

# Internet of Products:

## Toward a Product Identity Layer of the Digital Economy

*A Theory-Building Conceptual Paper*

---

Minh Do

Information Technology Institute  
Vietnam National University, Hanoi  
[minhdo@vnu.edu.vn](mailto:minhdo@vnu.edu.vn)

---

**Keywords:** Internet of Products • Product Identity Layer • Digital Infrastructure • 4E Model • Supply Chain Transparency • Tokenization • Real-World Assets • Transaction Cost Theory • Information Asymmetry • Digital Economy • Digital Twin

## Abstract

In three decades, the internet has undergone successive waves of digitalisation that have progressively transformed the global economy. The first wave, the *Internet of Information*, digitised content and data, enabling near-zero-cost global circulation of information. The second wave, the *Internet of Value*, digitised financial assets, enabling money and ownership rights to flow natively through networked systems via blockchain and cryptographic protocols. Each wave extended the scope of the digital economy by conferring persistent, verifiable, network-accessible digital identity on a new class of entities.

Despite these transformations, a large and economically significant domain remains comparatively under-digitised: the physical economy of manufactured goods. Each year, trillions of dollars of products move through global supply chains, yet once they leave manufacturing environments, their individual identities dissolve into fragmented logistics records. This absence imposes structural inefficiencies rooted in information asymmetry (Akerlof, 1970) and high transaction costs (Coase, 1937) that pervade global product markets, sustaining USD 467 billion in annual counterfeit trade (OECD, 2023) and inhibiting circular economy transitions.

This paper proposes the **Internet of Products (IoP)** as the third stage of internet infrastructure evolution: a digital infrastructure layer in which every physical product possesses a persistent, verifiable, and network-accessible digital identity. The IoP framework operates at the *infrastructure level*, theorising the generative conditions under which product identity systems of any kind can become open, evolving platforms, rather than at the application level of specific compliance systems such as Digital Product Passports (DPP) or logistics traceability tools.

Drawing on transaction cost theory (Coase, 1937; Williamson, 1981), information asymmetry theory (Akerlof, 1970), and digital infrastructure theory (Hanseth and Lyytinen, 2010; Tilson et al., 2010), the paper develops a conceptual framework describing how product identity functions as a new infrastructure layer connecting the physical economy, digital platforms, and financial systems. The paper introduces the **4E Model of Product Digitisation**, Encode, Enrich, Exchange, and Execute, with explicit boundary conditions and failure-mode analysis. Six research propositions are developed; five are refined into testable hypotheses with operationalised constructs and identified data sources, while a sixth proposition is positioned as a priority item in the empirical research agenda pending development of cross-industry measurement infrastructure. Critical counter-arguments, including the oracle problem, governance failure risks as illustrated by TradeLens, and privacy-transparency tensions, are systematically examined. An illustrative case analysis across three deployed systems grounds the framework in observable practice, with explicit proposition-case mapping.

**Theoretical contributions** are threefold: (1) positioning IoP as a new digital

infrastructure layer that is categorically distinct from prior phases due to the physical-digital linkage challenge; (2) the 4E Model as a staged framework for product digitisation with specified boundary conditions and failure modes; and (3) integration of previously fragmented research streams, IoT, supply chain, smart products, DPP, digital twins, and tokenisation, under a unified infrastructure-theoretic lens that produces cross-stream predictions not derivable from any constituent theory alone. The paper concludes with a structured research agenda targeting empirical validation.

## Contents

---

<b>1</b>	<b>Introduction</b>	<b>5</b>
<b>2</b>	<b>Research Design and Methodological Approach</b>	<b>7</b>
<b>3</b>	<b>Literature Review</b>	<b>7</b>
3.1	The Evolution of Digital Infrastructure: Two Precedents and a Missing Third	7
3.2	Internet of Things and Connected Physical Objects . . . . .	9
3.3	Digital Supply Chains and Traceability . . . . .	10
3.4	Smart Products, Digital Twins, and Product-Service Systems . . . . .	10
3.5	Digital Product Passports and Sustainability Infrastructure . . . . .	10
3.6	Tokenisation and Real-World Assets . . . . .	11
3.7	Prior Uses of the Term ‘Internet of Products’ . . . . .	12
3.8	Theoretical Foundations and the Categorical Distinctiveness of Physical Identity . . . . .	12
3.9	Research Gap . . . . .	13
<b>4</b>	<b>Conceptual Framework: The Internet of Products</b>	<b>14</b>
4.1	From Anonymous Objects to Identifiable Entities: An Ontological Shift . .	14
4.2	Defining the Internet of Products . . . . .	14
4.3	The Product Identity Layer . . . . .	17
4.4	Architecture of the Internet of Products . . . . .	18
4.5	The 4E Model of Product Digitisation . . . . .	19
4.5.1	Boundary Conditions of the 4E Model . . . . .	21
4.5.2	Failure Modes of the 4E Model . . . . .	21
4.6	Bridging the Three Economies . . . . .	23
<b>5</b>	<b>Critical Analysis: Counter-Arguments and Limitations</b>	<b>23</b>
5.1	The Oracle Problem: The Constitutive Limitation of Physical Identity Infrastructure . . . . .	23
5.2	Governance Failure: The TradeLens Cautionary Tale . . . . .	24
5.3	Privacy-Transparency Tension . . . . .	26
5.4	Scalability and Economic Viability for Low-Value Products . . . . .	26
<b>6</b>	<b>Illustrative Case Analysis</b>	<b>27</b>
6.1	EU Falsified Medicines Directive (FMD) Pharmaceutical Serialisation: Mandated E1–E2 at Scale . . . . .	27
6.2	Parmigiano Reggiano Microchip Authentication (2022): E1 with Immediate E3 Value . . . . .	29

6.3	Homie Washing Machine-as-a-Service (Netherlands): E1–E4 as Theoretical Illustration . . . . .	29
<b>7</b>	<b>Research Propositions and Testable Hypotheses</b>	<b>30</b>
<b>8</b>	<b>Discussion: Theoretical Contributions and Implications</b>	<b>33</b>
8.1	Theoretical Contributions . . . . .	33
8.2	Implications for Research . . . . .	35
8.3	Implications for Practice . . . . .	36
<b>9</b>	<b>Research Agenda for Future Studies</b>	<b>38</b>
9.1	Empirical Validation of the 4E Model . . . . .	38
9.2	Oracle Problem and Trust Architecture Research . . . . .	38
9.3	Governance Models and Adoption Dynamics . . . . .	38
9.4	Financial Integration and Cost of Capital . . . . .	39
9.5	Product Data Ecosystems and Generativity: Developing Measurement Infrastructure for H4 . . . . .	39
9.6	Sustainability and Circular Economy Outcomes . . . . .	39
9.7	AI Agents and Agentic Commerce . . . . .	39
<b>10</b>	<b>Conclusion</b>	<b>40</b>

## Introduction

---

Over the past three decades, digital technologies have fundamentally reshaped the global economy through successive expansions of what entities can possess identity within networked systems. Each such expansion has followed a recognisable pattern: a new class of entities acquires persistent, verifiable, network-accessible digital identity; identity infrastructure enables generative applications that designers cannot foresee; and the resulting economic transformation is far larger than any initial application suggested.

The first expansion, widely described as the *Internet of Information*, established digital identity for informational objects through the Domain Name System (DNS) and Uniform Resource Locators (URLs) (Castells, 2010). Domain names transformed internet addresses from routing codes into economic entities: ‘amazon.com’ is not merely a technical address but a valuable asset with history, brand equity, and market value. The DNS infrastructure was not designed to enable e-commerce; it was designed to organise information. Yet its *generativity*, the capacity of an open infrastructure to produce unanticipated innovations (Zittrain, 2006), made e-commerce possible as an emergent application.

The second expansion, the *Internet of Value*, established digital identity for financial assets through blockchain wallet addresses, token protocols, and distributed ledger systems (Nakamoto, 2008; Tapscott and Tapscott, 2016). Wallet addresses transformed cryptographic keys into economic entities: an Ethereum address holds assets, executes contracts, and participates in decentralised finance protocols as an identifiable actor. Again, the infrastructure was generative: smart contracts, decentralised finance, and tokenised real-world assets emerged as applications of identity infrastructure that Bitcoin’s designers did not anticipate.

Despite these profound developments, a large segment of the global economy remains comparatively under-digitised at the level of individual entities: the physical economy of manufactured goods. Each year, trillions of dollars of products move through complex global supply chains. Yet once products leave manufacturing environments, their individual identities often dissolve into fragmented logistics records and organisation-specific databases. A batch of pharmaceuticals, a luxury handbag, or an EV battery becomes, from the perspective of most digital systems, just another anonymous item in a shipment.

This anonymity imposes significant economic costs. Drawing on Akerlof (1970), markets for physical goods are characterised by buyers’ inability to verify product provenance, authenticity, and condition, precisely the conditions that generate counterfeit markets, unsafe products, and under-investment in quality. Globally, counterfeit goods account for approximately 3.3% of world trade, representing USD 467 billion annually (OECD, 2023). Applying Coase (1937)’s transaction cost lens, anonymous products generate high costs of

search, verification, and contract enforcement across supply chains, costs that persist because no infrastructure exists to make product histories credibly observable. These are not failures of individual firms; they are structural consequences of the absence of product identity infrastructure.

Recent regulatory and technological developments suggest this condition may be changing. The EU's Ecodesign for Sustainable Products Regulation mandates Digital Product Passports for most physical goods by 2030. Major financial institutions including BlackRock and JPMorgan have launched tokenised real-world asset products. GS1 and W3C have developed interoperable standards for item-level product identification. Together, these developments suggest the emergence of a new infrastructure phase: the **Internet of Products**.

This paper proposes the Internet of Products (IoP) as the third stage of internet infrastructure evolution, in which each physical product is assigned a persistent digital identity allowing it to accumulate lifecycle data, interact with digital platforms, and potentially participate in financial systems. The concept is grounded in digital infrastructure theory (Hanseth and Lyytinen, 2010; Tilson et al., 2010).

A critical distinction is necessary at the outset. Recent scholarship has made important advances in designing Digital Product Passport systems for specific regulatory contexts, for example, Heeß et al. (2024)'s conceptualisation of DPPs for hydrogen supply chains in this journal. That work addresses a vital application-level question: how should a product data-sharing instrument be designed to meet specific stakeholder and regulatory requirements? The IoP framework addresses a different, more foundational question: what infrastructure conditions must hold for product identity systems of *any* kind, including DPPs, but also tokenisation platforms, servitisation models, and circular economy marketplaces, to function as generative economic infrastructure rather than as compliance instruments? In Hanseth and Lyytinen (2010)'s terms, DPPs are applications; IoP theorises the infrastructure layer that makes such applications possible and that enables unanticipated innovations beyond them.

The paper makes three theoretical contributions. First, it defines the IoP and its **Product Identity Layer** as a new infrastructure layer of the digital economy, one that is categorically distinct from prior phases because physical products, unlike information or financial assets, can be damaged, moved, or substituted independently of their digital representation, creating the *oracle problem* as a constitutive design challenge. Second, it proposes the **4E Model**, Encode, Enrich, Exchange, Execute, with explicit boundary conditions and failure modes. Third, it integrates previously fragmented research streams under a unified infrastructure-theoretic lens and proposes testable hypotheses to guide empirical investigation.

## Research Design and Methodological Approach

---

This paper adopts a **theory-building conceptual approach**, positioned within the interpretive tradition of information systems research (Gregor, 2006; Weick, 1995). The study does not present primary empirical data; instead, it develops a novel theoretical framework by synthesising and extending existing theoretical traditions, transaction cost theory, information asymmetry theory, and digital infrastructure theory, and applying them to a phenomenon that prior literature has not theorised at the infrastructure level.

This methodological choice is consistent with established IS research practice. As Gregor (2006) argues, theory-building papers that *analyse and explain* phenomena by specifying constructs, propositions, and scope conditions represent a distinct and valued contribution type, independent of empirical testing. The present paper corresponds to Gregor's Type II theory (theory for explanation): it identifies the mechanisms through which product identity creates economic value, specifies the conditions under which these mechanisms operate (boundary conditions), and identifies the points at which they fail (failure modes). The six research propositions and associated hypotheses are the primary output of this theory-building process, designed to guide subsequent empirical investigation rather than to be tested within the paper itself.

The illustrative case analysis in Section 5 serves an *abductive* function: it grounds abstract theoretical claims in observable patterns, identifies anomalies that refine boundary conditions, and generates theoretical insights that purely deductive reasoning would not produce (Dubois and Gadde, 2002). Three cases were selected to provide non-redundant theoretical coverage across 4E stages, governance models, and industry contexts, not as a representative sample for generalisation. This distinction between *analytical generalisation* (from cases to theory) and *statistical generalisation* (from sample to population) is fundamental to the methodological logic of the paper (Yin, 2018).

The epistemological stance is *critical realist*: the paper assumes that product identity infrastructure has generative mechanisms that can be theorised prior to observation, while acknowledging that their empirical manifestations are contextually contingent and subject to falsification. This stance is consistent with the broader IS tradition of developing mid-range theories that are abstract enough to travel across contexts but specific enough to be tested (Merton, 1968).

## Literature Review

---

### The Evolution of Digital Infrastructure: Two Precedents and a Missing Third

The development of the internet can be understood as a progressive expansion of the range of entities that can possess identity and act within digital networked environments.

Digital infrastructure theory, as developed by [Hanseth and Lyytinen \(2010\)](#) and [Tilson et al. \(2010\)](#), provides the primary theoretical lens for this paper. Hanseth and Lyytinen define information infrastructures as shared, open, heterogeneous, and evolving socio-technical systems characterised by path dependency, network effects, and unbounded scale. Three properties are particularly relevant to IoP analysis.

**First, generativity** ([Zittrain, 2006](#)): the capacity of an infrastructure to produce unanticipated innovations from diverse actors without central permission. DNS was designed to resolve domain names; its generativity enabled the Web, e-commerce, and cloud computing as unanticipated applications. Bitcoin was designed for peer-to-peer electronic cash; its generativity enabled smart contracts, decentralised finance, and tokenised assets.

**Second, the bootstrap problem:** new infrastructures must create value for early adopters before achieving network scale, but network value depends on scale. Identity data for one product has little value; identity data for a billion products creates a transformative infrastructure. This creates a chicken-and-egg adoption challenge requiring regulatory mandates, market power of large buyers, or natural experiments to overcome.

**Third, the installed base problem:** new infrastructures must evolve from existing systems rather than replacing them. IoP must build on GS1 barcodes, RFID networks, and existing supply chain IT, not replace them.

Two additional theoretical lenses sharpen the IoP analysis. [Rochet and Tirole \(2003\)](#)'s platform economics predicts that product identity infrastructure will exhibit two-sided market dynamics: the infrastructure creates value only when both product data *providers* (manufacturers, logistics actors, certifying regulators) and product data *consumers* (buyers, financial institutions, circular economy actors, AI agents) participate simultaneously. [Zuboff \(2019\)](#)'s analysis of surveillance capitalism provides the counterpoint: the same product identity infrastructure that reduces information asymmetry for buyers also creates behavioural data at product-level granularity that platform operators can extract and monetise without product owners' awareness, what might be called 'product behavioural surplus'. This concern directly motivates the governance typology in Proposition 6 and the privacy-transparency analysis in Section 5.3.

Against this theoretical backdrop, the Internet of Information phase established digital identity for informational objects, enabling what [Shapiro and Varian \(1999\)](#) describe as the economics of information goods: near-zero marginal cost of reproduction, strong network effects, and winner-takes-most dynamics. The Internet of Value phase established digital identity for financial assets, enabling programmable financial systems with atomic settlement and automated enforcement ([Tapscott and Tapscott, 2016](#); [Buterin, 2013](#)). A third stage, product identity infrastructure, would extend the logic of these precedents to

physical goods, but faces a categorically different design challenge discussed in Section 3.8.

### Internet of Things and Connected Physical Objects

The Internet of Things (IoT) literature, originating with Ashton (2009)'s coinage, describes ecosystems in which physical devices are connected to the internet through sensors and communication technologies. Foundational surveys document the technical architecture and application landscape of IoT (Atzori et al., 2010). Research has documented IoT applications in manufacturing, logistics, and supply chains (Gubbi et al., 2013; Lee and Lee, 2015), including RFID-based product tracking (Zhou et al., 2020).

However, IoT research primarily focuses on device connectivity and operational data collection rather than the economic identity of products themselves. Physical objects in IoT systems are treated as data-generating nodes, *objects* being monitored, rather than as economic *entities* with persistent identities capable of acting within market systems (Xu et al., 2014). This distinction, between object (tracked, passive) and entity (identified, active), is the core conceptual gap that IoP addresses. IoT answers the question 'How can a physical device transmit operational data?'; IoP answers 'How can a physical product participate as an identifiable economic actor across its full lifecycle?'

A related clarification is required with respect to two adjacent research streams that reviewers may reasonably expect IoP to address: the *Industrial Internet of Things* (IIoT) and *Industry 4.0*. IIoT focuses on operational connectivity *within* manufacturing environments, sensor networks, machine-to-machine communication, and real-time process monitoring in factories and industrial facilities (Lee and Lee, 2015). Its unit of analysis is the production system, and its primary value proposition is operational efficiency and predictive maintenance within organisational boundaries. Industry 4.0 broadens this to the automation of production processes through cyber-physical systems, advanced robotics, and AI-driven manufacturing (Ivanov et al., 2019). Its concern is the *transformation of production*, not the post-production economic life of the product.

IoP is conceptually downstream from both: it theorises what happens to products *after* they leave the manufacturing environment and enter complex, multi-actor economic systems. The IoP's unit of analysis is the *individual product entity* across its full lifecycle, manufacturing, logistics, retail, use, secondary markets, and end-of-life, rather than the production process or the connected machine. This distinction means that IIoT and Industry 4.0 can *contribute to* IoP (by generating high-quality E2 data at the point of manufacture) but are not substitutes for it: they do not address the cross-organisational identity persistence, verifiability, or financial integration that define the Product Identity Layer.

## Digital Supply Chains and Traceability

Research on digital supply chain management has examined RFID, blockchain, and cloud-based systems as mechanisms for improving supply chain transparency (Christopher, 2016; Ivanov et al., 2019). Traceability systems have been studied extensively in pharmaceuticals and food (OECD, 2021). Blockchain has been proposed for tamper-resistant supply chain records (Sabeti et al., 2019; Treiblmaier, 2018), though its limitations for organisational trust, particularly the gap between technical immutability and organisational data accuracy, have also been identified (Iansiti and Lakhani, 2017).

While these systems enhance supply chain visibility, most solutions remain organisation-centric. Products rarely possess persistent digital identities that exist independently across organisational boundaries (Nguyen et al., 2018). The Walmart Food Trust blockchain, which reduced produce traceability time from seven days to 2.2 seconds (IBM, 2018), illustrates both the potential and the governance challenge: such systems require a powerful orchestrating actor to mandate participation.

## Smart Products, Digital Twins, and Product-Service Systems

Porter and Heppelmann (2014, 2015) demonstrate that smart connected products enable new forms of value creation by enabling continuous interaction between manufacturers and products after sale. The Rolls-Royce Power by the Hour model, Signify's Light-as-a-Service, and Michelin's tyre management services represent large-scale instantiations of product-service systems (PSS) in which manufacturers shift from one-time transactions to ongoing service provision (Baines et al., 2007; Vandermerwe and Rada, 1988).

A closely related body of research concerns digital twins: computational representations of physical products that mirror their real-world state through continuous data feeds (Grieves and Vickers, 2017). Digital twin research has advanced significantly, demonstrating how product-level data enables real-time simulation, predictive maintenance, and lifecycle optimisation (Tao et al., 2019). Critically for the IoP framework, digital twins presuppose a persistent product identity to anchor the data accumulation that makes simulation possible, they are, in effect, a rich instantiation of the E2 (Enrich) and E3 (Exchange) stages of the 4E Model. However, digital twin research focuses on computational and operational capabilities, not on the infrastructure conditions that make such data trustworthy, portable, and economically generative across organisational boundaries.

## Digital Product Passports and Sustainability Infrastructure

The EU's Ecodesign for Sustainable Products Regulation (ESPR, EU 2024/1781) mandates Digital Product Passports (DPP) for most physical goods sold in EU markets,

with implementation timelines from battery passports in 2027 through textiles, electronics, and construction materials to 2030. DPP systems record detailed lifecycle information including materials, environmental footprint, repairability, and recyclability (European Commission, 2022; WEF, 2023).

DPP systems represent the most significant regulatory instantiation of product identity infrastructure to date. Recent contributions have made important advances: Heeß et al. (2024), in this journal, develop design principles for DPPs in hydrogen supply chains, demonstrating how decentralised data processing, privacy-preserving architecture, and cross-organisational interoperability can address information asymmetry. Their work reveals critical design requirements that align closely with the IoP Trust and Data Layers discussed in Section 4.4.

However, as currently designed, DPP systems function primarily as application-layer compliance instruments, providing specified data fields to regulators and consumers within defined regulatory frameworks, rather than as generative economic infrastructure. They answer a specific question: ‘Does this product meet specified sustainability standards?’ The IoP framework addresses a different, foundational question: ‘Under what infrastructure conditions can product identity systems become open, generative platforms that enable unanticipated economic applications, including but extending far beyond DPP compliance?’ In Hanseth and Lyytinen (2010)’s infrastructure theory terms, DPPs are a valuable application layer; the IoP framework theorises the infrastructure layer beneath them. This is not a critique of DPP research, it is a recognition that the infrastructure-level question requires a different theoretical treatment.

### Tokenisation and Real-World Assets

A growing body of research examines the tokenisation of real-world assets (RWA) using blockchain technologies. Tokenisation converts ownership rights over physical or financial assets into digital tokens on distributed ledgers, enabling programmable ownership transfers, fractional ownership, and integration into decentralised financial systems (Yermack, 2017; Casino et al., 2019). The Bank for International Settlements (BIS, 2023) characterises tokenisation as a potentially transformative development for global securities markets. BCG and ADDX (2022) project that tokenised illiquid assets could represent USD 16 trillion by 2030.

Major financial institutions have moved from analysis to deployment: BlackRock’s BUIDL fund, JPMorgan’s Onyx platform, and Franklin Templeton’s tokenised money market fund represent institutional validation. However, each faces a common challenge: reliable linkage between the digital token and the underlying physical asset across the asset’s lifecycle. Without trusted identity mechanisms for physical goods, the credibility of tokenised asset systems remains structurally limited. This ‘oracle problem’ represents the

critical missing link between the IoT/supply chain literature and the tokenisation literature.

### Prior Uses of the Term ‘Internet of Products’

The term Internet of Products has appeared in earlier discourse. [Sarma et al. \(2000\)](#) at the MIT Auto-ID Center envisaged a communications network for tracking physical goods through RFID-based inventory visibility. [Neumann \(2013\)](#) subsequently developed the concept in *The Internet of Products: An Approach to Establishing Total Transparency in Electronic Markets*. Industry publications such as [Qliktag \(2018\)](#) have used the term for digital representations of consumer products accessible through QR codes.

While these prior contributions are valuable, they primarily conceptualise IoP as a product information or logistics tracking system, an application-level advancement in supply chain IT. None addresses product identity as generative digital infrastructure in [Hanseth and Lyytinen \(2010\)](#)’s sense: open, evolving, shared, and capable of enabling unanticipated innovations beyond initial purposes. None integrates transaction cost and information asymmetry theory as the economic mechanisms through which product identity creates value. And none addresses the physical-digital linkage challenge, the oracle problem, as a constitutive design problem. These are the theoretical gaps that the present study addresses.

### Theoretical Foundations and the Categorical Distinctiveness of Physical Identity

Three theoretical traditions provide the foundational grounding for the IoP framework, and together they generate a prediction that prior treatments have not made explicit: physical product identity poses a categorically harder infrastructure design problem than information or financial identity.

**Transaction Cost Theory** ([Coase, 1937](#); [Williamson, 1981](#)). Coase’s foundational insight is that the structure of economic organisation is determined by transaction costs: the costs of searching for information, negotiating contracts, and enforcing agreements. [Williamson \(1981\)](#) extends this by identifying asset specificity, uncertainty, and transaction frequency as determinants of governance form. IoP directly addresses three categories of transaction costs pervading physical goods markets: *search costs* (who made this product, from what materials, under what conditions?); *verification costs* (is this product genuine, compliant, in the claimed condition?); and *enforcement costs* (can warranty and compliance obligations be automatically executed?). By making product histories credibly observable, IoP reduces information asymmetry and lowers transaction costs, potentially enabling market governance to replace organisational governance in contexts where high transaction costs currently force vertical integration.

**Information Asymmetry Theory** (Akerlof, 1970). Akerlof’s ‘market for lemons’ demonstrates that information asymmetry between buyers and sellers can cause market failure through adverse selection: low-quality goods drive out high-quality ones as buyers discount all goods by the average probability of receiving a ‘lemon’. In physical goods markets, buyers cannot observe provenance, manufacturing conditions, or usage history, generating counterfeit markets, unsafe products, and under-investment in quality. The 4E Model’s Execute stage, where smart contracts automate enforcement of verifiable product claims, represents the mechanism through which information asymmetry reduction translates into economic value.

**Digital Infrastructure Theory** (Hanseth and Lyytinen, 2010; Tilson et al., 2010; Zittrain, 2006). As discussed in Section 2.1, infrastructure theory generates predictions about generativity, bootstrap dynamics, and governance challenges that domain-specific framings cannot produce. Applied to IoP, it predicts: (a) that early design decisions about openness and standards will have long-lasting, path-dependent consequences; (b) that the bootstrap problem will be acute; and (c) that generative product identity infrastructure will enable applications that current designers cannot anticipate.

**The categorical distinctiveness of physical identity.** Both DNS (for information) and blockchain wallet addresses (for financial value) assign digital identity to inherently *intangible* entities, entities whose existence is, by definition, confined to digital systems. A domain name cannot be physically stolen; an Ethereum wallet address cannot be physically damaged. Physical product identity poses a qualitatively different challenge: it requires maintaining a credible, verifiable link between a digital identity and a physical object that can be damaged, moved, stolen, or substituted independently of its digital representation. This physical-digital linkage problem, which manifests as the oracle problem in blockchain-based systems, has no parallel in prior internet infrastructure phases. Recognising this distinction is the primary theoretical contribution that distinguishes the IoP framework from both prior IoP conceptualisations and from DPP research.

## Research Gap

Taken together, existing research streams reveal an important conceptual and theoretical gap. IoT research treats products as data sources rather than economic entities. Supply chain research provides logistics visibility within organisation-centric data silos. Smart product and digital twin research demonstrates the value of product-level data without addressing the infrastructure conditions for cross-organisational trust and generativity. DPP systems create valuable compliance instruments without theorising generative product identity infrastructure. Tokenisation research addresses financial representations without the trusted physical-digital identity linkage that makes such representations

credible.

More fundamentally, no existing framework addresses how individual products might function as persistent digital entities, possessing durable, generative identities that allow them to interact across digital platforms, financial infrastructures, and regulatory systems throughout their entire lifecycle. This gap is not merely descriptive but *ontological*: existing systems treat products as objects (passive, tracked, managed), whereas the Internet of Products proposes reconceptualising them as entities (active, identifiable, participating). This ontological shift has significant implications for economic organisation, market design, and infrastructure governance.

## Conceptual Framework: The Internet of Products

---

### From Anonymous Objects to Identifiable Entities: An Ontological Shift

In the traditional industrial economy, physical products exist as anonymous objects once they leave manufacturing facilities. They are identified by category codes, SKUs, barcodes, rather than individual identities. Any two units of the same model are functionally indistinguishable in digital systems, even if they have different production histories, material compositions, or operational experiences.

The Internet of Products proposes a fundamental ontological shift: from treating products as *objects* being tracked within systems to treating products as *entities* possessing independent digital identities within networked economic systems. An entity, unlike an object, can accumulate a history, establish relationships, and participate in transactions as a recognisable party. An entity can carry forward its lifecycle narrative, serve as a credible basis for contractual claims, and interact with financial systems as an identifiable asset.

This ontological reframing parallels transformations already achieved in prior infrastructure phases. DNS transformed routing addresses into economic entities: ‘amazon.com’ is a valuable asset with history, relationships, and market value. Blockchain wallet addresses transformed cryptographic keys into economic actors: an Ethereum address holds assets, executes contracts, and participates in DeFi protocols as an identifiable actor. The Internet of Products proposes an analogous transformation for physical goods, but, as argued in Section 3.8, the physical-digital linkage challenge makes this transformation categorically harder to achieve and constitutes the primary design problem that product identity infrastructure must solve.

### Defining the Internet of Products

The Internet of Products is defined as a digital infrastructure layer in which physical goods are assigned persistent digital identities that allow them to exist as identifiable entities within networked economic systems. Through these identities, products

accumulate lifecycle data, interact with digital platforms, and participate in financial systems throughout their lifecycle.

This definition implies three necessary conditions:

1. **Persistence**, product identity continues to exist independently of any single organisation's database and survives organisational changes, mergers, and bankruptcies;
2. **Verifiability**, the claimed identity and associated data can be cryptographically authenticated by any party, without relying on the manufacturer's or any single intermediary's attestation;
3. **Network accessibility**, the identity and associated data can be retrieved and used by authorised parties across different systems without proprietary barriers.

These three conditions mark a categorical boundary between IoP and prior product identification systems. Traditional identification technologies, GS1 barcodes, EAN codes, and even first-generation RFID systems, satisfy none of the three conditions in full. Table 1 maps the progression from conventional product identification to the IoP standard across six analytically distinct dimensions, and clarifies how the EU Digital Product Passport (DPP) relates to the IoP infrastructure layer.

The table reveals three important points. First, GS1 barcodes and traditional RFID represent *category-level* identification, they describe what a product is, not *which* product it is. IoP's fundamental advance is the shift to item-level persistent identity: the difference between knowing that a batch of products is Parmigiano Reggiano and knowing that *this specific wheel* was produced at *this facility* on *this date* and has been authenticated at every point in its supply chain journey. This shift is what makes E3 and E4 economically possible: financial transactions and automated enforcement require individual asset identity, not category codes.

Second, the EU DPP sits at the application layer *above* the IoP infrastructure layer rather than at the same level. DPP solves a specific compliance question (“Does this product meet sustainability standard X?”) and is designed to answer it within a defined regulatory framework. IoP theorises the infrastructure conditions that would make DPP, and any other product identity application, generative: open namespaces, shared resolvers, and no-permission-required application development. A DPP built on IoP infrastructure principles would be interoperable with inventory financing systems, circular economy marketplaces, and AI procurement agents without redesign; a DPP built on a closed compliance platform would not.

Third, the progression from barcode to RFID to DPP to IoP is not merely technological but ontological: each stage changes what a product *is* in digital systems, from an

**Table 1.** From conventional product identification to the Internet of Products: a comparative analysis.

<b>Dimension</b>	<b>GS1 Barcode / EAN</b>	<b>RFID (first-gen)</b>	<b>EU DPP (ESPR)</b>	<b>IoP Product Identity Layer</b>
<b>Scope of identification</b>	Product <i>category</i> (SKU); all units of a model share the same code	Item-level possible but typically batch or pallet in practice	Item-level for regulated categories (batteries from 2027)	Item-level for any product; globally unique per physical unit
<b>Identity persistence</b>	Exists only in manufacturer/retailer databases; no cross-organisational continuity	Depends on reader infrastructure; identity lost across organisational boundaries	Persists within EU regulatory scope; dependent on registry operator continuity	Cryptographically anchored (W3C DID); survives organisational change and jurisdictional scope
<b>Verifiability</b>	None, any printer can reproduce a barcode; no cryptographic authentication	Limited, chip cloning possible; no end-to-end cryptographic proof	Regulatory attestation within DPP framework; third-party audit required	Cryptographic proof by any party; no reliance on issuer attestation
<b>Data richness</b>	Static: product code and price only	Moderate: adds location and movement data	Rich but compliance-scoped: mandated fields for specified sustainability attributes	Unbounded: accumulates any structured data across full lifecycle from any authorised source
<b>Economic participation</b>	Passive: tracking object only; cannot participate in transactions	Passive: tracking object only	Compliance instrument: answers regulatory questions; not designed for financial or market transactions	Active entity: can be priced, traded, tokenised, and used as collateral; participates in programmable commerce (E3–E4)
<b>Generativity</b>	None: closed, application-specific	Low: proprietary reader ecosystems	Low: application-layer compliance tool; not designed for unanticipated third-party innovation	High (by design): open standards enable unanticipated applications without infrastructure designer's approval

anonymous category member to a regulated compliance subject to a persistent economic entity. This ontological progression is the core theoretical claim of the paper.

The three stages of internet infrastructure evolution can be expressed through this lens:

- *Internet of Information*: persistent, verifiable, network-accessible digital identity for informational objects (DNS, URLs).
- *Internet of Value*: persistent, verifiable, network-accessible digital identity for financial assets (wallet addresses, token protocols).
- *Internet of Products*: persistent, verifiable, network-accessible digital identity for physical goods (DIDs, product identity infrastructure).

### The Product Identity Layer

The central theoretical construct of this paper is the **Product Identity Layer**, a digital infrastructure layer in which physical goods possess persistent, verifiable, and network-accessible digital identities. This layer functions as an intermediary connecting three previously separate economic domains: the Physical Economy (manufacturing, logistics, retail, and product use); the Digital Economy (data platforms, AI systems, and digital marketplaces); and the Financial Economy (tokenised asset markets, inventory financing, and programmable commerce).

The Product Identity Layer is analogous to the Domain Name System for the Internet of Information. DNS created infrastructure that allowed information to be universally identified and accessed across networks, enabling the Web as an open, generative platform. Crucially, DNS succeeded because it was governed as neutral, open infrastructure: no single commercial actor controlled it, and any developer could build applications on top of it without permission. A Product Identity Layer with comparable openness and neutrality could enable product data services, product-service systems, and product-backed financial instruments to emerge as generative applications.

Generativity in the IoP context requires specificity beyond the general infrastructure theory claim that ‘open systems are generative.’ DNS achieved generativity through three design properties: (1) an open namespace, any party could register a domain without permission from existing holders; (2) a shared resolver, any device could resolve any domain name through a common protocol; and (3) no permission required to build applications, email, the Web, streaming, and e-commerce all emerged without DNS designers’ approval. IoP achieves analogous generativity when: (1) any product can be assigned a persistent identity using open standards (W3C DIDs, GS1 Digital Link) without permission from manufacturers or platform operators; (2) any party can resolve a product identity and retrieve associated lifecycle data through a common protocol without proprietary API licences; and (3) any developer can build applications on top of

product identity data, circular economy marketplaces, automated insurance, tokenised asset platforms, AI agent procurement systems, without the identity infrastructure designer's approval.

Supporting technologies currently include: GS1 identification standards (with GS1 Digital Link enabling migration to item-level web-addressable identity); Decentralised Identifiers (DIDs, W3C Recommendation 2022) providing cryptographically verifiable, organisationally independent item-level identity; QR codes and NFC tags as physical-digital interfaces; blockchain systems providing tamper-resistant data records; and IoT sensors and Digital Twins (Grieves and Vickers, 2017; Tao et al., 2019) providing continuous data enrichment and computational product state representation.

### Architecture of the Internet of Products

The Internet of Products can be conceptualised as a layered architecture consisting of four interconnected strata:

**Identity Layer.** The foundation of IoP is the assignment of a unique, persistent, organisationally independent digital identity to each physical product. This identity must be globally unique, cryptographically verifiable, and not dependent on any single organisation's continued operation, requirements that W3C DID standards address. The challenge of establishing item-level rather than category-level identity represents the primary departure from existing GS1 barcode infrastructure.

**Data Layer.** Once a product possesses a persistent identity, it can accumulate structured data throughout its lifecycle: manufacturing information, logistics events, ownership transfers, environmental metrics, maintenance records, and end-of-life data. The critical design challenge is interoperability: data generated by different organisations using different systems must be linkable to the same product identity. EU Digital Product Passport standards represent a regulatory mandate for a minimum viable version of this layer, while digital twin architectures represent its most data-rich instantiation.

**Trust Layer.** Product identity and lifecycle data require mechanisms to ensure integrity, authenticity, and trust across organisational boundaries. Distributed ledger technologies, cryptographic signatures, and shared data governance frameworks provide tamper-resistant records. This layer must confront the oracle problem directly: the Trust Layer can guarantee that data, once recorded, has not been modified, but it cannot guarantee that data was accurate when first recorded. Trusted hardware with tamper-proof attestation, third-party verification systems with economic skin-in-the-game, and statistical sampling combined with reputational mechanisms represent practical mitigation strategies, though none fully eliminates the oracle problem.

**Application Layer.** Built on the previous three layers, the Application Layer enables

new digital services and economic interactions centred on individual products: automated product authentication systems, smart warranties and insurance contracts, circular economy marketplaces with verified product histories, tokenised ownership of physical goods, AI-driven supply chain optimisation, and digital twin applications. The *generativity* of this layer, the extent to which innovators can develop unforeseen applications using the identity infrastructure without central permission, is the primary measure of IoP's long-term economic potential.

**Table 2.** Architecture of the Internet of Products with design challenges and current instantiations.

Layer	Primary Function	Key Technology	Critical Design Challenge	Current Instantiation
Identity	Assign unique, persistent product identity	GS1 Digital Link; W3C DIDs; NFC/RFID	Item-level uniqueness; organisational independence	EU Battery Regulation unique battery ID; pharma serialisation (EU FMD)
Data	Accumulate structured lifecycle data	Cloud databases; DPP standards; IoT; Digital Twins	Cross-organisational interoperability	EU Digital Product Passport; GBA Battery Passport
Trust	Ensure data integrity and provenance	Blockchain; cryptographic signatures; trusted hardware; auditors	Oracle problem; initial data accuracy	IBM Food Trust; Walmart Food Safety blockchain
Application	Enable new economic interactions	Smart contracts; AI; tokenisation platforms; digital twins	Generativity; open vs. closed architecture	Homie Product-as-a-Service; Parmigiano RFID authentication

### The 4E Model of Product Digitisation

To describe the process through which physical products evolve into programmable economic entities within the Internet of Products, this study proposes the **4E Model of Product Digitisation**. The model describes four sequential stages through which products acquire progressively deeper digital identity.

The sequential logic reflects both technical and economic dependencies: E2 requires a persistent identity anchor from E1 to link data across systems; E3 requires sufficient data richness from E2 to support credible valuation; E4 requires the exchange mechanisms of E3 to provide the economic context in which automated execution creates value.

The sequentiality warrants clarification on three potential edge cases. **First**, can E2 exist without E1? Batch-level traceability systems accumulate rich lifecycle data at the batch

**Table 3.** The 4E Model of Product Digitisation.

<b>Stage</b>	<b>Process</b>	<b>Economic Outcome</b>	<b>Enabling Technologies</b>	<b>Minimum Viable Indicator</b>
E1, Encode	Assign unique persistent digital identity to each physical product item	Product becomes individually identifiable in digital systems across organisational boundaries	QR codes; NFC/RFID; serialised GS1 barcodes; W3C DIDs	Unique identifier linkable across $\geq 2$ independent organisational systems
E2, Enrich	Accumulate structured lifecycle data around the product identity	Product gains a traceable digital history; enables AI-driven service and analytics	IoT sensors; ERP/PLM integration; DPP standards; blockchain records; Digital Twins	Data spans $\geq 3$ lifecycle stages from $\geq 2$ independent sources
E3, Exchange	Enable product to participate in digital economic transactions as a verifiable entity	Product becomes tradeable or financeable digital asset; price premiums from verified provenance	RWA tokenisation platforms; smart contracts; digital marketplaces; inventory financing APIs	At least one financial or commercial transaction references product identity
E4, Execute	Automate economic interactions via smart contracts directly linked to specific product identities	Product becomes a node in programmable digital commerce; enforcement costs approach zero	Smart contracts; IoT triggers; oracle networks (Chainlink); DeFi protocols	At least one economic action executes automatically on verified product event

level without item-level persistent identity. The IoP framework classifies these as *pre-E1* systems: they achieve E2-like data richness at the category/batch level but not at the entity level that defines IoP. The distinction matters because batch-level data cannot support E3 financial transactions (which require individual asset identity) or E4 automated execution. **Second**, can E1 skip directly to E3 without E2? The Parmigiano Reggiano case demonstrates this: minimal authenticated identity data at E1 creates immediate price premium value at E3 without requiring full E2 data richness. This reveals that the E3 value threshold varies by product: for authentication-scarce goods, E1 alone exceeds the minimum viable E3 threshold. **Third**, can E4 exist without E3? Automated execution (E4) could be triggered by verified product events without the product itself being traded (E3), for example, an automated warranty repair dispatch triggered by an IoT fault signal. The framework treats this as E4-within-E2 rather than a violation of the model. These clarifications reinforce rather than undermine the model's staged logic.

### *Boundary Conditions of the 4E Model*

The 4E Model does not apply uniformly to all product categories. Its applicability and the economic value generated at each stage vary systematically with product characteristics, as summarised in Table 4.

### *Failure Modes of the 4E Model*

Each stage of the 4E Model is subject to characteristic failure modes. Recognising these failure modes is essential for governance design and empirical research.

**E1 Failure, Identity fragmentation:** Multiple incompatible identity systems are deployed for the same product category, preventing cross-system aggregation and destroying the value of interoperability. *Mitigation:* open standards adoption (GS1 Digital Link, W3C DID) enforced through regulation or market power of large buyers.

**E2 Failure, Data provenance corruption (the oracle problem):** Inaccurate or fraudulent data is entered at source and then permanently preserved by blockchain's immutability guarantee. The Trust Layer's cryptographic guarantees apply after recording, not to recording accuracy. *Mitigation:* trusted hardware attestation (IoT devices with secure enclaves), mandatory third-party audits with liability for inaccurate attestation, and insurance structures that align incentives for accurate recording.

**E3 Failure, Token-asset linkage breakdown:** The physical asset underlying a tokenised product changes state (damaged, stolen, or substituted) while the digital token continues to represent original claimed value. This 'physical-digital divergence' is particularly acute in long-lived assets. *Mitigation:* continuous IoT monitoring integrated with token smart contracts, periodic physical verification requirements, and insurance products covering divergence risk.

**Table 4.** Boundary conditions for the 4E Model by product characteristic.

<b>Product Characteristic</b>	<b>Applicability</b>	<b>Most Relevant Stage</b>	<b>Rationale</b>
High unit value (>~\$500)	High, identity ROI clearly positive	E1 through E4	Identity and data costs are small relative to product value; financial applications (E3–E4) create compounding returns
Regulated (pharma, food, EV batteries)	High, regulatory mandate creates immediate installed base	E1 and E2 (compliance-driven); E3 optional	Regulation solves bootstrap problem; compliance investment creates foundation for value-creating applications
Durable/long lifecycle (>5 years)	High, lifecycle data creates compounding value	E2 through E4 increasingly valuable over time	Lifecycle data richness grows with time; service models and secondary markets create sustained E3–E4 value
High counterfeit risk	High, authentication value is immediate and large	E1 and E2 sufficient for core value creation	Counterfeit premium lost without authentication; even minimal identity data creates large economic value
Low unit value, fast-moving consumer goods	Low, identity ROI positive only at aggregate/category level	E1 only (category tracking) typically viable	Per-unit identity costs exceed per-unit value created; batch/category tracking more economical
Perishable/single-use	Moderate, time-bound and context-specific value	E2 (cold chain, condition monitoring) highest value	Identity value concentrated in transit and point-of-use verification; long-term lifecycle narrative has limited value

**E4 Failure, Oracle manipulation:** Smart contract execution is triggered by oracle data that has been manipulated to produce a desired outcome. The Mango Markets exploit of 2022, in which USD 117 million was extracted by manipulating price oracles, illustrates the severity of this risk in financial contexts. *Mitigation:* multi-source oracle aggregation requiring consensus across independent data sources, circuit breakers that pause execution on anomalous data patterns, and decentralised oracle networks (e.g., Chainlink) that distribute the attack surface.

### Bridging the Three Economies

The Internet of Products introduces infrastructure that bridges three previously separate economic domains. The Product Identity Layer connects these domains through a progression of integration: when products acquire persistent digital identities (E1: Encode), product data flows into digital analytics systems, connecting the physical and digital economies. When identities are enriched with verified lifecycle data (E2: Enrich), goods can serve as credible inputs to financial systems, connecting the physical and financial economies. When smart contracts automate economic interactions on verified product events (E4: Execute), the three economies operate as an integrated system with information, physical, and financial transactions synchronised in real time.

This integration does not occur automatically or costlessly. As both the TradeLens failure and the Walmart Food Trust success illustrate, the governance architecture, who controls the identity layer, what standards govern data sharing, who has authority to validate data provenance, determines whether the connecting infrastructure generates broad economic value or replicates existing power structures in digital form.

## Critical Analysis: Counter-Arguments and Limitations

---

### The Oracle Problem: The Constitutive Limitation of Physical Identity Infrastructure

The most significant theoretical limitation of the IoP framework is the oracle problem: the impossibility of guaranteeing the accuracy of data at the point of entry into an immutable system. Blockchain technology provides strong guarantees about data integrity, once recorded, data cannot be altered. But it provides no guarantee about data provenance, whether data was accurate when first recorded (Breidenbach et al., 2021). This limitation is not merely technical but constitutive: it explains why physical product identity infrastructure has not emerged as a natural extension of prior internet infrastructure phases, and it specifies the fundamental design challenge that distinguishes IoP from the Internet of Information and Internet of Value.

In product contexts, the oracle problem manifests as the possibility that a manufacturer

records false sustainability certifications, a logistics provider records false temperature compliance, or a repair service records false maintenance history, all permanently preserved on-chain with the appearance of cryptographic authority. The more authoritative product identity systems become, the greater the incentives for fraud at the point of data entry.

This limitation does not invalidate the IoP framework but significantly shapes its design requirements and governance architecture. Infrastructure-level trust cannot be achieved through blockchain immutability alone: it requires trusted hardware with tamper-proof attestation at critical lifecycle junctions, third-party auditors with economic skin-in-the-game, and statistical sampling verification combined with reputational mechanisms. Importantly, the oracle problem is not unique to blockchain; it applies to any product data system that relies on organisational self-reporting without independent verification. IoP architectures that acknowledge this limitation and design for it are more robust than those that treat blockchain's technical properties as a substitute for organisational trust.

Each mitigation approach involves distinct trade-offs across four dimensions, as summarised in Table 5. No single approach eliminates the oracle problem; effective IoP governance requires combining them based on product category, data criticality, and available infrastructure.

The trade-off analysis has a direct implication for governance design: low-value, high-volume products (FMCG, commodities) can only practically employ statistical sampling, which provides weaker guarantees. High-value, regulated products (pharmaceuticals, EV batteries) justify trusted hardware attestation. The governance architecture of IoP must therefore be *tiered*, applying different verification regimes to different product categories rather than imposing a uniform standard that would be either economically infeasible for low-value goods or insufficiently rigorous for high-stakes assets.

### **Governance Failure: The TradeLens Cautionary Tale**

The failure of TradeLens, Maersk and IBM's blockchain-based maritime logistics platform, shut down in November 2022 after four years and significant investment, provides the most important real-world test of supply chain identity infrastructure to date. The technical implementation was sound; the failure was economic and political.

Competing container shipping companies refused to participate because they were unwilling to share sensitive operational data with a system controlled by their largest competitor. The governance architecture of TradeLens, controlled by Maersk and IBM, created justified concerns about competitive advantage asymmetry: firms sharing data would provide intelligence to a rival while receiving limited information in return. The

**Table 5.** Trade-off analysis of oracle problem mitigation strategies.

<b>Approach</b>	<b>Implementation Cost</b>	<b>Scalability</b>	<b>Protection Level</b>	<b>Best-Fit Product Type</b>
Trusted hardware attestation (IoT secure enclaves, tamper-evident seals)	High upfront; hardware per unit	Limited to products where hardware is economically justified	High, makes falsification physically difficult at point of capture	High-value, durable goods; regulated assets (pharma, EV batteries)
Third-party audit with liability (credentialed inspectors; insurance-backed certification)	Moderate; periodic not continuous	Moderate, scales with auditor network; not real-time	Medium, deters systematic fraud; does not prevent one-time errors	Products with discrete, verifiable lifecycle events (origin, condition change)
Statistical sampling and reputational mechanisms (spot checks; penalty scoring; market exclusion)	Low, probabilistic rather than per-unit	High, applicable at commodity scale	Low to medium, deters but does not prevent fraud; relies on detection probability	Commodity goods; contexts where per-unit verification is economically infeasible
Multi-source oracle aggregation (cross-referencing independent data streams)	Low-to-medium for data; high for coordination	High if standardised data APIs exist	High against single-actor manipulation; weaker against coordinated fraud	Complex supply chains where multiple independent actors generate overlapping data

contrast between TradeLens (competitor-controlled, failed) and Walmart Food Trust (buyer-orchestrated with aligned incentives, succeeded) reveals that governance architecture, not technical design, is the primary determinant of multi-actor IoP adoption.

The DNS analogy is instructive. DNS succeeded as a global infrastructure because it was governed by a multi-stakeholder body (ICANN) with no single commercial interest, and because participation created clear value for all parties through a shared namespace.

Product identity infrastructure faces an analogous governance challenge: it must be designed as neutral, shared infrastructure that no single commercial actor controls, if it is to achieve the multi-stakeholder participation necessary for network effects. This insight directly informs Proposition 6 and Hypothesis H6.

### Privacy-Transparency Tension

The Internet of Products, by making product histories verifiable and potentially public, creates surveillance infrastructure for goods, and, by extension, for the behaviour of people who use those goods. Vehicle IoT data reveals location patterns; smart appliance data reveals domestic routines; wearable device data reveals health behaviours. In supply chains, detailed product data may reveal commercially sensitive manufacturing processes, supplier relationships, and cost structures.

These concerns are not hypothetical. The EU's Data Act (2023), entering into force in September 2025, explicitly addresses IoT data ownership and access rights, recognising that product data generated by users belongs, in important senses, to those users rather than exclusively to manufacturers. GDPR imposes strict requirements on personal data generated by connected products. Zuboff (2019)'s analysis of surveillance capitalism suggests that product-level data infrastructures, if governed primarily by private commercial interests, may create new forms of behavioural prediction and manipulation.

Architectural responses include tiered access systems, with public data (product category, general provenance), restricted data (detailed supply chain events, accessible to verified supply chain participants), and private data (usage patterns, accessible only with explicit consent); selective disclosure using zero-knowledge proofs enabling parties to verify specific claims without revealing underlying data; and privacy-preserving computation (federated learning, homomorphic encryption) allowing analytics on encrypted product data without exposing raw records.

### Scalability and Economic Viability for Low-Value Products

The economic case for IoP is strongest for high-value, regulated, or long-lived products where lifecycle data creates compounding value. For low-value, fast-moving consumer goods, packaged food commodities, basic apparel, common hardware, the per-unit economics of item-level identity assignment and data management may be negative at

current technology costs.

This creates a stratified adoption pattern where IoP progresses fastest where regulatory mandates (EU Battery Regulation, EU ESPR) or business case clarity (pharmaceutical serialisation, luxury authentication) overcome economic barriers. Broad adoption in commodity sectors likely requires dramatic reduction in identity assignment costs, regulatory mandate (EU DPP progressively extending to most product categories by 2030), or emergence of platform models in which identity infrastructure costs are distributed across applications generating value.

## Illustrative Case Analysis

---

While empirical validation of the full IoP framework awaits future research, three existing deployments illustrate key propositions and reveal conditions for success and failure. The cases serve an *abductive* function consistent with the theory-building research design described in Section 2: they ground abstract theoretical claims in observable patterns, generate insights that refine boundary conditions, and expose mechanisms that deductive reasoning alone would not surface (Dubois and Gadde, 2002). They are not a representative sample; analytical generalisation rather than statistical generalisation is the operative logic. Table 6 maps each case to the propositions it primarily illustrates.

Cases were selected to satisfy three criteria: (1) each case must illustrate a meaningfully different subset of the 4E stages, to provide non-redundant theoretical coverage; (2) cases must span different governance models and industry contexts; and (3) cases must include both successes and failures within the broader analysis. On criterion (3), the IoP framework analysis in Section 4.2 explicitly incorporates TradeLens as a governance failure case, and the Walmart Food Trust as a success case with contrasting governance architecture. These are analysed as counter-arguments rather than supporting cases to prevent cherry-picking of confirmatory evidence. Future empirical research should apply systematic sampling across IoP deployments to test framework propositions with appropriately representative coverage.

### EU Falsified Medicines Directive (FMD) Pharmaceutical Serialisation: Mandated E1–E2 at Scale

The EU Falsified Medicines Directive (FMD, EU 2011/62) mandated item-level serialisation for all prescription medicines sold in EU markets, fully operational since February 2019. Every prescription pack must carry a unique identifier (Datamatrix code encoding a product code, serial number, batch number, and expiry date), registered in a national medicines verification system and verified at the point of dispensing by pharmacists. The system covers over 14 billion packs annually across 31 European countries, making it the largest deployed instance of E1–E2 product identity

**Table 6.** Case-Proposition Mapping.

Case	4E Stages	Primary Propositions	Key Theoretical Insight
EU FMD Pharmaceutical Serialisation (deployed 2019)	E1–E2 (mandated, operational)	P1, P2, P3	Regulatory mandate with federated governance achieves bootstrap at scale; deployed system enables cross-organisational E2 data; creates unintended E3–E4 foundation
Parmigiano Reggiano RFID Authentication (deployed 2022)	E1→E3 (E2 minimal)	P1, P3, P5	E1 alone creates immediate E3 economic value when authentication problem is severe; validates boundary conditions for high-counterfeit-risk products
Homie Product-as-a-Service (operational)	E1–E4 (full cycle)	P1, P2, P4, P5	Illustrates the <i>logical possibility</i> of complete 4E integration; servitisation business model is structurally dependent on persistent product identity in ways that reveal IoP’s enabling role

infrastructure in existence ([European Medicines Verification Organisation, 2023](#)).

This case illustrates P1 (product identity as digital infrastructure) and P3 (supply chain transformation) through a *deployed, measurable* system. The FMD represents fully operational E1 (each pack carries a globally unique serialised identifier) and E2 (lifecycle data, manufacture, distribution, verification at dispensing, is recorded and accessible across organisational boundaries). The regulatory mandate addressed the bootstrap problem decisively: simultaneous adoption was mandated across all EU market participants, eliminating the chicken-and-egg dynamic that stalls voluntary adoption.

The case also generates important theoretical insights beyond the compliance mandate. First, cross-organisational interoperability at scale was achieved through a federated governance model, national verification organisations (NVOs) in each member state, coordinated by the European Medicines Verification Organisation (EMVO), rather than a single centralised platform. This multi-stakeholder governance architecture is consistent with the prediction in H6 that standards-based and government-mandated governance achieves higher adoption than platform-controlled alternatives. Second, the FMD system has created an *unintended* foundation for E3–E4 value creation beyond compliance: serialisation data now enables pharmaceutical inventory financing with improved collateral verification, and creates the infrastructure base for automated dispensing and smart contract-based supply chain payments, illustrating the generativity mechanism of infrastructure theory.

### **Parmigiano Reggiano Microchip Authentication (2022): E1 with Immediate E3 Value**

The Parmigiano Reggiano Consortium began embedding edible RFID microchips (1 mm diameter, cost <USD 0.01 per unit) in cheese rinds during production. The microchip enables immediate verification of authentic geographic origin, production facility, production date, and conformance with DOP (Protected Designation of Origin) standards, scannable by consumers and supply chain actors throughout the product lifecycle.

This case illustrates P1, P3, and P5, and validates the boundary conditions table in an important way. E1 (Encode) creates immediate E3 (Exchange) value without requiring full E2 data richness: even minimal authenticated identity data, this is genuine Parmigiano Reggiano, from this facility, on this date, is sufficient to command price premiums and prevent counterfeit substitution. The cheese counterfeit market was estimated to exceed the authentic market in some segments, making the USD 0.01 microchip cost negligible relative to the price premium captured through verified authenticity.

The case demonstrates the boundary condition that high counterfeit risk products justify immediate E1 investment even without full E2 infrastructure, and that minimal identity data can create substantial economic value when the information asymmetry problem is severe. It would not translate directly to commodity cheeses where counterfeit risk is low and margins thin, confirming the boundary conditions in Table 4.

### **Homie Washing Machine-as-a-Service (Netherlands): E1–E4 as Theoretical Illustration**

Homie, a Dutch start-up, deploys washing machines as a service, customers pay approximately EUR 20–30 per month rather than purchasing the machine, with IoT sensors monitoring cycle count, energy consumption, and fault indicators. Homie retains ownership and full responsibility for maintenance and replacement.

This case is presented as a *theoretical illustration* rather than an empirical benchmark: it demonstrates the logical structure of a complete 4E Model instantiation and reveals the mechanisms through which persistent product identity enables economic structures that are qualitatively unavailable without it. The case is not used to validate specific quantitative claims about performance outcomes.

The 4E logic is as follows. E1: each machine carries a persistent identity linked to the service contract, enabling the asset to be tracked, valued, and managed throughout its operational life regardless of changes in customer or location. E2: continuous IoT-generated operational data (cycle count, energy profile, fault history) enables predictive maintenance scheduling and real-time asset condition assessment, this data is

the foundation of the service model's economics (P2). E3: because each machine has a persistent identity and a verifiable operational history, the fleet functions as a financeable asset base; contracted future cash flows, backed by asset condition data, can serve as collateral for inventory financing (P5), a financial structure that would be impossible without E1–E2 infrastructure. E4: maintenance dispatch is automatically triggered by IoT fault signals without human intervention, eliminating enforcement costs and creating programmable commerce.

The theoretical insight this case illustrates (P4) is structural: the complete 4E implementation enables a servitisation business model, and its associated sustainability outcomes of extended product lifespans and high component reuse rates, that is *economically infeasible* without persistent product identity. Without E1, asset tracking across customer relationships is impossible. Without E2, predictive maintenance cannot replace reactive repair, destroying the economics of manufacturer-retained ownership. This structural dependency, rather than any specific performance figure, is the theoretically significant claim. Independent research on product-service systems in consumer appliances (Baines et al., 2007) confirms the general mechanism: manufacturer-retained ownership with continuous monitoring substantially extends product lifespans and improves resource efficiency relative to conventional retail models.

## Research Propositions and Testable Hypotheses

---

Based on the conceptual framework and critical analysis developed above, six research propositions are formulated. Five propositions (P1, P2, P3, P5, P6) are accompanied by testable hypotheses (H1–H3, H5–H6) with operationalised constructs and identified data sources. Proposition 4 is positioned as a priority research agenda item (Section 8.5) because the cross-industry measurement infrastructure required for empirical testing does not yet exist. Each proposition is grounded in the theoretical mechanisms identified in Section 3.8.

### Proposition 1: Product Identity as Digital Infrastructure

Products with persistent digital identities will function as identifiable entities within digital economic networks, enabling participation in digital systems beyond traditional supply chain management. *Theoretical grounding*: transaction cost theory predicts that persistent identity reduces search and verification costs; infrastructure theory predicts that open identity standards generate positive externalities beyond any single application.

**H1:** Supply chain actors who adopt item-level product identity (operationalised as: unique identifier linkable across  $\geq 3$  independent organisational systems) will demonstrate significantly lower inter-organisational transaction costs (measured as: time and cost of product verification, traceability, and dispute resolution) compared to actors using

category-level identification only, after controlling for firm size, supply chain complexity, and industry sector.

### Proposition 2: Product Lifecycle Data and Transparency

Products with persistent digital identities will generate continuous lifecycle data that significantly increases supply chain transparency, enabling more effective risk management, compliance monitoring, and circular economy initiatives. *Theoretical grounding*: information asymmetry theory predicts that verifiable lifecycle data reduces adverse selection; transaction cost theory predicts that observable product histories reduce verification costs across all supply chain actors.

**H2:** Higher completeness of product lifecycle data records (operationalised as: percentage of lifecycle stages with verified data from  $\geq 2$  independent sources) will be positively associated with: (a) reduction in counterfeit detection time (days from introduction to identification); (b) reduction in product recall scope (products recalled relative to actual contaminated units); and (c) material recovery rate at end of life (percentage of material mass recovered for secondary use).

### Proposition 3: Supply Chain Transformation

IoP adoption will increase supply chain transparency and traceability at the product level, leading to measurable reductions in counterfeiting, improved recall efficiency, and enhanced supply chain coordination.

**H3:** Firm-level adoption of product-level identity and lifecycle data systems (operationalised as: 4E stage reached, E1 through E4, assessed through structured audit) will predict: (a) reduction in counterfeit claims by product category; (b) reduction in product recall time-to-containment; (c) improvement in on-time-in-full delivery rates. Relationships will be moderated by supply chain complexity (number of tiers and countries) and mediated by data sharing breadth across supply chain tiers.

### Proposition 4: Emergence of Product Data Ecosystems

Product identity proliferation will enable digital ecosystems centred on lifecycle data, in which platform-based services, analytics applications, and circular economy markets develop as unanticipated applications of identity infrastructure. *Theoretical grounding*: digital infrastructure theory predicts that generative infrastructure enables unanticipated innovations; the generativity mechanism (Zittrain, 2006) predicts that open identity standards will support more third-party innovation than proprietary alternatives.

*Note on testability*: This proposition cannot yet be operationalised as a testable hypothesis in the standard sense because the cross-industry measurement infrastructure required, standardised registries of third-party services consuming product identity data,

does not yet exist. Accordingly, Proposition 4 is positioned as a **priority research agenda item** (see Section 8.5) rather than a near-term testable hypothesis. The proposition is nonetheless theoretically grounded and falsifiable in principle: as EU DPP registries and pharmaceutical serialisation systems generate longitudinal data on third-party ecosystem development, the relationship between identity infrastructure maturity and innovation rates can be examined empirically. Initial exploratory work is feasible in pharmaceutical serialisation (EU FMD, operational since 2019), where compliance databases and developer ecosystems are already documented.

### Proposition 5: Product-Based Financial Systems

IoP will facilitate integration of physical goods into digital financial systems through product-level identity and tokenisation, enabling new financial instruments and reducing the cost of capital for inventory financing. *Theoretical grounding:* transaction cost theory predicts that verifiable product identity reduces verification costs in financial contracting; information asymmetry theory predicts that credible product data reduces adverse selection in inventory financing markets.

**H5:** Products with E3-level digital identity (operationalised as: product identity referenced in at least one financial transaction, with lifecycle data accessible to the financing party) will demonstrate significantly lower cost of capital for inventory financing (measured as: interest rate spread versus unsecured financing of equivalent maturity and credit quality) compared to equivalent products without verified digital identity. The relationship will be moderated by data completeness (percentage of lifecycle stages with verified data) and asset liquidity in the tokenised market (bid-ask spread for tokenised product claims). Data for testing this hypothesis can be drawn from: (1) emerging tokenised RWA platforms (BlackRock BUIDL, JPMorgan Onyx, Franklin Templeton) that disclose financing terms alongside product identity data; and (2) pharmaceutical serialisation-linked inventory financing, where EU FMD compliance creates a quasi-experimental comparison. Researchers should pre-register designs and plan for sequential analysis as market data accumulates.

### Proposition 6: Governance of Product Identity Infrastructure

IoP development will produce competing governance models with significantly different implications for adoption rates, data sharing breadth, and third-party innovation development. *Theoretical grounding:* Williamson (1981)'s transaction cost theory of governance, combined with Hanseth and Lyytinen (2010)'s infrastructure governance analysis, predicts that governance architecture, specifically, whether a single commercial actor controls identity infrastructure, will be the primary determinant of multi-stakeholder adoption.

The four-category governance typology is grounded in these theoretical foundations. *Platform-controlled governance* (e.g., TradeLens under Maersk/IBM) represents Williamson’s hierarchical governance by a market participant with competing interests. *Industry standards-based governance* (e.g., GS1, W3C DID) represents market-like coordination through neutral standards, with no single actor holding control rights. *Blockchain-decentralised governance* distributes control across network participants through cryptographic consensus, a novel form that Williamson’s framework did not anticipate but which can be analysed through its transaction cost properties. *Government-mandated governance* (e.g., EU ESPR/DPP) represents public regulatory authority that can mandate participation and set interoperability standards. These four categories are exhaustive of the primary actors who can hold ultimate authority over identity standards and data access rules, and mutually exclusive in terms of who holds that ultimate authority, though hybrid arrangements are common in practice.

**H6:** The governance architecture of product identity infrastructure (operationalised as: the four-category typology, assessed through institutional analysis of who controls identity standard-setting and data access rules, coding based on charter documents, governance board composition, and API access terms) will significantly predict: (a) adoption rate among competing supply chain actors (highest for standards-based and government-mandated; lowest for platform-controlled by a market participant in that market), measured as percentage of market participants with active product identity records after 36 months; (b) data sharing breadth across supply chain tiers (measured as: number of supply chain tiers with data contributing to product identity record, assessed through supply chain mapping studies in pilot industries); and (c) third-party innovation development on the infrastructure (measured as: number of distinct organisations consuming identity data APIs, tracked through developer registration logs where available). Data for testing H6 exists in comparably structured natural experiments: EU DPP implementations across battery (2027), textile (2028), and electronics (2030) sectors involve the same regulatory mandate but are expected to produce different governance architectures, enabling difference-in-differences analysis.

---

## Discussion: Theoretical Contributions and Implications

### Theoretical Contributions

This study advances theory in three directions, each building on but extending prior contributions in ways that address the gaps identified in Section 2.9.

**First contribution, IoP as infrastructure, not application.** By framing the Internet of Products as a digital infrastructure layer, with explicit grounding in [Hanseth and Lyytinen \(2010\)](#)’s infrastructure design theory, rather than as an application system,

this paper generates predictions about generativity, bootstrap dynamics, and governance challenges that domain-specific framings do not produce. DPP research addresses how to design a particular application of product identity; IoP infrastructure theory addresses what conditions make *any* application of product identity generative and economically transformative. This distinction carries substantive theoretical weight. Infrastructure theory predicts that early design decisions about openness and standards will have long-lasting, path-dependent consequences, a prediction absent from existing IoP-adjacent literatures. Concretely, it predicts that a DPP system designed as a closed compliance instrument (answering only “Does this product meet regulation X?”) and one designed as open, generative infrastructure (allowing any developer to build applications on product identity data) will produce radically different long-run economic outcomes, even if they appear functionally equivalent at the point of initial deployment. This prediction cannot be derived from DPP literature, IoT research, or supply chain traceability research; it requires infrastructure theory. The implication for current EU DPP governance design is therefore not merely academic but urgently practical.

**Second contribution, the categorical distinctiveness of physical identity.** Prior infrastructure theory has been applied to information and financial identity systems. The IoP framework identifies a constitutive difference between physical and intangible identity: physical products can be damaged, stolen, or substituted independently of their digital representations, creating the oracle problem as a design challenge without precedent in prior internet infrastructure phases. This theoretical insight, that physical identity is categorically harder than intangible identity, explains both why product identity infrastructure has not emerged naturally from prior phases and what specific design problems it must solve.

The categorical distinction has a further implication that the prior literature has not drawn: it establishes an *asymmetric trust architecture requirement* for IoP that does not apply to information or financial identity infrastructure. For DNS, cryptographic proof of identity is both necessary and sufficient, there is no “physical object” that can be swapped or degraded independently of its DNS record. For blockchain wallet addresses, possession of the private key is both necessary and sufficient. For IoP, cryptographic proof of identity is *necessary but not sufficient*: it must be supplemented by physical verification regimes (trusted hardware, third-party audit, statistical sampling) that have no direct analogue in prior infrastructure phases. The trade-off analysis in Table 5 is a direct theoretical consequence of this asymmetry, it specifies what “sufficient” trust requires across product categories, and why no single verification mechanism can serve all contexts.

**Third contribution, theoretical unification with novel cross-stream predictions.** By integrating transaction cost theory (Coase, 1937) and information asymmetry theory (Akerlof, 1970) as the primary economic mechanisms through which product identity

creates value, the paper provides micro-foundations absent in existing digital supply chain and IoT research. These theoretical lenses do more than label existing research streams: they generate *cross-stream predictions* that no single stream produces in isolation.

Three such predictions are worth making explicit. *First*, the integration of information asymmetry theory with infrastructure theory predicts that authentication-scarce products (those where buyers face severe “lemon” risk) will exhibit disproportionately high returns to E1 investment, a prediction confirmed by the Parmigiano case but derivable *prior* to any empirical observation from the theoretical framework alone. *Second*, the integration of transaction cost theory with infrastructure governance theory predicts that platform-controlled IoP governance will fail in horizontal industry sectors (where participants are competitors) but succeed in vertical sectors (where a powerful buyer can mandate participation), a prediction that retrospectively explains the TradeLens failure and Walmart Food Trust success, and prospectively guides governance design for EU DPP. *Third*, the integration of all three theories yields the boundary conditions in Table 4, specifying which product categories will adopt which stages first, a prediction that would not follow from any single theoretical tradition. Infrastructure theory alone cannot predict that high-counterfeit-risk products reach E3 faster than durable goods; it requires Akerlof’s adverse selection mechanism. Transaction cost theory alone cannot predict that governance neutrality is a precondition for generativity; it requires Hanseth and Lyytinen’s path dependency logic. The theoretical integration is therefore not merely taxonomic but generative: it produces testable predictions that are not derivable from constituent theories individually.

### Implications for Research

The research agenda emerging from this paper spans multiple disciplines and levels of analysis, and the three theoretical contributions above generate specific, non-obvious directions that prior surveys of this domain have not identified.

For **information systems researchers**, the generativity mechanism of the Product Identity Layer presents a distinctive opportunity. Prior IS infrastructure research has examined generativity in information (the Web) and financial (DeFi) contexts; IoP offers a third empirical setting that differs categorically because of the oracle problem. The central IS research question is: which governance and design choices (openness of identity standards, resolver access, data portability requirements) are necessary and sufficient conditions for generativity, and do the same conditions that produced generativity in DNS and blockchain identity hold for physical product identity? This question requires theory-driven field studies and natural experiments, not replication of prior infrastructure studies, but a new theoretical application.

For **operations management researchers**, H1–H3 offer empirically tractable

hypotheses using variation already present in the environment. The phased implementation of EU DPP across product categories (batteries 2027, textiles 2028, electronics 2030) creates a staggered difference-in-differences design with a common regulatory treatment and heterogeneous governance arrangements, a natural experiment that did not exist when prior supply chain IT adoption studies were conducted. The critical methodological challenge is measurement: operationalising “transaction costs” at the product identity level requires new instruments that existing SCM surveys do not provide. Developing and validating these instruments is itself a priority research contribution.

For **financial economics researchers**, H5 sits at the intersection of digital asset finance and inventory finance, two literatures that have not been brought into dialogue. The theoretical prediction is that E3-level product identity reduces the adverse selection problem in inventory lending by making collateral quality verifiable at item level rather than category level. Testing this requires matched samples of serialised and non-serialised products entering inventory financing, with EU FMD pharmaceutical data providing the most immediate quasi-experimental setting.

For **governance and policy scholars**, H6’s governance typology generates a comparative institutional prediction that is directly relevant to ongoing EU DPP design decisions. The prediction, that platform-controlled governance systematically underperforms standards-based and government-mandated governance in horizontal markets, is falsifiable through the TradeLens/Walmart Food Trust archival comparison and, prospectively, through EU DPP implementation monitoring. Policy scholars can contribute by developing the institutional analysis methods needed to code governance architectures consistently across industries and jurisdictions.

Cross-disciplinary collaboration is particularly important for two challenges. The oracle problem requires both cryptography expertise (trusted hardware attestation, zero-knowledge proofs) and organisational trust theory (incentive design, liability structures, reputational mechanisms). The privacy-transparency tension requires both technical privacy engineering (zero-knowledge proofs, homomorphic encryption) and sociological analysis of surveillance and power asymmetry. Neither challenge can be adequately addressed within a single disciplinary tradition.

### Implications for Practice

For **business leaders**, the 4E Model and boundary conditions table (Table 4) provide a staged diagnostic framework for assessing current IoP position and investment priorities. The key insight is that the appropriate entry point and investment level varies systematically by product characteristic, there is no universal “IoP strategy.”

Organisations in regulated industries (pharmaceuticals, EV batteries, food safety) face imminent compliance mandates that require E1–E2 capability by defined deadlines. The EU FMD case demonstrates that compliance investment creates an infrastructure foundation that generates value well beyond the compliance mandate itself:

pharmaceutical serialisation data now enables inventory financing improvements and creates the base for E3–E4 applications that were not contemplated when the regulation was designed. The strategic recommendation is therefore not merely to meet compliance requirements but to design compliance infrastructure for *generativity*, using open standards (W3C DIDs, GS1 Digital Link) rather than proprietary systems, and ensuring data portability for future applications.

High-value product manufacturers facing significant counterfeit exposure have the most immediate and self-financing business case: the Parmigiano case demonstrates that E1 investment at USD 0.01 per unit is recoverable through price premiums within a very small number of product cycles when the authentication problem is severe. For these manufacturers, the governance design of identity infrastructure is a competitive asset, proprietary identity systems that prevent competitor access create lock-in, while open systems that allow consumer-level verification create brand equity.

Industrial equipment manufacturers and consumer durable producers should assess E2–E4 investment through the lens of service model transformation rather than product tracking. The Homie case demonstrates not that IoT-enabled services are “nice to have” but that persistent product identity is the *enabling condition* for manufacturer-retained ownership models, and that these models deliver substantially better resource efficiency outcomes than conventional retail. The strategic question is therefore whether the transition from product sales to service provision is economically viable given current identity infrastructure costs, and the boundary conditions in Table 4 provide the framework for answering it by product category.

For **policymakers**, the governance implications of H6 have direct relevance to current EU DPP design decisions. The oracle problem trade-off analysis (Table 5) implies that DPP governance architecture must be *tiered* by product category, applying trusted hardware attestation requirements to high-value regulated goods (pharma, EV batteries) while relying on statistical sampling for commodity sectors where per-unit verification is economically infeasible. A uniform verification standard across all product categories would either impose unacceptable costs on commodity sectors or provide inadequate assurance for high-stakes assets.

More broadly, the TradeLens failure and the theoretical prediction that governance neutrality is a precondition for multi-stakeholder adoption together suggest that DPP infrastructure governed as an open, multi-stakeholder system, analogous to ICANN for

DNS, will achieve higher adoption rates, broader data sharing, and more third-party innovation than a platform controlled by any single commercial actor or industrial consortium. The EU's ESPR/DPP represents a rare historical opportunity to establish "the DNS for products" as public infrastructure rather than as private platform: the governance architecture choices being made now will have path-dependent consequences that persist for decades. Policymakers designing DPP governance architecture are, in effect, deciding whether the Internet of Products becomes an open platform for economic innovation or a compliance burden.

## **Research Agenda for Future Studies**

---

### **Empirical Validation of the 4E Model**

The most immediate research priority is empirical validation of the 4E Model's stage structure and boundary conditions. Mixed-methods designs, combining quantitative analysis of IoP adoption data with qualitative case studies of implementation journeys across regulated industries (batteries, pharmaceuticals, food), would test whether stage progression is indeed sequential, identify conditions under which stages are skipped or combined, and quantify the economic value created at each stage transition. Longitudinal designs are particularly valuable given that lifecycle data accumulation value grows over time.

### **Oracle Problem and Trust Architecture Research**

Empirical research on data provenance mechanisms is urgently needed. Which combination of trusted hardware attestation, third-party audit with liability, and reputational mechanisms most effectively addresses the oracle problem across different product categories and supply chain configurations? Natural experiments provided by EU DPP implementation, in which different industries adopt product identity systems under common regulatory requirements but different governance arrangements, offer particularly clean identification opportunities. This research requires collaboration between information systems, cryptography, and organisational trust scholars.

### **Governance Models and Adoption Dynamics**

Comparative institutional analysis of competing governance models can develop and test H6. The contrast between TradeLens (failed, competitor-controlled) and Walmart Food Trust (succeeded, buyer-orchestrated) provides the initial comparative case. EU DPP implementation across multiple product categories, batteries, textiles, electronics, with different governance arrangements will create a natural panel for testing governance effects on adoption rate, data sharing breadth, and innovation development. Researchers should design longitudinal data collection protocols in advance of full DPP implementation.

## Financial Integration and Cost of Capital

The emergence of tokenised RWA markets provides a natural experiment for testing H5. Event studies around DPP compliance announcements and tokenisation launches could identify market-level responses to product identity infrastructure improvements. Research can examine whether product-level identity data (E3-level) reduces cost of capital for inventory financing relative to category-level data, and whether the reduction is moderated by data completeness and audit quality.

## Product Data Ecosystems and Generativity: Developing Measurement Infrastructure for H4

Proposition 4, that product identity proliferation will enable digital ecosystems centred on lifecycle data, with third-party innovation as a measurable outcome, is theoretically grounded but cannot yet be tested empirically because standardised cross-industry measurement infrastructure does not exist. The priority research task is therefore to *build* that infrastructure: specifically, to develop protocols for tracking the number and diversity of third-party organisations consuming product identity data APIs, across industries with different levels of identity infrastructure maturity.

The EU FMD pharmaceutical serialisation system (operational since 2019) provides the most immediate data source: compliance databases document the installed base of serialised products, while the developer ecosystems around EMVO APIs provide a proxy for third-party innovation. A longitudinal study tracking API consumers from FMD implementation (2019) through the first years of EU DPP battery passport deployment (2027+) would provide a natural before-and-after comparison across two regulated industries. Researchers should engage EU DPP monitoring bodies early to design data collection instruments compatible with this research agenda.

## Sustainability and Circular Economy Outcomes

H2's proposition that product lifecycle data improves circular economy outcomes requires sector-specific empirical investigation. Studies examining material recovery rates before and after DPP implementation in battery, textile, and electronics sectors would provide direct evidence. Difference-in-differences designs exploiting the phased EU DPP implementation timeline across product categories could provide causal identification.

## AI Agents and Agentic Commerce

An emerging research frontier concerns the interaction between product identity infrastructure and AI agents. Schumacher et al. (2025) at McKinsey estimate that AI agents could mediate USD 3 to 5 trillion of global consumer commerce by 2030, with agents autonomously searching, comparing, and transacting based on machine-readable product data. McKinsey identifies product identity infrastructure, specifically,

standardised, machine-readable, verifiable product data accessible through open APIs, as a precondition for agentic commerce to function reliably. Research is needed on how product identity data quality and accessibility affects AI agent decision accuracy and market outcomes, and whether IoP infrastructure creates new forms of algorithmic market power for products with superior data quality, a question that connects platform economics (Rochet and Tirole, 2003) to emerging AI governance concerns.

## Conclusion

---

The internet has evolved through successive layers of digital identity infrastructure, each expanding the scope of the digital economy by conferring persistent, verifiable, network-accessible identity on a new class of entities. DNS gave identity to information, enabling e-commerce and the knowledge economy as unanticipated applications. Blockchain wallet addresses gave identity to financial value, enabling decentralised finance and programmable contracts as unanticipated applications. This paper has argued that giving identity to physical products represents the next such expansion, but one that is categorically harder to achieve because physical products, unlike information or financial assets, can be damaged, stolen, or substituted independently of their digital representations.

The Internet of Products framework is grounded in three theoretical foundations that together explain why product identity creates economic value, why its absence is structurally costly, and what design and governance conditions are necessary for it to be realised as generative infrastructure. Transaction cost theory explains the economic mechanism: persistent product identity reduces search, verification, and enforcement costs that pervade anonymous product markets. Information asymmetry theory explains the market failure that IoP addresses: without verifiable product histories, adverse selection drives out quality goods in a modern iteration of Akerlof's market for lemons. Digital infrastructure theory explains both the transformative potential and the adoption challenges: generativity requires open design; the bootstrap problem requires regulatory mandates or large-buyer orchestration; the installed base problem requires building on existing identification systems rather than replacing them.

The 4E Model, Encode, Enrich, Exchange, Execute, provides a staged framework for product digitisation with specified boundary conditions, failure modes (particularly the oracle problem, which blockchain immutability cannot solve), and explicit proposition-case mappings. Six testable hypotheses with operationalised constructs provide a structured empirical research agenda.

The governance insight may be the most practically significant contribution of this paper: the difference between DPP as a compliance reporting system and DPP as generative

infrastructure is primarily a governance design decision. Open, neutral, multi-stakeholder governance, analogous to ICANN for DNS, is a necessary condition for achieving the generativity that would make product identity infrastructure economically transformative rather than merely operationally useful. Policymakers designing DPP governance architecture are, in effect, deciding whether the Internet of Products becomes an open platform for economic innovation or a compliance burden.

For most of industrial history, products have been anonymous objects once they leave the factory. The Internet of Products offers the possibility that this condition may change, and that the next great infrastructure layer of the digital economy may be built not from data or digital dollars, but from the persistent, verifiable identities of the things we make.

## References

---

- Dubois, A., & Gadde, L.-E. (2002). Systematic combining: An abductive approach to case research. *Journal of Business Research*, 55(7), 553–560.
- European Medicines Verification Organisation (EMVO). (2023). *Annual Report 2022–2023: EU Falsified Medicines Directive Implementation*. Brussels: EMVO.
- Gregor, S. (2006). The nature of theory in information systems. *MIS Quarterly*, 30(3), 611–642.
- Merton, R.K. (1968). *Social Theory and Social Structure*. New York: Free Press.
- Weick, K.E. (1995). *Sensemaking in Organizations*. Thousand Oaks: Sage.
- Yin, R.K. (2018). *Case Study Research and Applications: Design and Methods* (6th ed.). Thousand Oaks: Sage.
- Akerlof, G.A. (1970). The market for lemons: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500.
- Ashton, K. (2009). That Internet of Things thing. *RFID Journal*, 22 June.
- Atzori, L., Iera, A., & Morabito, G. (2010). The Internet of Things: A survey. *Computer Networks*, 54(15), 2787–2805.
- Baines, T., Lightfoot, H., Evans, S., Neely, A., Greenough, R., Peppard, J., et al. (2007). State-of-the-art in product-service systems. *Proceedings of the Institution of Mechanical Engineers Part B*, 221(10), 1543–1552.
- Bank for International Settlements (BIS). (2023). *Tokenisation in the Context of Money and Other Assets: Concepts and Implications for Central Banks*. CPMI Papers No. 217. Basel: BIS.

- Boston Consulting Group & ADDX. (2022). *Relevance of On-Chain Asset Tokenization in 'Crypto Winter'*. Boston: BCG.
- Breidenbach, L., Cachin, C., Chan, B., Coventry, A., Ellis, S., Juels, A., et al. (2021). *Chainlink 2.0: Next Steps in the Evolution of Decentralized Oracle Networks*. Chainlink Labs White Paper.
- Buterin, V. (2013). *Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform*. ethereum.org.
- Casino, F., Dasaklis, T., & Patsakis, C. (2019). A systematic literature review of blockchain-based applications. *Telematics and Informatics*, 36, 55–81.
- Castells, M. (2010). *The Rise of the Network Society* (2nd ed.). Oxford: Wiley-Blackwell.
- Christopher, M. (2016). *Logistics and Supply Chain Management* (5th ed.). Harlow: Pearson.
- Coase, R.H. (1937). The nature of the firm. *Economica*, 4(16), 386–405.
- European Commission. (2022). *Digital Product Passport under the Circular Economy Action Plan*. Brussels: European Commission.
- European Commission. (2023). Regulation (EU) 2023/1542 (EU Battery Regulation). *Official Journal of the European Union*.
- European Commission. (2024). Regulation (EU) 2024/1781 (ESPR, Ecodesign for Sustainable Products Regulation). *Official Journal of the European Union*.
- Firmhouse. (2025). What is Product-as-a-Service? firmhouse.com. Accessed March 2025.
- Global Battery Alliance. (2023). *GBA Battery Passport Technical Framework*. World Economic Forum / GBA.
- Grieves, M., & Vickers, J. (2017). Digital Twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In *Transdisciplinary Perspectives on Complex Systems* (pp. 85–113). Cham: Springer.
- Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Generation Computer Systems*, 29(7), 1645–1660.
- Hanseth, O., & Lyytinen, K. (2010). Design theory for dynamic complexity in information infrastructures: The case of building internet. *Journal of Information Technology*, 25(1), 1–19.
- Heeß, P., Rockstuhl, J., Körner, M.-F., & Strüker, J. (2024). Enhancing trust in global

- supply chains: Conceptualizing Digital Product Passports for a low-carbon hydrogen market. *Electronic Markets*, 34(1), 1–20.
- IBM. (2018). *IBM Food Trust: Walmart Food Safety Initiative*. IBM.com.
- Iansiti, M., & Lakhani, K.R. (2017). The truth about blockchain. *Harvard Business Review*, 95(1), 118–127.
- Ivanov, D., Dolgui, A., & Sokolov, B. (2019). The impact of digital technology and Industry 4.0 on supply chain resilience. *International Journal of Production Research*, 57(3), 829–846.
- Lee, I., & Lee, K. (2015). The Internet of Things (IoT): Applications, investments, and challenges for enterprises. *Business Horizons*, 58(4), 431–440.
- Nakamoto, S. (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*. bitcoin.org.
- Neumann, R. (2013). *The Internet of Products: An Approach to Establishing Total Transparency in Electronic Markets*. Berlin: Springer.
- Nguyen, T., Zhou, L., Spiegler, V., Ieromonachou, P., & Lin, Y. (2018). Big data analytics in supply chain management: A state-of-the-art literature review. *Computers & Operations Research*, 98, 254–264.
- OECD. (2021). *Global Trade in Fakes: A Worrying Threat*. Paris: OECD Publishing.
- OECD. (2023). *Trade in Counterfeit and Pirated Goods*. oecd.org.
- Porter, M., & Heppelmann, J. (2014). How smart, connected products are transforming competition. *Harvard Business Review*, 92(11), 64–88.
- Porter, M., & Heppelmann, J. (2015). How smart, connected products are transforming companies. *Harvard Business Review*, 93(10), 96–114.
- Qliktag Software Inc. (2018). *The Internet of Products*. Qliktag.
- Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4), 990–1029.
- Saberi, S., Kouhizadeh, M., Sarkis, J., & Shen, L. (2019). Blockchain technology and its relationships to sustainable supply chain management. *International Journal of Production Research*, 57(7), 2117–2135.
- Sarma, S., Brock, D., & Ashton, K. (2000). *The Networked Physical World*. MIT Auto-ID Center.
- Schumacher, K., Roberts, R., & Giebel, K. (2025). The agentic commerce opportunity:

- How AI agents are ushering in a new era for consumers and merchants. McKinsey & Company, October 17, 2025. mckinsey.com.
- Shapiro, C., & Varian, H. (1999). *Information Rules: A Strategic Guide to the Network Economy*. Boston: Harvard Business School Press.
- Tao, F., Zhang, H., Liu, A., & Nee, A.Y.C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415.
- Tapscott, D., & Tapscott, A. (2016). *Blockchain Revolution*. New York: Portfolio/Penguin.
- Tilson, D., Lyytinen, K., & Sørensen, C. (2010). Digital infrastructures: The missing IS research agenda. *Information Systems Research*, 21(4), 748–759.
- Treiblmaier, H. (2018). The impact of the blockchain on the supply chain: A theory-based research framework and a call for action. *Supply Chain Management: An International Journal*, 23(6), 545–559.
- Vandermerwe, S., & Rada, J. (1988). Servitization of business: Adding value by adding services. *European Management Journal*, 6(4), 314–324.
- W3C. (2022). *Decentralized Identifiers (DIDs) v1.0*, W3C Recommendation. w3.org.
- World Economic Forum. (2023). *Digital Product Passport: Unlocking the Circular Economy*. Geneva: WEF.
- Williamson, O.E. (1981). The economics of organization: The transaction cost approach. *American Journal of Sociology*, 87(3), 548–577.
- Xu, L., He, W., & Li, S. (2014). Internet of Things in industries: A survey. *IEEE Transactions on Industrial Informatics*, 10(4), 2233–2243.
- Yermack, D. (2017). Corporate governance and blockchains. *Review of Finance*, 21(1), 7–31.
- Zhou, W., Piramuthu, S., & Chu, F. (2020). RFID-enabled traceability in the supply chain. *International Journal of Production Economics*, 220, 107463.
- Zittrain, J. (2006). The generative internet. *Harvard Law Review*, 119, 1974–2040.
- Zuboff, S. (2019). *The Age of Surveillance Capitalism*. New York: PublicAffairs.