Beyond the Buyer:

A Multi-Stakeholder Framework for Evaluating the True Economic Impact of Supply Chain Innovation

A White Paper

Editor, Reusable Packaging News
packagingrevolution.net

May 2026

Abstract

Supply chain innovations that distribute their benefits across multiple actors -- logistics providers, distribution center operators, retailers, maintenance teams, claims departments -- face a structural commercial disadvantage: the person making the purchasing decision typically sees only a fraction of the total value created. This paper argues that this is not a sales problem but a measurement problem, rooted in what economists call supply chain spillovers -- the propagation of economic value to parties outside the original investment decision. Existing frameworks -- Total Cost of Ownership (TCO), Activity-Based Costing (ABC), Cost-Benefit Analysis (CBA), and Life Cycle Assessment (LCA) -- each address adjacent problems but none was designed to attribute economic value across the heterogeneous, multi-stakeholder systems through which supply chain innovations create their effects. A growing empirical literature confirms that such spillovers are real and significant (Isaksson, Simeth & Seifert, 2016; Zheng, Chen & Zhao, 2025; Piermartini & Rubínová, 2014; Geng, Xiang, Zhang & Li, 2025), yet no practical methodology exists for measuring and attributing them at the level of a specific investment decision. This paper proposes the Multi-Stakeholder Benefit Attribution Model (MBAM), a five-step framework for identifying, measuring, and aggregating the full economic impact of supply chain innovations across all actors who experience their effects. The framework distinguishes between Tier A benefits -- measurable using existing data and standard conversion methods -- and Tier B benefits -- economically real but not yet fully amenable to monetary conversion -- positioning every MBAM analysis as a rigorous lower-bound estimate of total economic impact rather than a claim to completeness. We illustrate the framework through the IoT-enabled durable pallet -- a product category whose value proposition exemplifies the attribution problem -- and discuss implications for vendors, buyers, and the commercial structures needed to align costs with benefits. We conclude with a research agenda for empirical validation and industry standardization.

1. Introduction

Over the course of more than a decade editing Reusable Packaging News, I have had the same conversation more times than I can count. A vendor -- usually smart, usually passionate, usually sitting on a product that genuinely works -- describes the frustration of walking out of a procurement meeting empty-handed. The technology performed in the pilot. The case studies were real. The durability data was solid. And yet the deal didn't close, or worse, never got serious consideration in the first place. The reason, almost invariably, was some version of the same problem: the numbers didn't add up for the buyer, even when they clearly added up for the system.

This is not a sales problem. It is a measurement problem -- and underneath the measurement problem lies an incentive architecture problem. Even when the full benefit stream can be measured, the person being asked to approve the investment may not control the budgets where those benefits land. Both problems need to be named clearly before either can be addressed.

Supply chain innovation has long struggled with a structural disadvantage that has nothing to do with the quality of the innovation itself. When a new product or technology creates value that is spread across multiple actors in a logistics network -- across departments, across companies, across the chain -- the person making the purchasing decision typically sees only a fraction of that value. The rest lands somewhere else: in a claims department that didn't ask for the product, in a maintenance budget that benefits from fewer repairs, in a retail operation that experiences less shrink, in a janitorial crew that spends fewer hours sweeping up broken wood. These benefits are real. They are often substantial. But they are invisible to the procurement committee weighing them against a purchase price.

The history of reusable packaging illustrates this problem with particular clarity. When plastic pallets began displacing wooden ones in certain applications, the economic case was genuinely complicated. A plastic pallet cost more -- sometimes significantly more -- than its wooden equivalent. Its advantages were real but distributed: longer service life, fewer replacement cycles, better load stability, reduced product damage, lower waste handling costs, improved hygiene in food applications. A sophisticated buyer who controlled all of those cost lines and measured all of those outcomes could build a compelling total cost of ownership case. But many buyers didn't control all of those lines, didn't measure all of those outcomes, and evaluated the decision on the metrics in front of them. The plastic pallet lost deals it should have won.

Now consider the IoT-enabled pallet -- an asset embedded with sensors that track location, utilization, shock events, temperature, and dwell time in real time. The total value proposition is broader still. The data it generates can reduce asset loss for the pool operator, cut detention charges for the shipper, lower damage claims for the logistics provider, improve forecasting accuracy for the retailer, accelerate exception handling across the network, and provide audit trail documentation that reduces dispute resolution costs for everyone. Each of those benefits is real, measurable in principle, and meaningful in magnitude. And each of them accrues to a different actor, measured by a different team, on a different reporting cycle, in a different budget.

The salesperson is typically talking to one of them.

This progression -- from wooden pallet to plastic pallet to IoT-enabled pallet -- is not merely a product evolution. It is an illustration of a deepening and directional measurement problem. At each step, the innovation creates more total value for the system. And at each step, the fraction of that value visible to the purchasing decision-maker shrinks. This is not random noise in the evaluation process -- it is a systematic bias. Conventional analysis does not simply undercount the value of sophisticated supply chain innovations; it produces a directional error that consistently disadvantages the most advanced solutions relative to simpler alternatives. The more sophisticated the innovation, the more the analysis underestimates its true value. The problem worsens precisely as the innovation improves.

Economists would recognize this pattern immediately. The benefits that fall outside the buyer's analysis are what economists call spillovers -- economic effects that accrue to parties who played no role in the original investment decision. A growing body of empirical research confirms that supply chain spillovers are real, substantial, and measurable (Isaksson, Simeth & Seifert, 2016; Zheng, Chen & Zhao, 2025; Piermartini & Rubínová, 2014; Geng, Xiang, Zhang & Li, 2025). What is missing is not recognition that these effects exist but a practical methodology for attributing them at the level of a specific product or technology investment decision. That is the gap this paper addresses.

The existing literature offers partial remedies. Total Cost of Ownership (TCO) frameworks have significantly improved on purchase-price thinking by expanding the cost horizon for the buying firm. Activity-Based Costing (ABC) has surfaced hidden costs within organizations that conventional accounting obscured. Life Cycle Assessment (LCA) methodology -- well established in the reusable packaging space through rigorous studies commissioned by organizations like IFCO and Euro Pool System and conducted by institutions like the Fraunhofer Institute -- has demonstrated how to measure physical and environmental impacts systematically across a product's full life cycle. Each of these approaches represents genuine progress. None of them, individually or in combination, solves the specific problem of attributing economic value across the heterogeneous, multi-stakeholder systems through which supply chain innovations create their effects.

This paper proposes a framework designed to address that gap directly. We call it the Multi-Stakeholder Benefit Attribution Model (MBAM). It is intended as a practical tool for two audiences: vendors seeking to build a more complete and credible business case for innovations whose value is distributed across a supply chain network, and buyers seeking to evaluate those innovations against a more complete picture of the value they create. The framework draws on elements of TCO, ABC, and LCA methodology, but is purpose-built for a problem none of those frameworks was designed to solve.

We develop the argument in five steps. We first examine how supply chain innovation decisions are typically made today, and why the silo structure of most logistics organizations systematically disadvantages innovations with diffuse benefit streams. We then survey the existing literature, assessing each major framework against the specific requirements of multi-stakeholder benefit attribution. From that survey we derive a formal statement of the gap and a set of requirements any adequate framework must meet. We then present the MBAM framework itself, followed by an illustrative application using the

IoT-enabled pallet as a worked example. We conclude with a discussion of implications for vendors and practitioners, the limitations of the proposed framework, and an agenda for empirical validation.

The goal is not to make the procurement decision easier by making the numbers look better. It is to make the procurement decision better by making the numbers more complete. The measurement gap documented here is not random noise. It is a directional error -- one that systematically disadvantages the most sophisticated supply chain innovations precisely because their value is most broadly distributed. And even with perfect measurement, the problem does not fully resolve itself without commercial structures that realign who pays with who benefits. Both are necessary. This paper addresses the first, and points toward the second.

2. Background and Context

2.1 How Supply Chain Innovation Decisions Get Made

In most organizations, supply chain investment decisions follow a familiar path. A vendor presents a product or technology. A procurement team evaluates it against a set of criteria -- typically dominated by purchase price, payback period, and compatibility with existing systems. A business case is assembled, usually by the vendor, occasionally by an internal champion, and submitted for approval to a committee whose members may have little direct exposure to the operational realities the innovation addresses. If the numbers clear a hurdle rate, the project moves forward. If they don't, it doesn't.

This process is not irrational. It is a reasonable response to the complexity of organizational decision-making, the scarcity of internal analytical resources, and the legitimate need to impose discipline on capital allocation. The problem is not the process itself but the inputs it relies on. When the business case captures only the benefits visible to the buying organization -- and sometimes only the benefits visible to the specific department doing the buying -- the analysis is structurally incomplete. A decision made on incomplete analysis is not a well-informed decision, regardless of how rigorously the available numbers are interrogated.

The incompleteness is not random. It has a shape. Benefits that are direct, immediate, and owned by the budget holder get counted. Benefits that are indirect, delayed, or owned by someone else tend not to. This means the business case systematically overweights certain categories of value -- typically purchase price reduction and easily measurable efficiency gains -- and systematically underweights others -- typically damage reduction, labour savings in adjacent functions, downstream quality improvements, and risk mitigation. For innovations whose value proposition is concentrated in the first category, the process works reasonably well. For innovations whose value is distributed across the second, it produces a distorted picture.

2.2 The Silo Problem

Supply chain management has made significant strides over the past three decades in advocating for integrated, systems-level thinking. The academic literature, the major professional bodies, and the leading consultancies have all, at various points, argued that functional silos are the enemy of supply chain performance. And yet the silo structure of most logistics and distribution organizations remains remarkably persistent in practice. Procurement, warehousing, transportation, quality, maintenance, finance, and customer service typically operate with separate budgets, separate KPI frameworks, separate reporting lines, and separate planning cycles. Collaboration across these functions exists but is often episodic rather than structural.

For the purposes of innovation evaluation, this matters enormously. When a new pallet design reduces product damage in transit, the financial benefit may show up in the claims budget, the quality team's shrink numbers, the customer service team's complaint volume, and the retailer's shelf availability metrics -- simultaneously, and in amounts that none of those teams individually would attribute to a pallet procurement decision. The

procurement team that bought the pallet is unlikely to have visibility into any of those outcomes. They are measured elsewhere, by different people, against different targets, long after the purchasing decision has been made and forgotten.

This structural reality was identified in the academic literature as far back as 1992, when Cavinato described the problem of "functional chimneys" -- organizational structures in which the full cost and value of a product moving through a firm are fragmented across functions in ways that make holistic measurement practically impossible (Cavinato, 1992). That the problem persists more than three decades later, despite significant advances in supply chain analytics and integration, suggests it is not merely a data problem. It is an organizational and methodological one.

KPIs are assigned to functions, not to assets. Budgets are owned by departments, not by systems. The result is that the economic impact of a decision made in one part of the organization routinely lands, uncaptured and unattributed, in another. The problem is compounded when the benefit stream crosses organizational boundaries entirely. A reusable packaging solution that reduces a retailer's waste handling costs generates value that the packaging supplier cannot claim, the logistics provider cannot measure, and the brand owner may not even know about. Each party in the chain experiences a fragment of the total impact. No single party has both the incentive and the standing to aggregate the fragments into a complete picture.

2.3 The KPI Paradox

There is a certain irony in the fact that supply chain management is one of the most heavily measured disciplines in modern business. Organizations track hundreds of KPIs -- on-time delivery, order fill rate, inventory turns, damage rates, dwell time, cost per pallet move, claims frequency, and many more. The data infrastructure in leading logistics operations is sophisticated and growing more so with every passing year. And yet the existence of abundant measurement has not solved the attribution problem. If anything, it has made it more acute.

The reason is that KPIs are designed to measure performance within defined boundaries, not to trace the economic consequences of decisions across boundaries. A damage rate KPI tells you how often product is damaged. It does not tell you whether a change in pallet specification three steps upstream in the supply chain is responsible for an improvement in that rate. A claims frequency metric tells you how often compensation is sought. It does not tell you which asset characteristics are driving the pattern. The measurements exist. The connections between them and their upstream causes are rarely drawn, because drawing them requires both cross-functional data access and analytical intent that most organizations do not routinely apply to procurement decisions.

This creates what might be called the KPI paradox: the more rigorously an organization measures its own performance, the more clearly it can see the effects of supply chain decisions -- but the less likely it is to connect those effects back to their causes, because the measurement architecture was not designed to support that kind of attribution. The data to build a compelling multi-stakeholder business case often already exists within the organizations that would benefit from it. It is simply never assembled for that purpose.

It is worth noting that artificial intelligence tools are beginning to address one dimension of this problem. AI-assisted analytics platforms can surface patterns across fragmented organizational data sets -- connecting a change in pallet specification to downstream damage rates, or linking carrier dwell time improvements to upstream asset characteristics -- in ways that manual analysis rarely achieves at scale. But AI addresses the data assembly challenge, not the framework challenge. A machine learning model trained on single-actor KPI data will surface patterns within the existing measurement architecture, not beyond it. It will find connections that were always there but hidden. It will not redefine the system boundary to include stakeholders whose data was never collected. The KPI paradox requires both better tools for assembling existing data and a framework that specifies which data, from which actors, needs to be assembled in the first place. MBAM is designed to provide the second of those two things. AI is an increasingly capable partner for the first.

2.4 The LCA Parallel: A Partial Analogy

Practitioners in the reusable packaging space will be familiar with Life Cycle Assessment methodology. LCA provides a standardized framework -- governed by ISO 14040 and ISO 14044 -- for evaluating the environmental impacts of a product or system across its full life cycle, from raw material extraction through end of life. Studies commissioned by organizations such as IFCO and Euro Pool System, and conducted by institutions including the Fraunhofer Institute in Germany and Franklin Associates in the United States, have demonstrated that it is possible to measure the environmental impact of packaging decisions systematically and credibly across a product's full life cycle and across multiple organizational actors, as reported in Reusable Packaging News.

The economic problem we are addressing in this paper is structurally similar to the environmental problem that LCA was designed to solve. In both cases, the impacts of a decision are distributed across a system of actors and cannot be fully captured by any single participant. In both cases, conventional single-actor analysis produces a distorted picture. And in both cases, the solution requires a methodology that explicitly defines system boundaries, maps impacts to actors, and aggregates across the system using a common unit of measurement.

The difference -- and it is an important one -- is that physical and environmental flows are in some respects easier to measure than economic value flows. Carbon emissions, material weights, and energy consumption do not change depending on who owns them or what contracts govern their transfer. Economic value is more complex. A cost saving for one party may represent a revenue loss for another. A risk transfer may create value for the buyer while creating a liability for the seller. Benefits that appear in one organization's budget may have been paid for by another's. Any economic equivalent of LCA must grapple with these complications in ways that environmental assessment does not. Ellram's foundational work on Total Cost of Ownership acknowledged precisely this complexity, observing that the "true cost" of a procurement decision extends well beyond purchase price into dimensions that conventional accounting was not designed to capture (Ellram, 1993) -- though as we discuss in the following section, even TCO stops short of the full multi-stakeholder picture.

It is precisely this gap -- between what LCA demonstrated was possible for environmental impacts and what remains undone for economic impacts -- that motivates the framework proposed in this paper.

3. Literature Review: Existing Frameworks and Their Limitations

The challenge of evaluating supply chain innovations whose benefits extend beyond the immediate buyer is not new, and it has not gone entirely unaddressed. Over the past three decades, several frameworks have emerged -- from academic research, from management accounting practice, and from the environmental sciences -- that have meaningfully advanced our ability to measure the full costs and benefits of supply chain decisions. Each represents genuine progress. And each, as we will show, falls short of what is needed to solve the specific problem this paper addresses.

3.1 Total Cost of Ownership: Beyond Purchase Price, Not Beyond the Buyer

The most directly relevant framework in the existing literature is Total Cost of Ownership, or TCO. Developed as a purchasing and procurement tool, TCO represents a significant conceptual advance over purchase-price thinking. Rather than evaluating a procurement decision on the basis of acquisition cost alone, TCO attempts to account for all costs associated with a product or asset over its useful life -- including maintenance, operation, disposal, and the transaction costs of managing the supplier relationship.

The foundational academic work on TCO in a supply chain context was developed by Lisa Ellram through a series of papers in the early and mid-1990s. Ellram defined TCO as a philosophy and methodology designed to understand the "true cost" of a procurement decision, and developed taxonomies of cost categories that procurement teams could use to structure their analysis (Ellram, 1993; Ellram, 1995). Cavinato extended this thinking by arguing that cost analysis needed to move beyond the individual firm to capture interfirm cost advantages -- what he called a "total cost/value model" for supply chain competitiveness (Cavinato, 1992). The Belgian researchers Degraeve and Roodhooft further operationalized TCO by combining it with Activity-Based Costing methodology to create more rigorous supplier selection models (Degraeve & Roodhooft, 1999).

TCO has become genuinely influential. It is embedded in procurement training curricula, referenced in sourcing policy frameworks, and used -- at least nominally -- in many major purchasing decisions. And yet TCO has a structural limitation that its proponents have acknowledged but not fully resolved. It is, at its core, a single-actor framework. The question TCO is designed to answer is: "What does this decision truly cost -- or save -- me?" Even the more ambitious interfirm extensions of TCO are primarily oriented toward improving the buyer's decision-making by giving them a more complete picture of their own cost position. They do not attempt to measure the economic impact of the decision on other actors in the supply chain who are not party to the transaction. Ellram herself noted that TCO has limited application in certain outsourcing contexts, and that integrating logistics costs across the full supply chain remained an unresolved challenge (Ellram & Sifferd, 1998).

3.2 Activity-Based Costing: Surfacing Hidden Costs Within the Firm

Activity-Based Costing emerged from the management accounting literature in the late 1980s as a response to the recognized failure of traditional cost accounting to accurately reflect the true cost of products, services, and processes within organizations. Where conventional accounting allocated overhead costs broadly and often arbitrarily, ABC traced costs to the specific activities that consumed resources -- and then traced those activities to the products or decisions that drove them (Cooper & Kaplan, 1991).

Applied to supply chain management, ABC offered significant analytical power. By identifying the activities associated with receiving, handling, storing, and moving goods -- and assigning costs to those activities based on their actual drivers -- ABC could reveal cost structures that aggregate accounting obscured. Dekker and Van Goor demonstrated this potential in a case study applying ABC to an integrated supply chain, though they noted the need to keep the model simple enough to be practically implementable (Dekker & Van Goor, 2000). The combination of ABC and TCO, as developed by Degraeve and Roodhooft, represented the most rigorous attempt to bridge purchasing decision-making and activity-level cost analysis.

For the problem this paper addresses, however, ABC shares a fundamental limitation with TCO: it is designed to surface costs and benefits within an organization, not across organizations. No single firm has the incentive, the standing, or the data access to run an ABC study that honestly attributes costs and benefits across organizations it doesn't control. This means that while ABC can powerfully illuminate the hidden costs of supply chain decisions within a single firm, it cannot by itself solve the multi-stakeholder attribution problem.

3.3 Cost-Benefit Analysis: The Right Question, the Wrong Scale

Cost-Benefit Analysis has a longer history than either TCO or ABC, rooted in public economics and infrastructure investment appraisal. In its classic form, CBA attempts to identify and quantify all costs and benefits associated with a decision or project -- including those that fall outside the decision-maker's own budget -- and express them in common monetary terms to support comparison and choice. Unlike TCO, CBA is explicitly designed to capture externalities: impacts that fall on parties other than those directly involved in the decision.

In principle, CBA is the closest existing framework to what this paper proposes. Its explicit concern with externalities, its multi-stakeholder orientation, and its requirement to express all impacts in a common unit are exactly the characteristics an economic impact attribution framework needs. The problem is one of scale and application. CBA as practiced in supply chain contexts tends to operate at either the firm level (does this investment pay off for us?) or the macro level (what is the net benefit to society?). There is almost no published work applying rigorous multi-stakeholder CBA to a specific supply chain product or technology innovation -- the level at which the attribution problem actually lives.

3.4 Systems Dynamics: Tracing Propagation, Not Attribution

The systems dynamics tradition, originating with Jay Forrester's Industrial Dynamics (Forrester, 1961) and extensively applied to supply chain problems since the 1990s, offers a different kind of analytical power. Rather than calculating costs and benefits at a point in time, systems dynamics models how changes in one part of a system propagate through the network over time -- capturing feedback loops, delays, and non-linear effects that static analysis misses. The most celebrated application is the analysis of the Bullwhip Effect -- the phenomenon by which small variations in end-consumer demand are amplified as they propagate upstream.

For the problem this paper addresses, systems dynamics offers a genuinely useful perspective: it is explicitly designed to trace how a change in one part of a system creates effects elsewhere in the network. In practice, however, systems dynamics has not been applied to the problem of benefit attribution for supply chain innovations. It has focused overwhelmingly on risk, disruption, and demand signal distortion. Furthermore, systems dynamics models are typically built to understand system behaviour, not to assign economic value to specific actors. They can show that a change in pallet specification reduces damage events downstream. They cannot easily tell you how much of the resulting economic value lands with the logistics provider versus the retailer versus the end customer.

3.5 Supply Chain Spillovers: An Emerging Empirical Literature

Economists would describe the core phenomenon this paper addresses using the concept of spillovers -- benefits or costs that fall on parties outside the original investment decision. A growing body of empirical research has established that supply chain spillovers are real, substantial, and measurable, though it has not yet produced a practical methodology for attributing them at the level of a specific investment decision.

Isaksson, Simeth and Seifert (2016) provided foundational empirical evidence that buyer innovation positively and significantly influences supplier innovation in high-technology supply chains -- and that this effect strengthens with the duration of the buyer-supplier relationship. Their finding establishes that supply chain relationships are genuine conduits for economic value transfer: what one actor does reverberates through the network in measurable ways. Building on this directional insight, Zheng, Chen and Zhao (2025) reversed the lens to examine whether supplier innovation spills over to customer profitability, finding a positive and significant relationship that has become more pronounced in recent years as technology and employee mobility have increased. This is close empirical confirmation of the paper's central claim: value created at one node in a supply chain propagates to other actors who didn't originate it and may not even recognize its source.

At an international scale, Piermartini and Rubínová (2014) demonstrated using industry-level R&D and patent data across 29 countries that knowledge spillovers are stronger the more intense the supply chain linkage between the sender and the recipient -- and crucially, that supply chain participation is a more robust predictor of spillovers than geographic distance alone. This finding has important implications: it suggests that the supply chain relationship itself, independent of proximity, is the active mechanism of value transfer. Piermartini and Rubínová also introduce a useful distinction between what they

term "technology lending" -- conditional and temporary transfers -- and genuine spillovers that permanently improve productive capacity. This distinction maps directly onto the transfer payment problem we address in the MBAM aggregation step.

At a broader system level, Cai, Zhang, Gong, Li and Li (2025) demonstrated that logistics industry clustering generates not only local productivity gains but robust positive spillover effects in adjacent regions, with spillovers dissipating at approximately 350 kilometres. This finding establishes that logistics investments have spatial externalities with quantifiable boundaries -- a systems-level dimension of the argument we develop at the firm and transaction level.

Most directly relevant to the IoT dimension of our analysis, Geng, Xiang, Zhang and Li (2025) demonstrated empirically that digital transformation at one node in a supply chain generates spillover effects on both upstream suppliers and downstream customers, operating through knowledge spillover and resource access mechanisms. An IoT-enabled asset is not merely a more functional physical object -- it is a digital capability introduced into a supply chain network. If digital transformation generates spillovers through supply chain relationships, as Geng et al. demonstrate, then an IoT-enabled asset deployed at one node in a supply chain can be expected to generate measurable economic effects across the network -- effects that a single-actor analysis will systematically miss.

Taken together, this emerging literature establishes three things with reasonable empirical confidence: supply chain spillovers are real; they are directional (flowing both upstream and downstream through supply chain relationships); and they are intensified by digital technology and longer-term partnerships. What the literature does not yet provide is a practical methodology for measuring and attributing these spillovers at the level of a specific product or technology investment decision. That is precisely the gap MBAM is designed to fill.

3.6 Life Cycle Assessment: The Closest Methodological Cousin

Of all the frameworks reviewed here, Life Cycle Assessment is in some respects the most instructive -- not because it solves the problem this paper addresses, but because it demonstrates that the problem is solvable, and provides a partial template for how. LCA, governed by ISO 14040 and ISO 14044, provides a standardized methodology for evaluating the environmental impacts of a product system across its full life cycle and across all actors involved in that life cycle. Its defining characteristics are precisely those that an economic impact attribution framework needs: explicit system boundary definition, comprehensive impact category identification, multi-actor scope, and aggregation to a common unit of measurement.

The limitation of LCA for our purposes is that it was built to track physical and environmental flows, not economic value flows. Environmental flows -- carbon emissions, material weights, energy consumption -- are objective physical quantities that do not change depending on who owns them or what contracts govern their transfer. Economic value does not behave this way. A cost saving for one party may represent a revenue loss for another. These complications require explicit methodological treatment

that LCA does not provide. It is this gap -- between what LCA demonstrated is possible for environmental impacts and what remains methodologically undeveloped for economic impacts -- that most precisely defines the contribution this paper attempts to make.

3.7 Industry Practitioner Tools: The Calculator Problem

Beyond the academic literature, a parallel set of practitioner tools has emerged to help vendors and buyers evaluate the economics of supply chain innovations. Trade associations and industry bodies have developed ROI calculators, business case templates, and cost comparison tools intended to make the economic case for innovative products more accessible to procurement decision-makers.

One of the more comprehensive examples is the Reusable Packaging Association's Technology Value Assessment Calculator, developed to help companies quantify the return on investment from applying tracking and monitoring technologies to reusable transport packaging assets (Reusable Packaging Association, 2023). The tool takes as inputs the number and cost of reusable assets, technology deployment costs, shipment volumes, cycle times, and loss rates -- and produces estimated annual savings across five output categories: technology capital investment, asset loss reduction savings, asset utilization savings, labor efficiency gains, and food waste improvement value. This is a genuinely useful instrument for structuring a business case, and it represents meaningful progress over purely intuitive or anecdotal justification.

Yet the calculator's architecture illustrates precisely the limitation this paper identifies. Every input describes the primary stakeholder's situation. Every output describes the primary stakeholder's savings. The question the tool is designed to answer is: what does this technology investment return to the organization making it? Carrier cost reductions accruing to logistics partners, inventory financing improvements realized by downstream customers, or risk mitigation value captured by insurers do not appear in the output -- not because the tool's designers were unaware of these effects, but because the tool's purpose is to support a single organization's investment decision. The system boundary is drawn at the firm.

The iGPS Total Cost of Business (TCOB) framework represents a meaningful step beyond the RPA calculator and the conventional TCO tradition from which both draw. Where standard TCO focuses on asset acquisition and direct operating costs, TCOB explicitly incorporates downstream operational consequences of pallet selection, including product damage, worker safety incidents, equipment disruption, and retailer chargebacks. iGPS has documented that organizations switching from wood to plastic pallets can realize savings of $1 to $4 per pallet shipped when these broader cost categories are included in the analysis (iGPS, 2026). The framework’s core contribution is its insistence that the relevant unit of analysis is not the pallet itself but the system in which the pallet operates -- a meaningful conceptual advance that begins to surface costs and value that single-actor procurement analysis leaves invisible.

That framing aligns with what a number of industry practitioners acknowledge, at least informally. Conversations with vendors active in reusable packaging and pallet pooling markets reveal that indirect cost effects are widely recognized even where they are not

formally quantified. Several vendors have developed internal rules of thumb to estimate downstream impacts -- compressed claims rates, trailer damage reductions, and automated handling disruptions -- when proposing pallet transitions to prospective customers. These estimates are treated as supporting evidence rather than primary analysis, and they are rarely shared publicly. Their existence nonetheless confirms that practitioners understand value is being created and destroyed across organizational boundaries in ways that standard procurement tools do not capture.

What is missing from all of these approaches is a mechanism for systematically identifying, attributing, and aggregating those cross-boundary effects. The tools that exist are either scoped to a single actor, reliant on informal estimation, or both.

An instructive parallel is emerging in an adjacent domain. Greyparrot's Deepnest platform, recently deployed in a collaboration with consumer goods company Kenvue, uses AI-driven recognition technology installed inside material recovery facilities to analyze how packaging actually performs within real-world recycling streams (LeBlanc, 2026). The system tracks how specific packaging formats, materials, and physical attributes behave during sortation, generating operational data that flows back to packaging designers and brand managers who made the original upstream decisions. In a domain where packaging choices have long been evaluated primarily on design intent rather than downstream recovery performance, Deepnest creates a feedback loop across organizational and operational boundaries that previously did not exist.

The supply chain analogy is direct. Just as packaging decisions create consequences that manifest downstream in recycling infrastructure -- consequences that have historically been invisible to the decision-maker -- pallet selection decisions create consequences that manifest in retailer receiving operations, distribution center throughput, product damage rates, and secondary market value. In both cases, the measurement gap is not a permanent condition of the problem but a function of data infrastructure that did not previously exist at sufficient scale or granularity. The emergence of AI-enabled sensing and tracking in physical operations suggests that the boundary between what is measurable and what is not is shifting, with implications that the framework developed in this paper is designed to anticipate.

This observation is significant: the gap identified in the academic literature has been faithfully reproduced in the practical tools practitioners actually use. The structural assumption embedded in the academic TCO and ABC frameworks -- that the relevant economic boundary of an investment is the organization making it -- has been inherited by the industry tools built on those frameworks. MBAM does not replace tools like the RPA calculator; it provides the analytical structure to extend them, identifying which benefit streams exist beyond the primary stakeholder and connecting them to observable data before the investment decision is made.

3.8 Synthesis: Four Requirements, No Complete Solution

Taken together, the frameworks reviewed here address different aspects of a problem that none of them fully solves. From this survey, we derive four specific requirements that any adequate framework for multi-stakeholder economic benefit attribution must meet:

- Stakeholder mapping -- the framework must identify all actors who experience an economic effect from the innovation, including those with no direct contractual relationship to the buyer or seller
- Benefit stream identification -- for each stakeholder, the framework must provide a structured taxonomy of economic impact categories comprehensive enough to capture indirect and diffuse benefits
- Measurement methodology -- the framework must connect each benefit stream to observable data sources and KPIs, and provide a basis for converting them to a common economic unit
- Aggregation logic -- the framework must provide a principled method for summing economic impacts across actors without double-counting, and must explicitly address the treatment of transfers between parties versus net new value created in the system

No existing framework meets all four requirements simultaneously. The following section proposes one that does. The reason the gap has persisted is worth stating plainly: when the total return to a supply chain system from an innovation exceeds the cost of that innovation, but the private return to the investing firm does not, the investment will not occur -- even though it clearly should. This is not a failure of individual analysis or organizational competence. It is a structural problem with a structural consequence: systematic underinvestment in precisely the innovations that create the most system-wide value. Regulators have recognized a parallel version of this problem in the environmental domain. Extended Producer Responsibility legislation in the United States and the Packaging and Packaging Waste Regulation in Europe are both mechanisms for forcing private sector procurement decisions to account for cost and benefit streams -- waste handling, end-of-life recovery, material externalities -- that conventional analysis ignores. The economic equivalent of that regulatory recognition, applied voluntarily through better analytical frameworks, is what MBAM proposes. Closing the gap requires both a measurement framework and commercial structures that realign who pays with who benefits.

4. The Multi-Stakeholder Benefit Attribution Model (MBAM)

The Multi-Stakeholder Benefit Attribution Model (MBAM) is designed to meet all four requirements simultaneously. It is not a replacement for TCO, ABC, or LCA but a purpose-built complement to them -- one that draws on the best elements of each while addressing the specific problem none of them was designed to solve. MBAM is structured as a five-step process, sequential but iterative, designed to be applied at varying levels of analytical depth depending on the data available and the stakes of the decision.

MBAM is a framework -- a structured methodology for organizing and conducting an analysis -- not a calculator or a model in the computational sense. It tells the analyst what to look for, where to look, and how to aggregate what they find. The numbers must come from the organizations that hold them, which means that the practical application of MBAM requires a degree of cross-organizational collaboration that does not currently exist as a matter of routine. The framework is presented here as a normative ideal: what a complete analysis would look like if the data were available.

Step 1: System Boundary Definition

Every rigorous impact assessment begins with a decision about scope: what is inside the analysis and what is outside it. The system boundary for a multi-stakeholder economic impact assessment must answer three questions.

Which actors are in scope? The starting point is the direct transaction: buyer and seller. From there, the boundary expands to include all actors who interact with the innovation -- physically, financially, or informationally -- during its useful life. The key principle is that the boundary should be drawn around economic effect, not around physical contact. An actor who never touches the pallet but whose claims budget is affected by its performance characteristics is inside the system boundary.

Which time period is in scope? The time horizon must be set to match the expected useful life of the innovation being evaluated. Benefit streams must be mapped across that full horizon rather than truncated at an arbitrary payback period. A durable IoT-enabled pallet with a ten-year service life evaluated on a one-year payback criterion is not analyzed rigorously -- it is analyzed inappropriately.

Which benefit types are in scope? MBAM includes both quantifiable and harder-to-quantify benefit categories but distinguishes between them explicitly, requiring the analyst to flag the confidence level associated with each benefit stream rather than either excluding uncertain benefits or treating them with false precision.

Step 2: Stakeholder Mapping

MBAM uses a three-tier stakeholder taxonomy that distinguishes actors by their relationship to the innovation and to the economic impacts it generates.

Tier 1: Primary Stakeholders are direct parties to the transaction -- typically the buyer and seller. Their costs and benefits are most visible, most measurable, and most likely to

be captured in conventional analysis. They are the starting point of MBAM but not its endpoint.

Tier 2: Secondary Stakeholders are actors who interact directly with the innovation but are not party to the original transaction. They experience real and often significant economic effects -- in handling costs, damage rates, labour efficiency, and space utilization -- but are typically absent from the business case that justified the purchase.

Tier 3: Latent Stakeholders are actors who experience economic effects from the innovation without any direct interaction with it. They are the most commonly overlooked category in conventional analysis and often the most analytically interesting -- claims departments, insurance providers, quality assurance teams, and janitorial and maintenance staff whose labour hours are a direct function of asset durability and failure rates.

The output of Step 2 is a stakeholder map: a visual representation of all actors within the system boundary, organized by tier, with initial annotations identifying the nature of their interaction with the innovation and the categories of economic impact they are likely to experience.

Step 3: Benefit Stream Identification

MBAM organizes benefit streams into six categories, derived from the cost taxonomies of TCO and ABC but extended to cover the full range of impacts relevant to a multi-stakeholder analysis.

Category 1: Asset Cost Effects cover changes in the direct costs of acquiring, maintaining, and disposing of assets -- the domain of conventional TCO analysis, extended to secondary and latent stakeholders who own or manage related assets.

Category 2: Labour Effects cover changes in the human effort required to perform supply chain activities. Labour effects are frequently among the largest and most systematically undercounted benefit streams, including direct handling labour and indirect labour in adjacent functions: the maintenance technician who no longer needs to repair broken equipment, the claims administrator who processes fewer damage incidents, the janitorial worker who spends fewer hours cleaning up debris from failed assets.

Category 3: Damage and Loss Effects cover changes in the frequency and severity of product damage, asset loss, and associated compensation events. IoT-enabled assets are particularly significant in this category: real-time shock, tilt, and location data can not only reduce damage events but provide the evidentiary basis for resolving disputes more efficiently when they do occur.

Category 4: Risk and Insurance Effects cover changes in the risk profile of supply chain operations and associated insurance and contingency costs. These are harder to quantify than direct cost effects but are increasingly tractable as insurers develop more granular pricing models and as IoT data provides the loss history needed to support actuarial analysis. At the extreme tail, product recall risk -- the catastrophic event that supply chain risk managers most fear -- is partially a function of cold chain integrity and handling condition documentation. An IoT-enabled asset that provides a continuous condition

record across the supply chain reduces both the probability of undetected condition failures and the cost of responding to them when they occur. The economic value of tail risk reduction is real and substantial, even when it cannot be expressed as a precise dollar figure.

Category 5: Capital and Inventory Effects cover changes in working capital requirements, inventory levels, and asset utilization rates. A more durable asset with a longer service life reduces replacement capital requirements. An IoT-enabled asset with real-time location data reduces safety stock requirements and improves asset utilization.

Category 6: Revenue and Service Effects cover changes in revenue-generating capacity or customer service performance attributable to the innovation. This is the most difficult category to quantify and the most likely to be contested, but it is often real.

The output of Step 3 is a benefit stream matrix: a structured table mapping each stakeholder to the benefit categories relevant to their position in the system.

Step 4: Measurement Mapping

For each cell in the benefit stream matrix, MBAM requires the analyst to specify four elements.

The measurement indicator: what observable quantity corresponds to this benefit stream?

The data source: where does the data come from, and who owns it? In a multi-stakeholder analysis, the data needed to quantify a benefit stream often sits in an organization other than the one bearing the cost or capturing the benefit. Mapping data sources makes the organizational data-sharing requirements explicit.

The conversion method: how is the indicator converted to a monetary value? Where standard methods exist, MBAM applies them. Where they do not, MBAM requires the analyst to document the conversion assumptions explicitly.

The confidence level: MBAM uses a three-level classification: High (directly observed operational data, standard conversion methods), Medium (estimated indicators or non-standard conversion methods), and Low (modelled or assumed data with significant uncertainty).

The output of Step 4 is a completed measurement map. Gaps -- benefit streams that are real but for which adequate data does not currently exist -- are themselves informative: they identify where investment in measurement infrastructure would most improve the quality of future analyses.

Step 5: Aggregation, Validation, and Sensitivity Analysis

With benefit streams identified, measured, and expressed in monetary terms, the final step is aggregation. This step requires careful treatment of two complications that conventional analysis routinely ignores.

The transfer payment problem. Not all economic benefits captured in a multi-stakeholder analysis represent net new value created in the system. Some represent transfers between parties -- a cost reduction for one actor that corresponds to a revenue reduction for another. MBAM requires the analyst to distinguish between genuine net value creation -- benefits that represent real resource savings or output improvements at the system level -- and transfers, which redistribute existing value between actors without creating new value. Only the former should be included in the system-wide total.

The attribution boundary problem. At some point, the causal chain from the original asset specification decision to a downstream economic effect becomes too attenuated to support credible attribution. MBAM addresses this by requiring the analyst to define an attribution horizon -- a point beyond which downstream effects are acknowledged but not quantified.

With transfers excluded and the attribution horizon defined, the aggregation produces three outputs:

- The System-Wide Economic Impact Estimate -- a single monetary figure representing the total net new value created by the innovation, expressed as a net present value to account for the timing of benefit streams
- The Stakeholder Distribution Map -- a breakdown of the total impact by stakeholder, showing how the value is distributed across actors. An innovation that creates substantial total value but concentrates it in actors other than the buyer faces a commercial challenge that the total figure alone does not reveal
- The Sensitivity Analysis -- a structured assessment of how the system-wide estimate changes under different assumptions about uncertain benefit streams

The Role of IoT in Making MBAM Tractable

The framework proposed here would have been extremely difficult to apply in practice even a decade ago. IoT changes this materially. An IoT-enabled pallet is not merely a more durable or more functional asset -- it is a data-generating platform. The same sensors that track location and utilization for operational purposes generate precisely the data that MBAM requires for benefit attribution. Shock event logs provide the evidentiary basis for damage attribution. Dwell time records quantify detention and asset utilization effects. Location data supports loss rate calculation.

This convergence is not coincidental. The same characteristic that makes IoT-enabled assets more valuable -- their ability to generate information about their own performance and environment -- also makes them more measurable. For the first time, it is becoming practically possible to populate a multi-stakeholder benefit attribution analysis with real observed data rather than estimates and assumptions. The framework proposed here is designed for the analytical environment that IoT is creating.

Framework Summary

Step 1: System Boundary Definition -- define the actor scope, time horizon, and benefit types included in the analysis

Step 2: Stakeholder Mapping -- identify all Primary, Secondary, and Latent stakeholders within the system boundary

Step 3: Benefit Stream Identification -- map each stakeholder to the six benefit categories, producing a benefit stream matrix

Step 4: Measurement Mapping -- connect each benefit stream to observable indicators, data sources, conversion methods, and confidence levels

Step 5: Aggregation, Validation, and Sensitivity Analysis -- sum across stakeholders with explicit treatment of transfers and attribution boundaries, producing a system-wide impact estimate, a stakeholder distribution map, and a sensitivity analysis

5. Illustrative Application: The IoT-Enabled Pallet

The pallet is an almost invisible component of modern supply chains. According to research conducted by Virginia Tech's Department of Sustainable Biomaterials, there are an estimated 3.1 billion pallets in circulation in the United States, and approximately 80% of all products sold domestically are transported on pallets at some point in their journey (Hobbs, Horvath & Gething, 2025). In 2021 alone, the US pallet industry produced an estimated 919 million new pallets and 280 million repaired and remanufactured pallets. The economic flows running through this single asset category are substantial. The measurement gap documented here is not a theoretical curiosity -- it is a real commercial distortion affecting real investment decisions across one of the most economically significant infrastructure categories in the global supply chain.

This section walks MBAM through a realistic but illustrative application: the substitution of a conventional wooden pallet with an IoT-enabled durable pallet in a consumer goods supply chain. The numbers used are order-of-magnitude illustrations drawn from published industry data and practitioner experience rather than a formal empirical study. The goal is to demonstrate the structure of the analysis and the magnitude of the gap between what conventional evaluation captures and what MBAM reveals. It is worth noting at the outset that this application surfaces both positive spillovers -- value created for actors outside the primary transaction -- and a significant negative spillover in the form of disrupted pallet core revenue for distribution center receivers. MBAM is not an advocacy tool for innovation adoption. It is an honest accounting of where value flows, in both directions. That balance is part of what makes its findings credible.

Step 1: System Boundary Definition

Actor scope. The system boundary includes all actors in a representative consumer goods supply chain who interact with or are economically affected by pallet specification: the pallet supplier or pool operator; the brand owner or manufacturer (primary buyer); the contract logistics provider; the distribution center operator; the retailer; the claims and insurance functions; the maintenance and facilities teams at each handling node; and the janitorial and housekeeping teams at each node.

Time horizon. A conventional wooden pallet may have an effective service life of six to twelve months in a high-throughput supply chain. A durable IoT-enabled pallet may carry a warranted service life of eight to ten years or more. Setting the analysis time horizon at ten years is essential to capturing the full economic picture. A single-year payback analysis applied to a ten-year asset is not just incomplete -- it is systematically misleading.

Step 2: Stakeholder Mapping

Tier 1 -- Primary Stakeholders: Brand owner / manufacturer (buyer); Pallet supplier / pool operator (seller)

Tier 2 -- Secondary Stakeholders: Contract logistics provider; Distribution center operator; Retailer receiving and backroom operations

Tier 3 -- Latent Stakeholders: Claims and insurance functions; Maintenance and engineering teams at each node; Janitorial and housekeeping teams at each node; Quality assurance and shrink management teams; Sustainability and ESG reporting functions; IoT data consumers -- forecasting, planning, and logistics optimization teams who benefit from the data stream the asset generates

Step 3: Benefit Stream Identification

Asset Cost Effects

The brand owner captures the most visible benefit: an IoT-enabled durable pallet with a ten-year service life eliminates eight to nine replacement cycles that a wooden pallet program would require over the same period. Durability alone -- before any IoT benefit is considered -- significantly changes the asset economics. The distribution center operator also benefits: durable pallets with consistent dimensions and structural integrity reduce wear on racking systems and handling equipment, a connection that is well understood by facilities engineers who see the damage firsthand but rarely appears in procurement analyses.

Labour Effects

Wooden pallet failures generate splinters, broken boards, protruding nails, and structural collapses that damage handling equipment and require maintenance intervention. A durable pallet eliminates this maintenance demand almost entirely. In high-throughput environments, the accumulated maintenance labour saving over a ten-year asset life is significant. The janitorial dimension is arguably the most purely visible latent benefit stream in the entire analysis: wooden pallets shed debris constantly and fail catastrophically with some regularity, generating cleanup requirements that are a direct and measurable function of pallet volume and failure rate. This cost is tracked in every distribution center -- it simply never appears in a pallet procurement business case.

Damage and Loss Effects

Product damage attributable to pallet failure is a real and measurable cost for any brand operating at scale. Every damage claim generates processing costs -- investigation, documentation, negotiation, settlement -- that are independent of the claim value itself and sit entirely outside the pallet procurement budget. Asset loss -- pallets that go missing and are never recovered -- is a significant and chronically underestimated cost. Industry estimates suggest loss rates for untracked wooden pallets in open pool environments can reach twenty to thirty percent annually. An IoT-enabled pallet with real-time location tracking reduces loss rates dramatically -- and the economic value of that reduction, accumulated over a ten-year asset life and a fleet of meaningful scale, is among the largest single benefit streams in the entire MBAM analysis.

Risk, Capital, and Revenue Effects

A supply chain that generates fewer damage events, fewer asset losses, and better evidentiary documentation of handling conditions is a lower-risk supply chain. Real-time shock and condition data creates an evidentiary record that transforms damage disputes

from contested negotiations into data-supported determinations, reducing dispute resolution costs for all parties. On the capital side, IoT-enabled location tracking directly improves asset utilization by making the entire fleet visible in real time -- reducing the safety stock of spare pallets needed to buffer against loss and idle time. Improved on-shelf availability -- a downstream consequence of reduced damage and more reliable supply chain performance -- has a measurable revenue effect for brands and retailers, though we treat this as a lower-confidence benefit stream given the length of the causal chain.

Step 4: Measurement Mapping (Selected Benefit Streams)

Wooden pallet replacement costs → Asset cost saving for brand owner. Indicator: Replacement cycle frequency and unit cost. Data source: Procurement records. Conversion: Units replaced per year × unit cost × years in horizon. Confidence: High.

Maintenance labour reduction → Labour saving for maintenance teams. Indicator: Maintenance call-outs attributable to pallet failures, time per call-out. Data source: Maintenance management system logs. Conversion: Call-outs eliminated × average duration × loaded labour rate. Confidence: Medium.

Janitorial labour reduction → Labour saving for housekeeping teams. Indicator: Debris cleanup hours per thousand pallet movements. Data source: Facilities management records, time-motion study. Conversion: Hours eliminated per period × loaded labour rate. Confidence: Medium.

Asset loss reduction → Financial saving for pool operator / brand owner. Indicator: Loss rate before and after IoT deployment, fleet size, asset replacement cost. Data source: Asset tracking system reconciliation data. Conversion: Loss rate reduction × fleet size × asset unit cost × years in horizon. Confidence: High.

Damage claims reduction → Cost saving across logistics provider, claims functions. Indicator: Claims frequency and average claim value before and after. Data source: Claims management system records, insurance loss runs. Conversion: Claims eliminated × average total cost per claim. Confidence: High for frequency, Medium for processing cost attribution.

Fleet size reduction through improved utilization → Capital saving for pool operator. Indicator: Asset utilization rate before and after IoT deployment. Data source: Asset tracking system utilization reports. Conversion: Utilization improvement × fleet size × asset unit cost. Confidence: High.

Step 5: Aggregation and Key Findings

Handling the transfer payment problem. The most significant transfer payment to identify correctly is the damage claims reduction. A reduction in claims paid is a cost saving for the party making the claim and a revenue reduction for the insurer. In a system-wide analysis, these offset each other. What remains as genuine net value is the reduction in the resources consumed by the damage event itself -- the destroyed product, the wasted labour, the processing cost.

The system-wide picture. Without access to specific operational data from a real deployment, we cannot produce a validated system-wide figure. What we can observe is that the benefit streams falling outside the primary buyer's conventional analysis -- latent stakeholder labour savings, asset loss reduction, equipment wear reduction, claims processing costs, capital efficiency gains -- are collectively of comparable or greater magnitude than the benefit streams the buyer would typically count. A conventional procurement analysis of this innovation is likely capturing somewhere between thirty and sixty percent of the total economic value it creates.

The Fungible Reusable Asset Complication: GMA Pallets and Pallet Core Revenue

MBAM requires meaningful adaptation when applied to fungible reusable assets -- standardized items like GMA wood pallets and steel drums that circulate through secondary markets rather than being tracked as specific identified assets. The scale of this secondary market is substantial. According to Virginia Tech's pallet market research, the US industry produced approximately 280 million repaired and remanufactured pallets in 2021, of which an estimated 78% of received cores were repaired and sold -- the highest figure recorded in thirty years of survey data (Hobbs, Horvath & Gething, 2025). As reported by the Recycle Record in April 2026, typical core prices paid by recyclers range from approximately $2.00 to $3.00 per unit across US regions, with core availability described as tight to tight-average across virtually every market -- consistent with an active, supply-constrained secondary market, not a residual or informal one.

The economic mechanics of this market create a specific and practically important complication for MBAM. In a typical GMA pallet flow, the shipper purchases new or reconditioned pallets, ships product to the receiver, and after unloading, the receiver sells the used pallet cores to a reconditioner or broker. The receiver captures a per-unit pallet core revenue from that sale -- a meaningful and budgeted income line in many distribution center operations. At $2.00–$3.00 per unit, a large distribution center receiving 50,000 pallet loads per month could be generating $100,000–$150,000 per month in pallet core revenue -- a material budget line that a procurement decision about pallet specification directly affects.

When a shipper substitutes a proprietary IoT-enabled durable pallet -- which they retrieve after delivery rather than leaving with the receiver -- for a GMA wood pallet program, the receiver loses their pallet core revenue stream. This is a real economic cost imposed on a supply chain partner by a procurement decision that partner had no voice in. A complete MBAM analysis must capture this cost -- and a commercially savvy vendor or buyer must address it through pricing adjustments, retrieval logistics that create equivalent value, or some other mechanism.

These complications illustrate a broader principle: MBAM must account not only for new value created by an innovation but for existing value streams disrupted by it. The pallet core revenue stream is a clear example of a latent stakeholder benefit that could be negatively affected by a seemingly straightforward procurement decision. Making that disruption visible is as important as capturing the new value created.

This case also illustrates an important analytical distinction that MBAM is designed to make explicit. The receiver's lost core revenue from a proprietary retrieval program is a genuine negative spillover: a cost that crosses an organizational boundary and lands on a party outside the original transaction who had no voice in the pallet selection decision. By contrast, consider a DC operator who selects plastic pallets for its own downstream order picking and store delivery operations, replacing the GMA wood pallets that would otherwise make that punishing journey. The GMA cores that remain in the upstream flow experience less damage, grade out higher, and command better prices in the secondary market. That incremental improvement in core revenue accrues to the DC itself, the same organization making the plastic pallet selection. It does not cross an organizational boundary. It is not a spillover in the technical sense. It is an internal attribution failure: a real and material benefit landing within the decision-maker's own organization but is almost never connected back to the pallet selection decision that produced it, because no one has structured the business case to ask that question.

MBAM is designed to surface both categories and to distinguish between them. A negative spillover that crosses organizational boundaries requires a commercial response: pricing adjustments, retrieval logistics that create equivalent value, or some other mechanism that acknowledges the cost being imposed on a supply chain partner. An internal attribution failure requires an analytical response: deliberately tracing the benefit back to the decision that created it and including it in the business case on its own merits. Knowing which category you are dealing with is not just analytically important. It determines what you do about it.

6. From Framework to Practice: A Tiered Entry Point

The five-step MBAM process described above represents a normative ideal -- what a complete multi-stakeholder economic impact analysis would look like if all relevant data were available and all relevant actors were willing to share it. Practitioners reading this framework will reasonably ask: is this achievable in practice, or is it an analytical aspiration that bears little relationship to what a vendor or buyer can actually do before their next procurement meeting?

The honest answer is: both, depending on which benefit streams you are trying to measure and which data you already have access to.

MBAM is not an all-or-nothing proposition. It is a structured approach that delivers value at partial application -- and for the IoT-enabled pallet specifically, a meaningful subset of the full benefit stream is measurable today, with existing data, using standard methods, without requiring cross-organizational data sharing agreements or analytical infrastructure that does not yet exist. The remaining benefit streams are real but require either additional data collection effort or methodological development that is part of the research agenda. Knowing which is which is the first practical step.

What You Can Measure Today

The following benefit streams are immediately actionable for any organization deploying or evaluating an IoT-enabled durable pallet. Each relies on data that already exists in standard operational systems and requires only deliberate analysis -- not new data infrastructure.

Asset loss reduction is the most directly measurable benefit stream in the entire MBAM analysis, and often the largest. An IoT-enabled pallet generates its own evidence: before-and-after reconciliation of asset location data produces a documented loss rate that can be converted to a financial saving using asset replacement cost. No data sharing with other organizations is required. No assumptions about downstream effects are needed. This is observed data, generated automatically by the technology the asset already contains. For operators managing fleets of meaningful scale, this single benefit stream frequently justifies a substantial portion of the investment in IoT infrastructure alone.

Replacement cost savings from durability are straightforward arithmetic that any procurement team can perform with their own records. Purchase price of conventional pallets, replacement frequency, number of units in the program, and the service life of the durable alternative are all knowable from existing procurement data. This calculation requires no cross-organizational access and no IoT data -- it is a TCO calculation that most organizations are already equipped to perform but often don't complete over the full asset life horizon.

Fleet utilization improvement is directly observable from IoT tracking data. Dwell time reports, location records, and utilization dashboards produced by any asset tracking platform quantify how much of the fleet is actively in service versus idle, lost, or awaiting retrieval. Improved utilization translates directly to reduced fleet size requirements -- a capital saving that is calculable from utilization data the IoT system already generates.

Damage claims reduction is measurable from data that exists in every organization's claims management system and insurance loss runs. Before-and-after claims frequency and average claim value produce a documented saving that is attributable to the change in asset specification. The main analytical requirement is the deliberate step of connecting claims data to asset specification -- a connection that is almost never drawn in standard reporting but requires no new data collection to make.

Maintenance labour reduction is measurable with modest analytical effort from maintenance management system logs. Call-out records, time-per-incident data, and loaded labour rates are available in any facility that tracks maintenance activity. Attributing a portion of call-outs to pallet-related equipment damage requires review of incident records -- a few hours of analytical work -- but produces a Medium-confidence benefit estimate that meaningfully improves the completeness of the business case.

Janitorial and housekeeping labour reduction is perhaps the most consistently overlooked benefit stream in pallet procurement decisions, and yet it is among the most tractable. Facilities management records track cleaning hours. Time-motion studies -- simple observations of debris cleanup activities across a sample of shifts -- can establish a baseline cost per thousand pallet movements. The before-and-after comparison is straightforward once the baseline is established. This benefit stream requires deliberate measurement intent but no sophisticated data infrastructure.

What Requires Additional Effort

The following benefit streams are real but require either cross-organizational data sharing or analytical steps beyond what most organizations currently perform as a matter of routine.

Downstream retailer shrink reduction requires data from the retail partner -- shrink rates, availability metrics, inbound exception volumes. This data exists but sits outside the buyer's organizational boundary. In collaborative retailer-supplier relationships where vendor-managed inventory or joint business planning processes are already in place, accessing this data is feasible. In arm's-length commercial relationships, it typically is not without a specific analytical initiative.

Racking and equipment wear reduction at secondary nodes requires maintenance data from the distribution center operator or third-party logistics provider. Where the buyer operates their own distribution network, this data is internally accessible. Where third-party logistics providers are involved, data sharing agreements are needed.

Claims processing cost reduction at insurance and claims functions requires data on the administrative cost per claim -- investigation time, documentation effort, settlement overhead -- which is tracked by claims departments but rarely shared with supply chain partners whose decisions affect claim volume. This is achievable in integrated organizations with internal claims functions; it requires negotiated data access in externally insured arrangements.

What Remains Genuinely Difficult

The Tier B benefits described in Section 8.6 -- employee safety effects, organizational learning, relationship quality improvements, industry demonstration effects -- are real but remain outside the scope of what MBAM can currently operationalize with monetary precision. They should be acknowledged in any MBAM analysis as qualitative evidence that the Tier A estimate understates true total impact, but they should not be assigned monetary values without the methodological development that the research agenda identifies as a priority.

The Practical Implication

A practitioner applying MBAM does not need to achieve the full normative ideal to produce analysis that is materially more complete and more defensible than conventional single-actor TCO. Starting with the immediately actionable benefit streams -- asset loss reduction, replacement cost savings, fleet utilization improvement, damage claims reduction, maintenance labour, and janitorial labour -- and expressing each with appropriate confidence levels produces a business case that is demonstrably more complete than what currently reaches procurement committees, without requiring data infrastructure or cross-organizational agreements that do not yet exist.

For an IoT-enabled durable pallet deployed at meaningful scale, the immediately actionable benefit streams alone are likely to represent forty to sixty percent of the total system-wide value MBAM reveals. That is not everything. But it is typically enough to materially change the economics of the procurement decision -- and it is achievable with data that already exists, analysis that requires deliberate intent rather than new capability, and a framework that tells you exactly where to look.

The full MBAM analysis -- including cross-organizational benefit streams, the pallet core revenue complication, and the Tier B benefit acknowledgment -- represents the long-term standard toward which the industry should work. The immediately actionable subset represents where any practitioner can start, today, with what they already have. Early LCA studies did not measure every environmental impact perfectly. They measured what they could with available data and were transparent about limitations. The methodology gained credibility over time as data improved and methods were refined. MBAM is designed to follow the same trajectory -- building analytical confidence incrementally, starting with what is measurable now, and extending as data infrastructure and methodological development catch up with the framework's full ambition.

7. Discussion

7.1 What the Framework Reveals That Conventional Analysis Misses

The illustrative application of MBAM to the IoT-enabled pallet produces a finding that is both analytically significant and practically consequential: the benefit streams that fall outside the primary buyer's conventional analysis are not marginal. They are substantial, real, and in aggregate likely comparable to or greater than the benefit streams the buyer would typically count. A procurement decision made without them is not a slightly incomplete decision. It is a fundamentally distorted one -- and the distortion is not random. It runs in a consistent direction, systematically undervaluing innovations whose benefits are most broadly distributed and overvaluing those whose benefits are concentrated in the buyer's own budget. The more sophisticated and genuinely valuable the innovation, the more pronounced this directional bias tends to be.

When a vendor of an IoT-enabled durable pallet walks out of a procurement meeting empty-handed, the conventional interpretation is that the product failed to make its case. MBAM suggests a different interpretation: the numbers that were presented were an incomplete subset of the right numbers, selected in a way that structurally disadvantaged this category of innovation. In many organizations, the person evaluating the investment is structurally positioned to see only the fraction of value that lands in their budget -- which means they may be incapable of approving an investment that is clearly value-creating at the system level, not because of poor judgment, but because of how their role and their metrics are defined. The product may have made a compelling case for the full supply chain system while making a mediocre case for the specific budget being asked to fund it. These are very different commercial problems with very different solutions.

7.2 Why Would Buyers Care? The Incentive Alignment Problem

The framework proposed in this paper rests on an assumption that deserves direct examination: that making the full benefit stream of a supply chain innovation visible will change the behavior of the organizations evaluating it. This assumption is not self-evidently true, and intellectual honesty requires engaging with it head-on.

Most procurement managers, as currently incentivized, have no structural reason to care about economic benefits that land outside their budget. A procurement function is typically evaluated on cost savings within its own cost center, payback periods on capital it controls, and supplier performance against contracts it manages. Value that accrues to the janitorial department, the claims function, a downstream logistics partner, or a retail customer's shrink budget is, from the procurement manager's professional perspective, not their problem. This is not a failure of individual judgment. It is a rational response to a measurement and incentive architecture that was not designed to reward system-wide thinking.

Recognizing this does not undermine the case for MBAM. It clarifies it. The framework is not primarily a tool for procurement managers. It is a tool for the organizational levels and commercial contexts in which system-wide thinking is both possible and incentivized.

The C-suite sees a different picture. A Chief Financial Officer reviewing a consolidated P&L, or a Chief Supply Chain Officer accountable for end-to-end performance, is not constrained by the boundaries of a single procurement budget. If a MBAM analysis demonstrates that an innovation saves $900,000 in claims processing, $400,000 in maintenance labour, and $300,000 in janitorial hours while costing $600,000 more than the conventional alternative over a ten-year horizon, that executive has both the authority and the incentive to act on that finding.

Vertically integrated organizations capture more of the benefit stream. A retailer who owns their distribution centers, manages their own logistics operations, and self-insures their fleet may capture a substantially larger fraction of the total benefit stream than a fragmented supply chain participant. For those buyers, a MBAM analysis is immediately and directly useful.

Regulatory and ESG pressure is shifting the calculus. Increasingly, large buyers are accountable -- to investors, regulators, and customers -- for impacts that extend beyond their immediate operations. Supply chain sustainability reporting requirements, Scope 3 emissions accounting frameworks, and supplier due diligence regulations are all mechanisms that force buyers to internalize some portion of externalities they would otherwise ignore. Two regulatory developments are particularly relevant. Extended Producer Responsibility laws now enacted or in development across multiple US states are requiring brands and packaging users to account for end-of-life costs that previously fell on municipalities and consumers -- exactly the kind of diffuse external cost that conventional procurement analysis excludes. The European Union's Packaging and Packaging Waste Regulation (PPWR) is imposing similar requirements across the EU market, with significant implications for packaging specification decisions. Both frameworks are, in effect, regulatory mandates to do what MBAM proposes to do analytically: account for economic effects that fall outside the immediate transaction. Companies building compliance infrastructure for EPR and PPWR are already investing in cross-organizational cost tracking and externality attribution -- precisely the data infrastructure that makes MBAM more tractable. The regulatory and analytical imperatives are converging.

The vendor use case may be stronger than the buyer use case. A vendor who understands where their product's value actually lands can use that knowledge strategically -- to engage the right stakeholders in the buying process, to structure conversations at the appropriate organizational level, and to build a business case that speaks to a CFO rather than a procurement committee.

7.3 The Missing Piece: Commercial Structures That Realign Costs With Benefits

Better measurement is a necessary but not sufficient condition for correcting the commercial distortion that MBAM documents. The sufficient condition requires measurement plus commercial structures that redistribute the cost burden of an innovation to better reflect the distribution of its benefits.

Pooling and leasing models -- exemplified by the major pallet pool operators -- are partly a response to the attribution problem, whether or not they have been articulated in those terms. Rather than asking a single buyer to bear the full capital cost of a durable reusable asset whose benefits are distributed across a network of users, pooling spreads the cost across all participants while the pool operator captures and reinvests the system-wide efficiencies.

Outcome-based pricing models tie vendor compensation to measured outcomes rather than asset acquisition. If an IoT-enabled pallet reduces asset loss by a documented percentage, the vendor can price against that outcome. The buyer pays against realized value rather than anticipated value -- and crucially, the payment can be funded from the budget line that captures the saving rather than from procurement capital. IoT makes outcome-based pricing increasingly tractable by generating precisely the data needed to document outcomes continuously and at scale.

Gain sharing agreements explicitly split documented cost savings between buyer and seller, providing a mechanism for vendors to help surface and capture latent stakeholder value that would otherwise remain invisible.

Coalition or consortium purchasing arrangements, in which multiple supply chain participants who collectively capture most of a benefit stream jointly fund an innovation that none of them would fund alone, represent the most structurally complete response to the attribution problem.

To be direct about what MBAM is and is not: it is a measurement framework, not a commercial solution. It tells you where the value lands. What you do with that information is a commercial and organizational challenge that measurement alone cannot solve -- but that cannot be solved without it. Better measurement leads to better commercial structures. Better commercial structures lead to better alignment between who pays and who benefits. Better alignment leads to better investment decisions. And better investment decisions lead, eventually, to the adoption of innovations that genuinely deserve to be adopted.

7.4 Implications for Vendors

For vendors of supply chain innovations with diffuse benefit streams, MBAM offers both a diagnostic tool and a commercial framework. Before building a business case, a vendor should map the full stakeholder landscape, identify the complete benefit stream matrix, and assess honestly where the value actually lands. This exercise will frequently reveal that the primary buyer captures a minority of the total value -- which explains the commercial frustration but also points toward a solution.

Vendors whose innovations create substantial value for secondary and latent stakeholders have several strategic options. They can educate the buyer about downstream value using MBAM as an analytical foundation. They can engage multiple stakeholders in the buying process. They can restructure the commercial proposition through outcome-based pricing, pooling models, or gain sharing. And they can commission third-party validation studies that provide the credibility that vendor-produced analyses cannot. The LCA tradition in reusable packaging offers a direct precedent:

studies conducted by independent institutions were influential precisely because they were rigorous, independently executed, and transparent about methodology.

7.5 Implications for Buyers

For buyers, MBAM offers a structured challenge to the adequacy of the business cases they currently receive and evaluate. The most important question is the stakeholder distribution question: where does the value of this innovation actually land? A product that creates most of its value for the buyer's logistics partners or retail customers is not necessarily a bad investment -- but it may require a different kind of justification, a different funding mechanism, or a different commercial structure.

The second important question is the time horizon question. Buyers who evaluate ten-year assets on one or two year payback criteria are applying appropriate rigor to an inappropriate time frame. The third question is the latent stakeholder question. Every supply chain organization has functions that bear costs driven by upstream decisions they have no voice in. Making those costs visible -- connecting them to their upstream causes -- is an internal analytical exercise that most organizations have never undertaken deliberately. The data to make these connections typically already exists. It simply has never been assembled for this purpose.

7.6 The Measurable and the Intangible: A Two-Tier View of Benefits

A candid assessment of MBAM requires acknowledging something that the framework's five steps do not fully resolve: not all economically real benefits are measurable in the way MBAM's aggregation step requires. The benefit streams identified in Step 3 span a wide spectrum of tractability, and conflating them -- treating a precisely measured labour saving with the same confidence as a loosely estimated brand equity effect -- would undermine the framework's credibility.

It is useful to draw an explicit distinction between two tiers of benefit that any complete economic impact analysis must address, even when only one of them can be fully operationalized.

Tier A: Measurable Benefits are those for which observable indicators exist, data sources can be identified, and conversion to monetary value can be accomplished using standard methods with reasonable confidence. These are the primary domain of MBAM. They include direct cost savings, labour reductions, damage and loss avoidance, asset utilization improvements, and insurance premium effects. A MBAM analysis that captures only Tier A benefits is already a significant improvement over conventional single-actor analysis.

Tier B: Intangible and Hard-to-Measure Benefits are those for which the economic reality is clear but monetary conversion is genuinely difficult, contested, or dependent on assumptions that reasonable analysts would dispute. These benefits are real. Ignoring them produces a systematically low estimate of total economic impact. But including them without appropriate qualification would produce a systematically inflated one. The honest approach is to acknowledge them explicitly, describe their nature and likely direction, and

flag them as a lower-confidence layer of the analysis that requires methodological development beyond what MBAM currently provides.

Within the buying organization, Tier B benefits relevant to supply chain innovation decisions include:

Employee health and safety. Fewer pallet failures mean fewer workplace injuries -- splinter wounds, back injuries from collapsed loads, near-misses with forklift damage. Workers' compensation costs, lost-time incident rates, and safety compliance costs are all affected. These have real monetary dimensions but are rarely connected causally to a pallet specification decision in standard business case analysis.

Decision-making quality. IoT data does not merely track assets -- it improves the quality of supply chain decisions across the organization. Better real-time visibility supports more accurate demand forecasting, faster exception resolution, and more responsive inventory management. The economic value of better decisions is genuine but diffuse, accumulating across hundreds of micro-decisions over time rather than appearing as a discrete line item.

Organizational capability and learning. A company that deploys IoT infrastructure acquires capabilities -- data literacy, sensor management, real-time analytics -- that compound in value over time and extend to other applications beyond the original deployment. This capability building is a genuine economic benefit of the investment, but it is essentially impossible to attribute to a single procurement decision with precision.

Brand and reputational effects. Fewer damage incidents improve on-shelf availability, reduce consumer-facing quality failures, and contribute to sustainability reporting metrics. These have brand equity dimensions that affect long-term revenue and customer loyalty. The causal chain from pallet specification to brand equity is long, but it is not imaginary.

Across organizational boundaries, Tier B benefits include effects that accrue to the broader supply chain ecosystem:

Industry demonstration effects. An early adopter of IoT pallets or durable packaging bears the full cost of deployment and learning. The industry benefits as case studies, failure modes, integration patterns, and best practices become available to subsequent adopters at lower cost. The investing firm pays; the ecosystem captures a share of the value. This is a genuine externality that falls outside even a generous system boundary definition.

Relationship quality improvements. Better data transparency across supply chain partners -- real-time condition records, location accuracy, utilization visibility -- reduces dispute frequency, builds trust, and makes commercial relationships more resilient. These relationship quality effects have real economic value that manifests in lower transaction costs, more collaborative problem-solving, and reduced commercial friction over time. They resist clean monetary conversion but are experienced concretely by anyone who has managed a supply chain partnership before and after a significant improvement in data sharing.

Workforce capability spillovers. Employees who develop skills operating IoT systems, interpreting sensor data, or managing reusable asset programs carry those capabilities when they move between employers. The investing firm pays for capability development; the industry benefits from the resulting talent pool. This is a supply chain analog of the well-documented human capital externality in economics.

The practical implication of this two-tier structure is important for how MBAM analyses should be presented. A well-conducted MBAM analysis should clearly distinguish between Tier A findings -- expressed as monetary estimates with confidence levels as described in Step 4 -- and Tier B observations -- described qualitatively, with the direction of effect noted but no spurious monetary conversion attempted. This structure allows the analysis to be comprehensive without being misleading, and provides an honest answer to the question that any sophisticated evaluator will ask: what are you not measuring, and why?

The answer MBAM provides is clear: we are not measuring benefits for which credible monetary conversion is not currently possible. But we are naming them, describing their nature, and flagging them as evidence that the Tier A estimate is a lower bound on the true total economic impact -- not a complete picture of it. That is a more defensible and more useful position than either ignoring intangible benefits or inflating the analysis with poorly supported estimates. It is also a more honest one.

7.7 Limitations of the Proposed Framework

The data access problem. MBAM requires data from multiple organizations, some of which are in commercial relationships that create incentives to withhold rather than share information. IoT mitigates this problem for the data streams generated by the asset itself, but it does not solve it for the broader operational data required to fully populate a benefit stream matrix.

The counterfactual problem. MBAM calculates the economic impact of an innovation by comparing outcomes with and without it. In practice, constructing this counterfactual is rarely straightforward. Supply chains are dynamic systems in which many variables change simultaneously, making it difficult to isolate the effect of a single asset specification change from confounding factors.

The organizational incentive problem. Even where the data exists and the counterfactual can be constructed, the organizational incentives to conduct a MBAM analysis are not always present. Without a mechanism to share the costs of analysis and the benefits of the resulting insight, the analytical investment may be difficult to justify from any single actor's perspective.

The framework's current status as a normative proposal. MBAM is proposed here as a theoretical framework, not a validated methodology. It has not been applied to a real supply chain deployment with empirical data, subjected to peer review, or tested against alternative approaches. These are the next steps in its development, and until they are completed, the framework should be treated as a structured starting point rather than a proven tool.

8. Research Agenda

Priority 1: Empirical Validation in Real Deployments

The most urgent research need is to apply MBAM to real supply chain innovation deployments with real operational data, across real organizations, and with findings that can be subjected to peer review and independent replication. The ideal validation study would select a supply chain innovation with well-documented performance characteristics -- an IoT-enabled durable pallet program operated across a network of shippers, logistics providers, and retailers would be a natural candidate -- and apply the full five-step MBAM process with data collected directly from each stakeholder tier.

Such a study would require a trusted third-party research institution to serve as data custodian and analytical intermediary. Virginia Tech's Center for Sustainable Biomaterials, which has produced the most comprehensive empirical research on the US pallet market including the landmark survey by Hobbs, Horvath and Gething (2025), would be a natural institutional home for such a validation study.

Priority 2: Methodology Refinement for Fungible Reusable Assets

A dedicated research stream is needed to develop the framework's treatment of fungible reusable assets -- GMA wood pallets, steel drums, and comparable standardized containers that circulate through secondary markets. Key questions requiring methodological development include secondary market boundary definition, pallet core revenue attribution, and grade deterioration modeling.

The scale of the secondary market makes this a priority. At $2.00–$3.00 per unit as reported by the Recycle Record in April 2026, and with approximately 280 million repaired and remanufactured pallets processed annually in the US (Hobbs, Horvath & Gething, 2025), the aggregate pallet core revenue flowing through distribution center budgets annually represents a material income stream that procurement decisions about pallet specification directly affect, and that almost never appears in a conventional business case.

Priority 3: Data Infrastructure Development

MBAM's practical tractability depends on data that currently exists in fragmented form across multiple organizations. Several specific infrastructure gaps require attention: cross-organizational data sharing protocols governed by appropriate standards bodies; IoT data standardization across platforms and vendors to enable meaningful cross-deployment comparison; and KPI linkage research that empirically establishes the relationships between supply chain asset characteristics and the performance metrics of downstream stakeholders.

Priority 4: Industry Standardization and Tool Development

Achieving standardization requires engagement from industry bodies with the standing and reach to drive it. The Reusable Packaging Association is well positioned to develop

MBAM-aligned evaluation guidelines for its member community. MHI has existing relationships with technology vendors whose products most directly enable the data collection MBAM requires. GS1 is the natural lead for cross-organizational data sharing protocols and IoT data standardization. Academic institutions with strong supply chain research programs -- including Virginia Tech, MIT's Center for Transportation and Logistics, Penn State's Smeal College of Business, and Cranfield University in the UK -- could provide the independent research capacity needed to validate and refine the methodology.

Priority 5: Extension to Other Innovation Types

The framework's logic applies to a much broader range of supply chain innovation types beyond physical assets. Supply chain software platforms, cold chain monitoring systems, sustainable packaging innovations -- in each of these cases, MBAM's core logic applies: define the system boundary, map the stakeholders, identify the benefit streams, measure what is measurable, and aggregate with principled treatment of transfers and attribution boundaries. The development of context-specific MBAM extensions for each of these innovation types represents a substantial and valuable research program.

Priority 6: Methodological Development for Tier B Benefits

The two-tier benefit structure described in the Discussion section represents a genuine methodological frontier. Tier A benefits -- those amenable to monetary conversion using standard methods -- are the primary domain of MBAM as currently proposed. Tier B benefits -- employee safety, organizational learning, relationship quality, industry demonstration effects, workforce capability spillovers -- are real but resist the kind of clean quantification that procurement committees typically require.

Developing credible approaches to Tier B benefit estimation is a significant research opportunity. Several methodological traditions offer partial templates. Health economics has developed sophisticated approaches to monetizing safety and mortality risk reduction -- these could be adapted for supply chain workplace safety applications. Behavioral economics and organizational learning research offer frameworks for valuing decision quality improvements and capability building. Network economics has tools for valuing relationship quality and trust effects in commercial partnerships.

None of these traditions was developed with supply chain innovation evaluation in mind, but each offers methodological building blocks that could be adapted and combined. A research program that systematically develops Tier B measurement approaches -- and validates them through empirical application -- would substantially extend the completeness of what MBAM can capture. Until that work is done, the honest position is that every MBAM analysis produces a lower-bound estimate of total economic impact, with the true figure higher by an amount that is real but not yet fully measurable.

9. Conclusion

This paper began with a frustration that anyone who has spent time in the reusable packaging or supply chain technology space will recognize immediately: innovations that create genuine, substantial, and demonstrable value repeatedly fail to clear the commercial hurdles that stand between a promising product and widespread adoption. The numbers, as presented in conventional procurement analyses, simply don't add up -- even when, in some broader and more complete sense, they clearly do.

The argument developed here is that this frustration is not a market anomaly or a failure of salesmanship. It is the predictable consequence of a structural measurement problem that economists would call a market failure: the frameworks that supply chain organizations use to evaluate innovations were not designed to capture value that is distributed across multiple actors, multiple functions, and multiple organizational boundaries. When the total return to a supply chain system from an innovation exceeds the cost of that innovation, but the private return to the investing firm does not, the investment will not occur -- even though it should.

The pallet spectrum illustrates this with unusual clarity. A wooden pallet, a durable plastic pallet, and an IoT-enabled durable pallet represent three points on a continuum of increasing sophistication, increasing total system value, and increasing diffusion of that value across stakeholders who are not party to the purchasing decision. The janitorial team sweeping up broken wood, the claims administrator processing damage disputes, the distribution center manager budgeting pallet core revenues, the logistics provider bearing liability for damage events caused by asset failures upstream -- none of these actors appears in a standard pallet procurement business case. And yet each of them experiences real economic consequences from the specification decision being made without their input or awareness. A growing body of empirical research confirms that these spillover effects are real and measurable (Isaksson et al., 2016; Zheng et al., 2025; Piermartini & Rubínová, 2014; Geng et al., 2025) -- but no practical methodology has existed for attributing them at the level of a specific investment decision. MBAM proposes to fill that gap.

The Multi-Stakeholder Benefit Attribution Model proposed here is designed to make those consequences visible, measurable, and aggregable. It is a five-step framework -- system boundary definition, stakeholder mapping, benefit stream identification, measurement mapping, and aggregation with principled treatment of transfers and attribution boundaries -- that draws on the best elements of Total Cost of Ownership, Activity-Based Costing, Cost-Benefit Analysis, and Life Cycle Assessment methodology while addressing the specific problem none of those frameworks was built to solve.

The framework distinguishes explicitly between two tiers of benefit. Tier A benefits -- direct cost savings, labour reductions, damage and loss avoidance, asset utilization improvements, risk reduction -- are measurable using existing data and standard methods, and are the primary domain of MBAM. Tier B benefits -- employee safety effects, organizational learning, decision quality improvements, relationship quality, industry demonstration effects -- are economically real but resist clean monetary conversion with current methods. MBAM captures Tier A rigorously and acknowledges

Tier B honestly, positioning every analysis as a lower-bound estimate of true total impact rather than a claim to completeness. This is the methodologically defensible position, and it is a considerably stronger one than the partial-capture approach that conventional analysis takes without acknowledging that it is partial.

The scale of the opportunity is not trivial. The US pallet market alone involves 3.1 billion units in circulation, with 919 million new pallets and 280 million repaired and remanufactured units produced annually (Hobbs, Horvath & Gething, 2025). The secondary market for pallet cores is active, regionally priced, and supply-constrained, with typical core values of $2.00 to $3.00 per unit across US regions as of April 2026 (Recycle Record, 2026). The economic flows running through this single asset category -- flows that are partially visible to some actors and invisible to others -- are substantial.

Beyond pallets, the same logic applies to every supply chain innovation whose value is distributed across a network of actors. In each case, the innovation is being evaluated on a partial benefit stream against a full cost. In each case, the measurement gap systematically disadvantages the most sophisticated and most genuinely valuable innovations -- precisely because their value is most broadly distributed.

Closing that gap requires better frameworks, better data infrastructure, better organizational incentives, and better commercial structures. This paper has attempted to contribute the first of those four elements. The remaining three are the work of the research agenda, the industry bodies, and the practitioners who recognize the problem from their own daily experience and are ready to do something about it.

The goal, in the end, is not to make supply chain innovations look more valuable than they are. It is to make them look as valuable as they actually are. That distinction matters -- analytically, commercially, and for the long-term health of a supply chain ecosystem that depends on the right innovations getting funded, adopted, and scaled.

References

Cavinato, J.L. (1992). A total cost/value model for supply chain competitiveness. Journal of Business Logistics, 13(2), 285–301.

Cai, Y., Zhang, Y., Gong, Y., Li, W., & Li, F. (2025). The impact of logistics industry clustering on green total factor productivity: Evidence from China. Sustainability, 17, 7978. https://doi.org/10.3390/su17177978

Cooper, R. & Kaplan, R.S. (1991). Profit priorities from activity-based costing. Harvard Business Review, 69(3), 130–135.

Dekker, H.C. & Van Goor, A.R. (2000). Supply chain management and management accounting: A case study of activity-based costing. International Journal of Logistics Research and Applications, 3(1), 27–43.

Degraeve, Z. & Roodhooft, F. (1999). Effectively selecting suppliers using total cost of ownership. Journal of Supply Chain Management, 35(1), 5–10.

Ellram, L.M. (1993). A framework for total cost of ownership. International Journal of Logistics Management, 4(2), 49–60.

Ellram, L.M. (1995). Total cost of ownership: An analysis approach for purchasing. International Journal of Physical Distribution & Logistics Management, 25(8), 4–23.

Ellram, L.M. & Sifferd, S.P. (1998). Total cost of ownership: A key concept in strategic cost management decisions. Journal of Business Logistics, 19(1), 55–84.

Forrester, J.W. (1961). Industrial dynamics. MIT Press, Cambridge, MA.

Geng, Y., Xiang, X., Zhang, G., & Li, X. (2025). Digital transformation along the supply chain: Spillover effects from vertical partnerships. Journal of Business Research, 189, 114842. https://doi.org/10.1016/j.jbusres.2024.114842

Hobbs, S., Horvath, L., & Gething, B. (2025). Investigation of the status of the wooden pallet market during the COVID-19 pandemic. BioResources, 20(2), 3047–3074. https://doi.org/10.15376/biores.20.2.3047-3074

ISO 14040 (2006). Environmental management -- Life cycle assessment -- Principles and framework. International Organization for Standardization, Geneva.

ISO 14044 (2006). Environmental management -- Life cycle assessment -- Requirements and guidelines. International Organization for Standardization, Geneva.

iGPS (2026). The true cost of a broken pallet. Retrieved from https://igps.net/the-true-cost-of-a-broken-pallet/

Isaksson, O.H.D., Simeth, M., & Seifert, R.W. (2016). Knowledge spillovers in the supply chain: Evidence from the high tech sectors. Research Policy, 45(3), 699–706. https://doi.org/10.1016/j.respol.2015.12.007

LeBlanc, R. (2026, May 20). Greyparrot and Kenvue push circular packaging beyond theory and into measurable recovery performance. Reusable Packaging News / packagingrevolution.net.

Piermartini, R., & Rubínová, S. (2014). Knowledge spillovers through international supply chains. WTO Staff Working Paper No. ERSD-2014-11. World Trade Organization. https://doi.org/10.30875/921eddee-en

Recycle Record (2026, April). United States mixed species recycled pallets -- regional market report. Pallet Profile. Retrieved from palletprofile.com

Zheng, L., Brooks, L., Chen, Y., & Zhao, Y. (2025). Do the benefits of innovations spill over from suppliers to customers? Accounting and Finance Research, 14(3), 1–17. https://doi.org/10.5430/afr.v14n3p1