold

Not every SKU belongs in an Extended KANBAN pool. The decision is anchored to the familiar ABC (value contribution) and XYZ (demand variability) classifications, and further tempered by vendor ranking. ABC segments SKUs by their contribution to purchase value: A items are the top ~70% of value in ~10–15% of SKUs, B items the next ~20% in ~20–25%, C items the tail. XYZ segments the same SKUs by demand variability, typically measured as the coefficient of variation of weekly consumption: X items are highly predictable (CV < 0.25), Y items moderately variable (0.25–0.6), Z items erratic (> 0.6). The intersection creates a nine-cell matrix. The vendor's own performance rank — from the Vendor Intelligence Index of Section 4.5 — becomes the third axis, because holding significant buyer inventory at a mediocre vendor is a poor idea regardless of what the SKU matrix says.

7.5.7 A Practical Holding Rule

The recommended rule, refined through pilot experience, holds the following percentages of the buyer's monthly consumption at the vendor's premises. AX (high-value, highly predictable) at a Tier-1 vendor: 25–35%. AY at Tier-1: 15–20%. AZ at Tier-1: 8–12%, and generally not with a Tier-2 vendor at all. BX at Tier-1 or Tier-2: 15–20%. BY at Tier-2: 10–15%. BZ: not recommended for Extended KANBAN
— the erratic demand negates the working-capital benefit. CX: 20–25% but only if consolidation with other C items justifies the logistics. CY and CZ: not recommended. The percentages assume a 7–10 day replenishment cycle from the held pool and monthly re-calibration of quantum against the trailing 90-day consumption.

7.5.8 ABC × XYZ × Vendor Rank — Reference case Table with FMCG and Food Examples

7.5.9 Reference case Case — A Beverage Concentrate Supplier

A carbonated beverage bottler in western India adopted Extended KANBAN with its primary concentrate supplier, an AX item at a Tier-1 vendor. The held quantum was set at 30% of monthly consumption, refreshed on a five-day cycle triggered by IPC consumption events at the bottling line. In the first twelve months, the bottler released roughly ₹4.2 crore of working capital previously locked in incoming-goods inventory, reduced concentrate warehouse space by 11%, and cut concentrate-related line stoppages from an average of 3.2 per month to 0.6 per month. The supplier, in turn, reported a 22% improvement in its own production planning stability because the IPC consumption stream gave it a demand signal it had never previously enjoyed. The arrangement survived the following year's raw-material price shock — the buy-out ceiling was invoked once, at modest cost — and remains in place.

7.5.10 Reference case Case — A Packaged-Food Packaging Supplier

A biscuit manufacturer in southern India extended KANBAN to its printed laminate packaging supplier for five AX and BX flavour variants. The pre-staged quantum sat at 20% of monthly consumption at the supplier's finishing warehouse, segregated by SKU and photographed daily to the IPC platform. During the festive quarter, the arrangement absorbed a 34% surge in demand for the two lead flavours without a single line stoppage; the corresponding non-Extended-KANBAN variant, sourced conventionally from a second vendor, stocked out for four production shifts. The commercial reconciliation at year-end showed a net gain to both parties: the buyer released roughly ₹1.8 crore of working capital and reduced packaging obsolescence write-off by 41%; the vendor gained a longer-term volume commitment and a preferred-supplier designation.

7.5.11 Reference case Case — A Dairy Ingredient Supplier

A confectionery manufacturer extended KANBAN to its milk-solids supplier, an AY item with moderate seasonality tied to festive gifting demand. The held quantum was set at 18%, and the IPC layer's pattern library automatically raised it to 26% for the eight weeks preceding Diwali based on the historic aggressive signature. The supplier's production plan absorbed the ramp gracefully; the buyer avoided the last-minute air-freight bills that had disfigured the previous year's festive P&L. The value of the MAPE-IPC-KANBAN combination showed up here: the planning posture, the visibility layer and the pull mechanism reinforced one another.

7.5.12 Integration with the IPC Network

Extended KANBAN is not a stand-alone technique in the IPC context; it is a specific expression of the model's core philosophy. The Inventory Health Score covers the vendor-held pool. The Vendor Intelligence Index is the gatekeeper for participation. The MAPE-IPC S&OP loop of Section 7.5 modulates the held quantum in response to posture shifts. The event-based tracking of Section 4.6 provides the tamper-evident audit trail. The customer intelligence of Section 4.4 closes the loop by feeding actual shelf-off-take back into the consumption signal that triggers the pull. The elegance of the arrangement is that it uses every layer of the IPC model without introducing a new one.

7.5.13 Cautions on Scale

Extended KANBAN scales sub-linearly. A programme covering fifteen SKUs across three vendors is manageable within existing planning bandwidth; the same programme across eighty SKUs and twelve vendors is a different beast. Organisations that expand too quickly find that the coordination overhead swamps the working-capital benefit. The sensible growth path is category by category, with each new category preceded by a fresh joint value-mapping and a fresh 90-day soft-launch. The IPC platform reduces but does not eliminate the human coordination cost.

7.5.14 Concluding Note on Extended KANBAN

Extended KANBAN, done well, changes something structural in the buyer-vendor relationship. It converts an arm's-length transaction into a shared operational commitment, underwritten by shared data. The IPC Network is what makes the sharing safe. For organisations that have exhausted the easy gains from ERP and WMS optimisation, this is one of the few remaining places where meaningful working capital, real responsiveness and durable supplier partnership can still be found together in the same programme.

The table below distils the holding-percentage guidance into a single view, with representative FMCG and packaged-food examples in the last column. Percentages refer to the buyer's monthly consumption held at the vendor's premises under Extended KANBAN.

ABC × XYZ	Vendor Rank	Recommended Hold in % based on rank	Refresh Cycle	Reference case FMCG / Food Example
AX	Tier-1	25 – 35 %	5 days	Concentrate for carbonated beverages; refined edible oil for a biscuit line
AX	Tier-2	15 – 20 %	7 days	Standard sugar grade for a confectionery plant

ABC × XYZ	Vendor Rank	Recommended Hold in % based on rank	Refresh Cycle	Reference case FMCG / Food Example
AY	Tier-1	15 – 20 %	7 days	Milk solids for festive confectionery; wheat flour for a bakery chain

AY	Tier-2	10 – 15 %	10 days	Standard flavour compounds for a snack range
AZ	Tier-1	8 – 12 %	10 days	Seasonal fruit pulp for a limited-edition beverage
AZ	Tier-2	Not recommended	—	Erratic-demand items with weaker vendor governance
BX	Tier-1	15 – 20 %	7 days	Printed laminate packaging for lead SKUs of a biscuit maker
BX	Tier-2	10 – 15 %	10 days	Secondary corrugated cases for a beverage plant
BY	Tier-1 / Tier-2	10 – 15 %	10 days	Flavour-variant labels for a dairy range
BZ	Any	Not recommended	—	Regional flavour test-launch packaging
CX	Tier-1	20 – 25 % (if consolidated)	14 days	Bundled small-value consumables (tape, strapping) at a co-packer
CY / CZ	Any	Not recommended	—	Low-value, low-predictability items — conventional replenishment

Table 10 — Extended KANBAN holding guidance by ABC × XYZ class and vendor rank. Source: constructed for this study from floor-level operating practice.

8. Conclusion

8.1 Limitations

This illustrated model is a conceptual contribution. It does not present empirical validation from a controlled deployment. The simulated figures are directional. The instruments proposed require calibration to each business. The framework depends on operating discipline as much as on technology; organisations that adopt the technology without the discipline should not expect the impacts described.

8.2 Future Research

- Controlled pilot studies of the framework in FMCG and F&B organisations, with pre- and post-implementation KPI comparison.
- Formal statistical validation of the proposed composites — IHS, CFI, VII — against outcome measures.
- Comparative studies between IPC-augmented environments and matched control environments.
- Extension of the framework to adjacent sectors — pharmaceuticals, consumer electronics, quick-service restaurants — with sector-specific adaptations.

- Deeper research on the human-centred aspects: how operators interact with alert streams, how model explanations affect decisions, how governance evolves.

8.3 Concluding Remarks

The chains that move food and consumer goods do not, on the whole, fail for want of software. They are failing, when they fail, for lack of a coherent operating layer that reads what the software already knows, adds what modern sensing can now inexpensively deliver, and presents to human operators — the people whose judgement is still, and should remain, the final authority — a picture on which they can act before problems become losses.

The IPC Network Model is offered as one such operating layer. It may be refined, criticised, extended or replaced by better proposals. What matters, in my view, is that the discussion moves from software to operating design. The next KFC UK, the next Maggi, the next Abbott, the next Sriracha will not be prevented by a new module in an ERP. They may be moderated by a way of operating that treats inventory, product and customer as continuously observable, digitally represented and analytically linked — and by people whose decisions are supported, not replaced, by the decision layer described in this paper.

The framework is intended to be implemented. Future validation is recommended.

Appendix A — Data Tables

The following tables consolidate the reference case datasets referenced in the main paper. All figures are indicative, based on anonymised FMCG and F&B industry benchmarks and simulated pilot data; they are intended to demonstrate model behaviour and are not attributable to any single organisation.

B.1 Baseline vs IPC-Enabled KPI Comparison (Reference case FMCG Plant)

KPI	Unit	Baseline	Month 3	Month 6	Target	Δ vs Baseline
OTIF	%	82.0	86.4	91.2	≥90.0	+9.2 pp
Forecast WAPE	%	28.0	24.1	19.6	≤20.0	−8.4 pp
Inventory DOH	days	46	41	34	≤36	−12
Waste / obsolescence	%	3.6	2.9	2.1	≤2.2	−1.5 pp
Cost-to-serve	index	100	96	91	≤92	−9
Changeover time	min	74	66	58	≤60	−16
Cash-to-cash cycle	days	61	55	48	≤50	−13
Supplier OTIF (in)	%	78.0	83.5	89.0	≥88.0	+11.0 pp

Table 11 — Baseline versus IPC-enabled KPI comparison (reference case FMCG plant). Source: constructed for this study from floor-level operating practice.

B.2 VII Scoring — Worked Example (Beverage SKU × 5 Suppliers)

Value (V): commercial contribution; Integrity (I): quality, compliance, ethical sourcing; Impact (Im): sustainability & resilience. Scale 1–10; composite = 0.4V + 0.3I + 0.3Im.

Supplier	V	I	Im	Composite	Tier	Action
S-01 (regional)	8.5	9.0	7.5	8.35	Strategic	Long-term contract, joint capacity
S-02 (national)	7.8	8.2	6.8	7.62	Preferred	Multi-year, quarterly review
S-03 (global)	9.0	7.5	6.0	7.65	Preferred	Volume anchor, ESG uplift plan
S-04 (local)	6.5	8.8	8.5	7.79	Preferred	Grow share, sustainability partner
S-05 (spot)	7.0	6.5	5.5	6.40	Transactional	Retain as buffer only

Table 12 — VII scoring worked example (beverage SKU × five suppliers). Source: constructed for this study from floor-level operating practice.

B.3 Deterministic Tracking vs Blockchain — Cost & Performance

Dimension	Blockchain (DLT)	IPC Deterministic Tracking	Delta
Event capture accuracy	99.2%	99.6%	+0.4 pp
Median write latency	1.8 s	0.12 s	−93%
Infra cost / 1M events (USD)	420	165	−61%
Integration effort (person-weeks)	32	11	−66%
Auditability	Cryptographic ledger	Signed hashes + immutable log store	Comparable
Vendor lock-in risk	Medium–High	Low	Lower

Table 13 — Deterministic tracking versus blockchain: cost and performance. Source: constructed for this study from floor-level operating practice.

B.4 Three-Year Financial Scenarios (USD, mid-size FMCG BU)

Scenario	Investment Y0	Net Benefit Y1	Y2	Y3	3-Yr NPV @10%	IRR
Conservative	1.20 M	0.65 M	1.10 M	1.35 M	1.42 M	62%
Base	1.20 M	0.95 M	1.55 M	1.85 M	2.31 M	94%
Ambitious	1.45 M	1.30 M	2.10 M	2.55 M	3.28 M	121%
Stress (−20% benefits)	1.20 M	0.76 M	1.24 M	1.48 M	1.62 M	72%

Table 14 — Three-year financial scenarios (USD, mid-size FMCG business unit). Source: constructed for this study from floor-level operating practice.

B.5 Comparative Framework Positioning

Attribute	SCOR	Lean	TOC	Six Sigma	IPC Network
Primary lens	Process ref.	Waste	Constraints	Variability	Integrated value flow
Decision cadence	Annual	Continuous	Event-driven	Project-based	Rolling IPC cycle
Human-centricity	Low	Medium	Medium	Low	High (Industry 5.0)
Data model	Descriptive	Visual	Throughput	Statistical	Deterministic + VII
Sustainability	Add-on	Implicit	Not core	Not core	Embedded in VII
Tech dependency	Low	Low	Low	Medium	Moderate, ERP-agnostic

Table 15 — Comparative framework positioning. Source: constructed for this study from floor-level operating practice.

Vocabulary and Abbreviations

The following terms and abbreviations are used throughout the main paper and this supplementary volume. Definitions marked (proposed) are original to the IPC Network Model; the remainder follow accepted industry usage per APICS/ASCM, ISO, and CSCMP glossaries.

Abbreviation	Full Term / Definition
APS	Advanced Planning & Scheduling
BOM	Bill of Materials
CI	Confidence Interval
CPFR	Collaborative Planning, Forecasting & Replenishment
CSR	Corporate Social Responsibility
DC	Distribution Centre

DLT	Distributed Ledger Technology (Blockchain)
DMAIC	Define–Measure–Analyse–Improve–Control
DOH	Days of Holding (Inventory Days on Hand)
DTC	Direct-to-Consumer
ERP	Enterprise Resource Planning
ESG	Environmental, Social & Governance
EWMA	Exponentially Weighted Moving Average
FMCG	Fast-Moving Consumer Goods
F&B	Food and Beverage
GT	General Trade
IoT	Internet of Things
IPC	Inventory–Product–Customer Network Model (proposed model)
IRR	Internal Rate of Return
ISO	International Organisation for Standardisation
KPI	Key Performance Indicator
Abbreviation	**Full Term / Definition**
MAPE	Mean Absolute Percentage Error
MDM	Master Data Management
MRO	Maintenance, Repair & Operations
MRP	Material Requirements Planning
NPV	Net Present Value
OEE	Overall Equipment Effectiveness
OTIF	On-Time In-Full
P1	Priority-1 Incident
RACI	Responsible, Accountable, Consulted, Informed
ROI	Return on Investment
S&OE	Sales & Operations Execution
S&OP	Sales & Operations Planning
SC	Supply Chain

SCOR	Supply Chain Operations Reference model
SKU	Stock Keeping Unit
TOC	Theory of Constraints
VII	Value–Integrity–Impact (proposed scoring model)
VMI	Vendor Managed Inventory
WAPE	Weighted Absolute Percentage Error
WMS	Warehouse Management System

Table 16 — Vocabulary and abbreviations used in this model. Source: constructed for this study from floor-level operating practice.

Key conceptual terms

Integrated Planning & Control (IPC): a rolling, cross-functional decision cadence that fuses demand, supply, capacity, finance, and sustainability signals into a single reconciled plan reviewed on a fixed drumbeat.

VII Scoring: a three-axis evaluation of SKUs, suppliers, and network nodes across Value (commercial contribution), Integrity (quality, compliance, ethics), and Impact (sustainability and resilience).

Deterministic Tracking: a lightweight, event-sourced traceability approach using signed hashes and immutable append-only logs, providing auditability comparable to DLT without its cost or latency penalties.

Control Tower (IPC): a role-based operating surface that visualises reconciled plans, exceptions, and VII-weighted trade-offs across the end-to-end network.

Acknowledgements and Source Attribution

This research paper synthesises ideas, data patterns, and industry examples drawn from a broad body of publicly available management literature, professional bodies, and my applied practice in FMCG and Food & Beverage supply chains. I gratefully acknowledge the following sources of inspiration, data reference points, and conceptual grounding. Any residual errors of interpretation remain mine.

D.1 Frameworks and Standards Bodies

Source	Contribution to this frame work illustration
ASCM / APICS (SCOR Digital Standard)	Reference process taxonomy against which IPC is positioned
Lean Enterprise Institute	Waste taxonomy and continuous-improvement grounding
Goldratt Institute (Theory of Constraints)	Constraint logic informing IPC bottleneck orchestration
ISO 9001, ISO 22000, ISO 14001	Quality, food-safety, and environmental control baselines
CSCMP Supply Chain Glossary	Terminology alignment for the vocabulary section
World Economic Forum — Industry 5.0 briefings	Human-centric, resilient, sustainable framing

European Commission — Industry 5.0 policy papers	Definition of Industry 5.0 pillars

Table 17 — Frameworks and standards bodies acknowledged. Source: constructed for this study from floor-level operating practice.

D.2 Industry Benchmarks and Data Reference Points

Source	Nature of Contribution
Gartner Supply Chain Top 25 reports	Benchmark ranges for OTIF, DOH, cash-to-cash
McKinsey & Company — Operations & Supply Chain insights	Cost-to-serve and digital-planning benchmarks
Deloitte — Global CPO / CSCO surveys	Supplier segmentation and resilience data points
Bain & Company — Consumer Products practice	Forecast accuracy and promotional-planning ranges
BCG — Operations practice	Change-management and pilot-scaling patterns
Source	**Nature of Contribution**
Kearney — Assessment of Excellence in Procurement	VII value-axis calibration reference
EY — Supply Chain Reinvention series	Sustainability KPI baselines
Nielsen / Circana (formerly IRI)	FMCG demand-signal patterns
FAO / OECD Food Outlook	F&B commodity and waste benchmarks

Table 18 — Industry benchmarks and data reference points acknowledged. Source: constructed for this study from floor-level operating practice.

D.3 Academic and Practitioner Literature

Source / Work	Influence
Christopher, M. — Logistics & Supply Chain Management	End-to-end value orientation
Chopra, S. & Meindl, P. — Supply Chain Management	Strategic-tactical-operational layering
Simchi-Levi, D. — Operations Rules	Resilience and risk quantification
Sheffi, Y. — The Resilient Enterprise	Disruption response logic
Womack & Jones — Lean Thinking	Value-stream framing
Goldratt, E. — The Goal	Constraint-based decision hierarchy

Deming, W. E. — Out of the Crisis	Statistical control and continuous improvement
Kaplan & Norton — Balanced Scorecard	Multi-axis performance measurement inspiration for VII

D.4 Reference case Cases and Anonymised Practice

Reference case examples in the main paper and in Appendix B are composite representations drawn from anonymised programmes across mid- and large-cap FMCG and Food & Beverage manufacturers in South & Southeast Asia, the Middle East, and Europe. Where a specific pattern (e.g. beverage line changeover, dairy cold-chain OTIF, snacks demand-shaping) is described, it aggregates behaviour observed across multiple organisations rather than any single company. No confidential data has been reproduced.

D.5 Personal Acknowledgements

I thank the supply chain, planning, procurement, manufacturing, and analytics practitioners whose day-to-day rigour shaped the IPC cadence; the operations leaders who challenged early drafts of the VII scoring model; and the academic reviewers who tested its coherence against SCOR, Lean, TOC, and Six Sigma. Particular gratitude is extended to peers in the ASCM, CSCMP, and CILT communities for constructive critique, and to my family for their patience during the writing of this work.

References

The list below contains only works cited in the text. Industry benchmark sources consulted for parameter ranges, but not cited as evidence for any claim, are listed separately in Appendix D.2.

Aung, M. M., & Chang, Y. S. (2014). Traceability in a food supply chain: Safety and quality perspectives. Food Control, 39, 172-184.

Barratt, M., & Oke, A. (2007). Antecedents of supply chain visibility in retail supply chains. Journal of Operations Management, 25(6), 1217–1233.

BBC News. (2018). KFC closures and supply chain disruption. [Public news reporting.]

Cachon, G. P., & Fisher, M. (2000). Supply chain inventory management and the value of shared information. Management Science, 46(8), 1032–1048.

Chai, J., Liu, J. N. K., & Ngai, E. W. T. (2013). Application of decision-making techniques in supplier selection: A systematic review of literature. Expert Systems with Applications, 40(10), 3872-3885.

Chen, F., Drezner, Z., Ryan, J. K., & Simchi-Levi, D. (2000). Quantifying the bullwhip effect in a simple supply chain. Management Science, 46(3), 436–443.

Christopher, M. (2016). Logistics and Supply Chain Management (5th ed.). Pearson Education.

Dickson, G. W. (1966). An analysis of vendor selection systems and decisions. Journal of Purchasing, 2(1), 5-17.

European Commission. (2021). Industry 5.0: Towards a sustainable, human-centric and resilient European industry. Publications Office of the European Union.

Fisher, M. L. (1997). What is the right supply chain for your product? Harvard Business Review, 75(2), 105–116.

Food Safety and Standards Authority of India. (2015). Order regarding recall of Maggi instant noodles. FSSAI, Government of India.

Frank, A. G., Dalenogare, L. S., & Ayala, N. F. (2019). Industry 4.0 technologies: Implementation patterns in manufacturing companies. International Journal of Production Economics, 210, 15-26.

Hopp, W. J., & Spearman, M. L. (2011). Factory Physics (3rd ed.). Waveland Press.

Ivanov, D., & Dolgui, A. (2020). Viability of intertwined supply networks: Extending the supply chain resilience angles towards survivability. International Journal of Production Research, 58(10), 2904–2915.

Jaakkola, E. (2020). Designing conceptual articles: Four approaches. AMS Review, 10, 18-26.

Kagermann, H., Wahlster, W., & Helbig, J. (2013). Recommendations for implementing the strategic initiative Industrie 4.0. Acatech.

Kamble, S. S., Gunasekaran, A., & Sharma, R. (2020). Modeling the blockchain enabled traceability in agriculture supply chain. International Journal of Information Management, 52, 101967.

Lambert, D. M., & Cooper, M. C. (2000). Issues in supply chain management. Industrial Marketing Management, 29(1), 65–83.

Lee, H. L., Padmanabhan, V., & Whang, S. (1997). Information distortion in a supply chain: The bullwhip effect. Management Science, 43(4), 546–558.

McKinsey & Company. (2021). Succeeding in the AI supply-chain revolution. McKinsey Insights.

Melnyk, S. A., Narasimhan, R., & DeCampos, H. A. (2014). Supply chain design: Issues, challenges, frameworks and solutions. International Journal of Production Research, 52(7), 1887–1896.

Meredith, J. (1993). Theory building through conceptual methods. International Journal of Operations & Production Management, 13(5), 3-11.

Ohno, T. (1988). Toyota Production System: Beyond Large-Scale Production. Productivity Press.

Olsen, P., & Borit, M. (2013). How to define traceability. Trends in Food Science & Technology, 29(2), 142-150.

Reuters / Bloomberg. (2022–2023). Public reporting on Huy Fong Sriracha supply shortages.

Reuters. (2018). Coverage of the KFC UK chicken supply disruption, February 2018. [Public news reporting.]

Saberi, S., Kouhizadeh, M., Sarkis, J., & Shen, L. (2019). Blockchain technology and its relationships to sustainable supply chain management. International Journal of Production Research, 57(7), 2117-2135.

Seuring, S., & Muller, M. (2008). From a literature review to a conceptual framework for sustainable supply chain management. Journal of Cleaner Production, 16(15), 1699-1710.

Simchi-Levi, D., Kaminsky, P., & Simchi-Levi, E. (2008). Designing and Managing the Supply Chain (3rd ed.). McGraw-Hill.

Supply Chain Council. (2012). Supply Chain Operations Reference (SCOR) Model, Version 11.0.

Swift, C., Guide, V. D. R., & Muthulingam, S. (2019). Does supply chain visibility affect operating performance? Evidence from conflict minerals disclosures. Journal of Operations Management, 65(5), 406-429.

Tang, C. S. (2006). Perspectives in supply chain risk management. International Journal of Production Economics, 103(2), 451–488.

U.S. Food and Drug Administration. (2022). Abbott Nutrition voluntary recall — public statements and updates.

Wang, Y., Han, J. H., & Beynon-Davies, P. (2019). Understanding blockchain technology for future supply chains: A systematic literature review and research agenda. Supply Chain Management, 24(1), 62-84.

Weber, C. A., Current, J. R., & Benton, W. C. (1991). Vendor selection criteria and methods. European Journal of Operational Research, 50(1), 2-18.

Womack, J. P., & Jones, D. T. (2003). Lean Thinking: Banish Waste and Create Wealth in Your Corporation (2nd ed.). Simon & Schuster.

World Economic Forum. (2020). The Global Risks Report 2020. WEF.

World Economic Forum. (2023). Global Lighthouse Network: Shaping the Next Chapter of the Fourth Industrial Revolution. WEF.

World Health Organization. (2020). COVID-19 and impact on global food supply chains — situation reports.

Xu, X., Lu, Y., Vogel-Heuser, B., & Wang, L. (2021). Industry 4.0 and Industry 5.0: Inception, conception and perception. Journal of Manufacturing Systems, 61, 530-535.

Zhong, R. Y., Xu, X., Klotz, E., & Newman, S. T. (2017). Intelligent manufacturing in the context of Industry 4.0: A review. Engineering, 3(5), 616–630.