International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(22s) 2026 pp. 772-788
https://doi.org/10.70917/ijcisim-2026-5476
Article

Beyond Checkout: A Payment-Centric Scalability Framework for Autonomous Retail

Abhigyan Mukherjee

Independent Researcher, India
Email: abhigyan.mukherjee@yahoo.com

Abstract: Autonomous retail is often described as a computer vision and sensor-fusion breakthrough that removes checkout lines. This paper argues that such a framing is incomplete. Perception accuracy is necessary for autonomous retail, but it is not sufficient: durable scalability also requires mature payment orchestration, customer trust, merchant economics, and ecosystem alignment among issuers, acquirers, processors, schemes, wallets, and merchants. Through a practitioner-informed, design-science-oriented analysis of stadiums and arenas, airports, corporate and university campuses, hospitals, convenience stores, and hotels, this paper offers three contributions: a format-specific Autonomous Retail Scalability Matrix with an explicit scoring rubric, a Frictionless Commerce Payments Maturity Model, and a payment-centric reference architecture. The conceptual analysis suggests that autonomous retail scales best when shopping missions are time-sensitive, baskets are simple, customer identity can be bound at entry, and payment flows can support initial/pre-authorization, incremental authorization, partial reversal, final capture, refunds, chargebacks, and reconciliation without undermining customer trust. The paper also explains why large grocery is more difficult and that real-time basket visibility and shopping-control expectations matter. The paper concludes that the future of autonomous retail is not a single universal checkout model, but a portfolio of format-specific systems in which payments become the invisible backbone of frictionless commerce.
Keywords: autonomous retail; frictionless commerce; payment orchestration; card-present payments; tokenization; authorization reversal; merchant readiness; customer trust; retail scalability

1. INTRODUCTION

Autonomous retail promises to remove one of the most persistent sources of in-store friction: checkout. In its most advanced form, a customer enters a store, selects items, exits, and receives a digital receipt after the system has identified the shopper, associated item interactions with a virtual basket, and charged a payment instrument. Public autonomous retail materials describe the technology as using artificial intelligence, sensors, computer vision, radio-frequency identification (RFID), and automated payment at exit [1], [2].

However, the industry debate has shifted from "Can autonomous checkout work?" to "Where, why, and under what payment and operational conditions can it scale?" Autonomous retail expansion in convenience stores, stadiums, airports, campuses, healthcare, and hospitality coexists with public reporting that many large-scale grocery stores moved away from autonomous checkout toward smart carts that provide real-time basket visibility. [3], [4]. This contrast suggests that autonomous retail is not a universal store replacement model. It is a format-specific operating model whose success depends on mission urgency, basket complexity, identity certainty, operational discipline, and payment resilience.

This paper makes a payment-centric argument: perception accuracy is a necessary condition for autonomous retail, but it is not a sufficient one. Trust and customer experience in automated retail settings have long been examined in operations research, human–computer interaction, and retail-systems literatures; the gap this paper addresses is narrower and specific , the autonomous retail debate in particular has concentrated on cameras and item recognition while underweighting another necessary condition among several interacting determinants: the ability to transform checkout from a visible terminal event into a continuous system of identity, basket confidence, risk, authorization, receipting, dispute handling, and settlement. In this model, payment is not the last step. It is the trust layer that makes the entire autonomous journey commercially acceptable.

The thesis can be stated directly. Autonomous retail will not scale on camera accuracy alone; it will scale where payment orchestration, customer trust, and merchant economics mature enough to convert checkout into a reliable, transparent, and reconcilable commerce layer. The future of frictionless commerce is therefore payment-centric, format-specific, and hybrid rather than universally cashier-less. The remainder of the paper develops this thesis: Section II describes the research approach and the author's positioning; Section III defines autonomous retail against adjacent checkout models; Section IV reviews the literature; Sections V and VI develop the payment-centric architecture and its deep design challenges; Sections VII and VIII present the two frameworks; Sections IX and X discuss scalability conditions and stakeholder implications; Section XI states limitations and a research agenda.

2. RESEARCH APPROACH, AUTHOR POSITIONING, AND DISCLOSURE

The paper uses a hybrid practitioner research approach combining literature review, public case analysis, and design-science-oriented framework construction in the tradition of Hevner et al. [27] and the design science research methodology of Peffers et al. [28]. Mapped to that methodology, the paper implements problem identification (Sections I and IV), objective definition (Section II), artifact design (Sections V, VII, and VIII), and a demonstration step in the form of a retrodictive check against held-out public cases (Section VII.B). Full artifact evaluation, multi-expert scoring with reported inter-rater agreement is beyond the scope of a single-author paper and is stated as future work in Section XI. The maturity model in Section VIII follows, in abbreviated form, the problem-driven construction procedure described by Becker et al. for maturity model development [29]. The goal is to construct generalizable frameworks using public evidence and payments-domain reasoning.

Author positioning. The author is a payments and commerce-technology practitioner with global-scale experience in autonomous retail payment integration and card-present checkout modernization, including direct work with card schemes, acquirers, and processors on the classes of transaction flow this paper analyzes. This practitioner vantage point underlies the paper's original contributions — the Autonomous Retail Scalability Matrix, the Frictionless Commerce Payments Maturity Model, and the payment-centric reference architecture; which are offered as generalizable, vendor-neutral frameworks rather than descriptions of any single deployment. Where the paper makes design assertions that rest on practitioner judgment rather than public sources, they are identified as such. The paper discloses no proprietary architecture, internal metrics, partner terms, or non-public implementation details; all named examples and figures are drawn from cited public sources.

Disclosure. The author has held payments and product roles in the broader commerce and retail-technology sector and therefore has professional proximity to the subject matter. This proximity is both a source of practitioner insight and a potential source of bias, and the paper is written to remain independent of any single vendor's interests. The analysis draws exclusively on publicly available information about autonomous retail, and to keep the frameworks vendor-neutral the paper: (a) relies on independent reporting (Associated Press, Reuters, CNBC) alongside vendor materials; (b) incorporates deployments from multiple providers (AiFi/Żabka, Zippin, Hudson Nonstop, Standard Cognition) as comparative evidence; and (c) presents the frameworks in vendor-neutral terms intended to apply to any autonomous retail stack.

3. THE PAPER CONTRIBUTES:

A Payment-Centric Scalability Thesis: autonomous retail scales where merchant economics, customer trust, and payment ecosystem requirements are jointly optimized, and stalls where any one of them is neglected, regardless of perception accuracy.

An Autonomous Retail Scalability Matrix: a format-level framework, with an explicit dimension rubric and aggregation rule, for evaluating where autonomous checkout is likely to succeed or fail.

A Frictionless Commerce Payments Maturity Model: a staged model from the traditional point of sale (POS) to autonomous and predictive payment orchestration, supported by a layered reference architecture emphasizing shopper identity, payment orchestration, risk and exception handling, receipt trust, and merchant reconciliation.

Framework construction procedure. For reproducibility, the construction of the Scalability Matrix followed five recorded steps. (1) Candidate scalability determinants were enumerated from the five layers of the reference architecture (Section V) plus the deployment context factors recurring in the public cases (mission timing, basket composition, operating economics). (2) Candidates were consolidated into seven dimensions by merging overlapping factors and discarding any factor that did not discriminate between at least two of the nine formats; seven is therefore not a designed constant but the count that survived consolidation. (3) An ordinal H/M/L scale was chosen deliberately over continuous scoring: the public evidence supports ordinal judgments, and interval-level scores would imply precision the evidence cannot carry. (4) Equal weighting was adopted as the least-assumptive default in the absence of an empirical basis for differential weights; it is flagged as a falsifiable assumption, and Section VII.A reports a sensitivity analysis of verdicts under alternative weightings. (5) Ratings were iterated against the construction cases only, with held-out cases reserved for the retrodictive check in Section VII.B. The maturity model followed the same procedure in abbreviated form, with levels defined by cumulative capability increments.

The conceptual validation is performed against priority formats: stadiums and arenas, airports, corporate campuses, university campuses, hospitals, convenience stores, and hotels, followed by broader retail formats with lower near-term fit.

4. DEFINING AUTONOMOUS RETAIL

Autonomous retail is a physical retail model in which item selection, basket calculation, payment initiation, and receipt generation are automated through sensing, identity, and payment infrastructure. It differs from self-checkout because the customer does not scan items at a kiosk. It differs from scan-and-go because the customer does not manually scan each item. It differs from smart carts because intelligence is embedded into the store or exit experience rather than primarily into the cart.

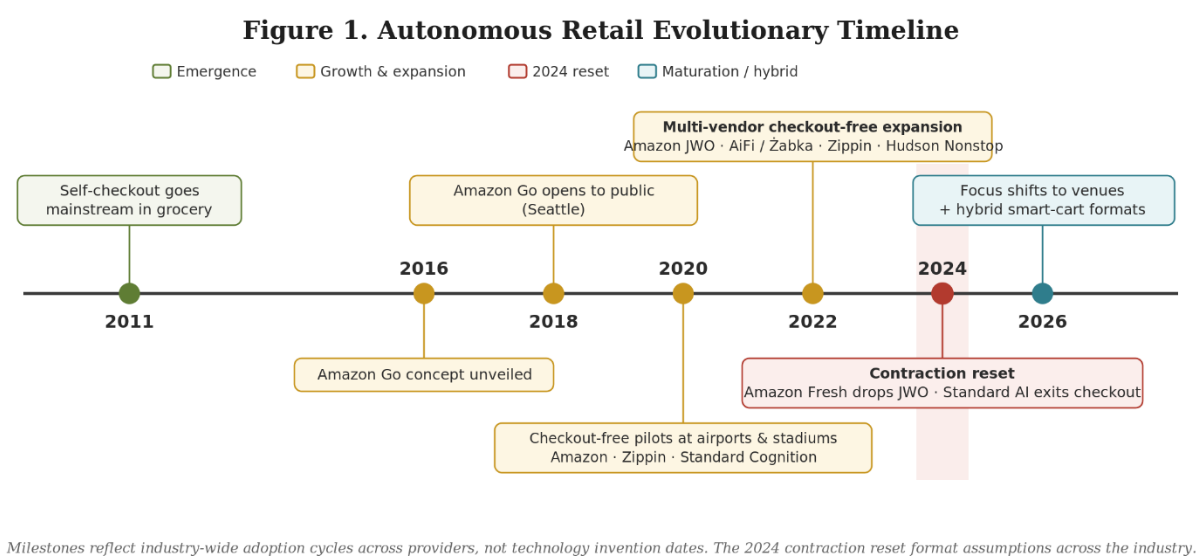

Figure 1. Autonomous retail evolutionary timeline. Milestones reflect adoption cycles rather than technology invention dates; the 2024 contraction reset format assumptions.

Table 1. Checkout model comparison

Model	Customer Work	Payment Moment	Strength	Limitation
Traditional checkout	Queue and cashier interaction	Visible terminal payment	Familiar, inclusive, controlled	Slow, labor-dependent, limited data
Self-checkout	Customer scans and pays	Visible kiosk payment	Labor leverage for small baskets	Shrink, frustration, interventions
Scan-and-go	Customer scans with app/device	Visible in-app payment	Lower capex and basket visibility	Requires scanning discipline and app adoption
Smart cart	Customer shops with instrumented cart	Cart/app payment or exit confirmation	Real-time basket visibility	Cart cost and maintenance
Autonomous checkout	Enter, select, exit	Invisible or delayed payment	Lowest checkout friction, high throughput	Payment, trust, and exception complexity

5. LITERATURE REVIEW

A. Checkout Evolution and the Limits of Self-Service

Checkout innovation has moved from cashier lanes to self-checkout, scan-and-go, smart carts, and fully autonomous stores. Self-checkout reduces cashier dependency but shifts work to shoppers and can create new loss-prevention and loyalty concerns. Drexel research reported stronger loyalty among customers using regular checkout compared with self-checkout users [12]. ECR Retail Loss Group research has also highlighted elevated loss risks in self-checkout and scan-and-go environments, especially where validation controls are weak [13]. These findings matter because autonomous retail should not be evaluated merely as "less checkout"; it should be evaluated as a redesign of control, responsibility, and trust.

Retail theft and shrink concerns further shape automation adoption, although the scale of the problem is contested: NRF reported increases in shoplifting incidents and dollar losses from 2019 to 2023 [14], but was compelled in December 2023 to retract a headline claim attributing nearly half of industry shrink to organized retail crime [19]. This paper therefore treats shrink as a directionally real but imprecisely measured pressure. The relevance for autonomous retail is twofold: retailers may seek more controlled, instrumented shopping environments, but they will reject systems that introduce new dispute, misidentification, or reconciliation risk.

B. SENSING AND AUTONOMOUS STORE TECHNOLOGY

Computer vision, RFID, shelf sensors, weight sensors, and multimodal sensing are core to autonomous retail. Recent sensing literature highlights occlusion, scalability of sensor deployment, theft prevention, and real-time processing as central technical challenges [15]. The Grab research prototype similarly showed how camera, weight, and RFID signals can be fused to associate shoppers with item interactions, while reducing compute requirements [16]. Public JWO materials emphasize AI, sensors, computer vision, RFID, and automated payment, as well as store analytics beyond ordinary POS data [1], [6].

The deployment landscape is not limited to a single vendor. AiFi-powered Żabka Nano stores constitute the largest autonomous store chain in Europe, expanding from 25 locations within six months of launch in Poland [20]. Zippin operates checkout-free stores across a substantial share of U.S. professional sports venues [21]. Any credible scalability framework must therefore generalize across vendors and geographies rather than describe one company's stack.

C. PAYMENTS AS THE UNDERSTUDIED LAYER

Most public discussion of autonomous retail centers on cameras and item recognition. Yet the payment system must decide whether the transaction is card-present, credential-on-file, wallet-tokenized, badge-linked, or closed-loop; whether authorization occurs at entry, exit, or through an estimated and incremental strategy; whether final capture is delayed; and how reversals, refunds, disputes, and settlement are reconciled. Visa's merchant guidance on estimated and incremental authorization requires that any excess between an estimated authorization and a lower final amount be released through a partial authorization reversal [8]. Scheme eligibility for estimated and incremental authorization was historically limited to designated merchant categories such as lodging, vehicle rental, and cruise, with eligibility expanded over time to additional categories [23]. Effective April 2023, Visa expanded the framework to permit merchants across categories to submit estimated and incremental authorizations for purchase transactions in which the final amount is not known when the transaction begins [25]; the requirements now appear in the published Visa Core Rules [26]. Mastercard rules similarly address processing rules such as partial approval support for prepaid and debit ranges [9]. These are not edge cases in autonomous retail; they are central design constraints.

Under the EMVCo tokenization framework, a card's primary account number is substituted with a surrogate value that stands in for it across the payment flow [10]. Tokenization is especially important when a physical store experience becomes account-bound or wallet-bound. The Payment Card Industry Data Security Standard (PCI DSS) remains relevant because autonomous retail links payment data flows with identity, telemetry, receipts, and reconciliation systems [11]. Formal verification research has, moreover, demonstrated exploitable flaws in the EMV standard itself [17], a reminder that when a card tap at entry anchors an entire shopping journey, the trustworthiness of that anchor is protocol-dependent, not assumed.

D. CUSTOMER TRUST AND CONTROL

Amazon Fresh provides an important case in format-specific fit. AP reported that customers liked skipping checkout lines but wanted to view receipts and savings as they shopped, leading to a shift toward smart carts in many U.S. grocery stores [3]. This case suggests that large-basket grocery trips require greater price and basket visibility than small-format, urgency-driven missions. An emerging empirical acceptance literature, including studies applying technology-acceptance frameworks to cashier-less formats , examines the satisfaction and trust drivers behind such behavior, and reports that perceived benefit and confidence in the system shape adoption intention [18]; the trust dimension of the framework in Section VII should be read alongside that literature rather than a single case.

The behavioral theory underlying that dimension is well established even where its application to autonomous retail is new. The technology acceptance tradition identifies perceived usefulness and perceived ease of use as adoption drivers [36], and diffusion theory adds observability and trialability [37], significant here because invisible checkout is low-observability by design: the shopper cannot watch the mechanism work, which diffusion theory predicts slows trust formation. The trust-in-automation literature frames the resulting requirement precisely: appropriate reliance depends on feedback that lets users calibrate trust to system performance [38], which in autonomous retail is exactly the role of receipts, running charge estimates, and correction flows. At the retail-systems level, omnichannel research describes the dissolution of boundaries between channels [39], and a service-dominant reading treats checkout not as a terminal transaction but as part of co-created service value [40] , both consistent with this paper's claim that payment becomes a distributed layer rather than an event.

Trust also depends on how the system's inner workings are understood. In 2024, media reporting alleged that Amazon’s autonomous retail depended heavily on offshore human review of transactions; Amazon publicly disputed the characterization, stating that human annotation supports model training rather than live checkout decisions [24]. Whatever the technical reality, the episode illustrates a structural point for this paper: when checkout becomes invisible, customers and journalists will scrutinize how charges are produced, and receipt latency, correction quality, and explanation quality become trust-critical properties of the payment layer , not customer-service afterthoughts.

6. PAYMENT-CENTRIC REFERENCE ARCHITECTURE

A payment-centric architecture treats autonomous checkout as a continuous commerce orchestration problem, not a terminal replacement. Fig. 2 depicts the five layers in shopper-journey order, with the feedback path from corrections and disputes back into orchestration and risk; the table that follows details each layer's function and its governing design questions. The architecture intentionally omits proprietary implementation detail and is described in public-domain terms.

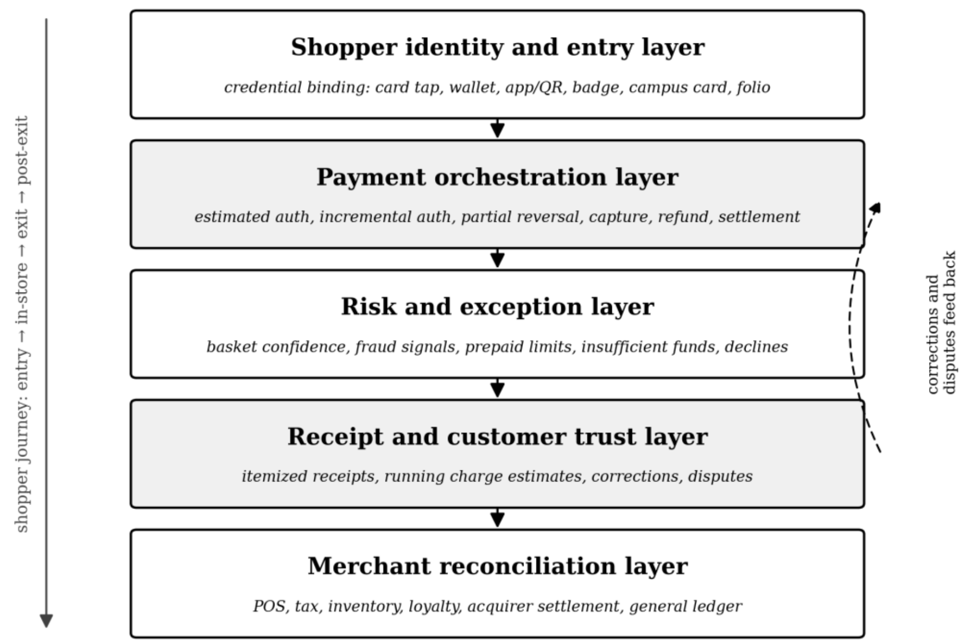

Figure 2. Payment-centric reference architecture for autonomous retail, in shopper-journey order with the correction/dispute feedback path.

Table 2. Payment-centric reference architecture

Layer	Function	Key payment and trust questions
Shopper identity and entry layer	Binds the shopper to a credential: card tap, wallet, QR/app, badge, membership, campus card, or hotel folio.	Is entry card-present, account-based, or wallet-tokenized? Can non-smartphone and prepaid users enter safely?
Payment orchestration layer	Routes authorization, token handling, incremental authorization, partial reversal, capture, refund, and settlement events.	Should the system pre-authorize at entry, authorize at exit, or estimate and adjust? Which schemes and acquirers support the needed transaction types?
Risk and exception layer	Monitors basket confidence, velocity, payment risk, fraud signals, false declines, misreads, and disputed item interactions.	How are prepaid abuse, insufficient funds, chargebacks, and friendly fraud handled without breaking the frictionless journey?
Receipt and customer trust layer	Delivers itemized receipts, corrections, self-service disputes, refunds, and explanations.	How quickly must receipts arrive? How should customers see basket,

		tax, promotions, holds, and reversals?
Merchant reconciliation layer	Reconciles autonomous carts with POS, tax, inventory, loyalty, acquirer settlement, network files, and general ledger.	Can item-level, payment-level, and settlement-level records be matched deterministically?

Boundary of the payment trust layer. Because later sections argue from this boundary, it must be stated precisely. In this paper, the payment trust layer comprises the functions that determine what a customer is charged and how the customer can anticipate, verify, correct, and contest that charge: authorization and holds, incremental and reversal handling, capture, the application of pricing, promotions, and tax to the tender amount, receipts and running charge estimates, corrections, refunds, disputes, and reconciliation. It excludes item recognition, person tracking, sensor fusion, and shopper-facing merchandising experience. Some artifacts are jointly produced: a running basket total on a smart-cart screen depends on perception for item identification, but the priced, promotion-adjusted, tax-inclusive expected charge, and its timely, trustworthy display, is a payment-layer output. The paper's claims about "payment trust" refer to this defined boundary and not to every customer-facing surface.

7. DEEP PAYMENT CHALLENGES IN AUTONOMOUS CHECKOUT

A. Card-Present And Credential-On-File Ambiguity

Autonomous stores often begin with a physical card, phone, badge, or app at entry, but the final basket may be known only after exit. This creates a classification and processing challenge: the customer performed a physical entry interaction, but the final amount may be computed later. A mature design must preserve the benefits of card-present acceptance where appropriate while ensuring that delayed calculation, credential storage, and token reuse are coded correctly under scheme, acquirer, and processor requirements. Where entry is instead app-bound or closed-loop , an account credential scanned at a gate, a stored-value or payroll-linked instrument , much of the scheme complexity in this section shifts to credential lifecycle management, account risk, and, for stored value, applicable e-money regulation. The architecture of Section V applies unchanged; the card-present analysis in this section binds only where card rails carry the tender.

B. Pre/Initial/Estimated Authorization, Incremental Authorization, and Partial Reversal

Stadiums, hotels, fuel, transit, and hospitality-like environments often use estimated authorizations and adjustments. Autonomous retail can resemble these variable-amount models because the final basket is unknown at entry. If a pre-authorized amount is too low, an incremental authorization may be needed. If it is too high, a partial reversal should release excess hold. Visa explicitly describes the requirement to reduce an estimated authorization through a partial authorization reversal when the final amount is lower [8], and scheme eligibility for these variable-amount transaction types was historically restricted by merchant category [23]; a general framework permitting estimated and incremental authorizations across merchant categories took effect only in April 2023 [25], [26]. The timing is instructive: autonomous stores had been open to the public for roughly five years before a general-purpose scheme framework existed for their transaction pattern. Extending variable-amount models to autonomous retail is therefore not a routine integration task; in the author's practitioner experience, it requires sustained engagement among merchants, acquirers, and card schemes to align transaction classification, authorization behavior, and reversal handling for a store model the rules did not originally anticipate. This is directly relevant to customer trust: holds that linger after the final basket is captured are among the most visible ways an invisible checkout can fail. The expanded framework also carries scheme pricing: approved estimated and incremental authorization requests incur a per-request fee, 0.02 percent under the 2023 Visa framework [25], [35], a recurring cost unique to this flow that enters the deployment-economics dimension of the Scalability Matrix in Section VII.

C. Scheme Support and Transaction-Type Fit

Card schemes support complex transaction types, but support is not uniform across merchant category, region, acquirer, issuer, and card product. Mastercard rules include requirements around partial approval support for prepaid and debit ranges [9]. Autonomous retail designs must therefore evaluate whether a proposed transaction flow is supported not only by the merchant's processor but by the full scheme and issuer ecosystem. A technically elegant flow can fail commercially if authorization, reversal, or partial-approval behavior is inconsistent. This assessment reflects the author's direct experience integrating autonomous checkout flows across schemes and acquirers; the pattern is corroborated by the scheme documentation cited above.

D. Prepaid Acceptance and Fraud Exposure

Prepaid cards create a distinct risk in autonomous stores. If entry is granted after a low or non-final authorization, a shopper may acquire goods whose final basket exceeds available funds. Partial approvals, insufficient funds, and issuer behavior can create gaps between access control and payment assurance; scheme rules on partial-approval support for prepaid and debit ranges define part of the available design toolkit [9]. This does not imply prepaid should be excluded; exclusion would harm accessibility. The design question, grounded in practitioner experience rather than published data , is how to construct prepaid-safe flows using authorization thresholds, real-time balance-aware acceptance where available, incremental auth support, basket limits, risk scoring, or alternative tender paths.

E. Chargebacks and Dispute Evidence

Autonomous checkout changes the evidence model for disputes. In a traditional store, the customer sees items scanned and receives a receipt immediately. In autonomous retail, the customer may challenge a charge after exit. Dispute evidence may need to include item-level logs, timestamps, cart-confidence scoring, receipt history, refund actions, and customer-facing explanation. This makes receipts and dispute UX part of the payment architecture, not merely customer support.

F. Offline Fallback and Latency

Autonomous retail often operates in environments where network conditions are constrained, including stadiums, airports, and hospitals. A resilient design must answer: can the store admit shoppers if the payment gateway is degraded? Can basket inference continue at the edge? Can receipts be delayed without harming trust? Can reversals be queued safely? The business continuity model should separate store access, basket detection, payment authorization, and customer notification so that one degraded component does not cascade into a poor checkout experience. Published scheme rules on deferred authorization and floor limits provide partial precedent for degraded-mode acceptance, though they were not written with autonomous stores in mind [26].

G. Regional Regulatory and Scheme Variation

The transaction flows above do not travel unchanged across jurisdictions, and a payments-centric framework must say so. In the European Economic Area and the United Kingdom, PSD2 strong customer authentication requirements, implemented through regulatory technical standards [30] , govern when a cardholder must actively authenticate, with cumulative limits on contactless exemptions. A card-tap entry followed by delayed basket finalization must therefore establish, per jurisdiction, where in the journey SCA applies and how incremental amounts above exemption thresholds are treated. Separately, the schemes' merchant-initiated transaction frameworks constrain how a merchant may charge a stored credential when the cardholder is no longer present, precisely the situation of post-exit capture and adjustment. Domestic debit schemes (for example, national networks outside the international card schemes) do not uniformly support estimated authorization, incremental authorization, or partial reversal at all. The scheme citations in this paper ([8], [9], [25], [26]) are drawn primarily from the international networks' U.S.-published rules; the framework's payment-complexity dimension must be re-scored per jurisdiction, and this paper's transaction-flow analysis should be read as scoped to international card scheme rails unless stated otherwise.

8. AUTONOMOUS RETAIL SCALABILITY MATRIX

A. Dimensions, Scoring Rule, and Ratings

The proposed matrix evaluates retail formats across seven dimensions. Each dimension is defined below, and each format is rated High (H), Medium (M), or Low (L) on favorability to autonomous checkout , so a format with manageable payment complexity rates H on that dimension, while a format whose shoppers demand continuous price visibility rates L on trust favorability. Boundary judgments, where the evidence sits between two levels, are recorded as M–H or L–M and scored at the midpoint.

THE SEVEN DIMENSIONS:

1. Mission urgency, how much the shopper's time pressure monetizes queue elimination.

2. Basket simplicity, item count, price variability, weighed goods, promotions, substitutions.

3. Identity certainty, how reliably a shopper can be bound to a credential at entry (badge, campus card, folio, card tap, anonymous walk-in).

4. Payment complexity favorability, whether the required transaction types (pre-auth, incremental auth, partial reversal, delayed capture, prepaid handling) are supportable in that format's payment mix.

5. Operational controllability, the operator's control over entrances, population, hours, and exceptions.

6. Deployment economics, revenue density and labor savings relative to sensor and integration capex.

7. Customer trust favorability, whether the format's shoppers accept delayed, invisible charging without continuous basket visibility.

Aggregation rule. The "Overall fit" verdict is derived, not separately judged. Ratings are scored H = 3, M = 2, L = 1 (midpoints 2.5 and 1.5 for boundary judgments), and overall fit is the unweighted mean of the seven dimension scores, mapped to bands: Very high (mean ≥ 2.70), High (2.40 ≤ mean < 2.70), Medium–high (2.00 ≤ mean < 2.40), Medium–low (1.25 ≤ mean < 2.00), and Low (mean < 1.25). One rule overrides the mean: a format rated L on any single dimension cannot exceed Medium–high regardless of its mean, because the dimensions are necessary conditions rather than fully compensatory. The override does not bind in the table below, both formats carrying an L already fall below Medium–high on mean, but it is stated as a guard for future scorings, including per-jurisdiction re-scorings of the payment dimension under Section VI.G. Equal weighting is itself a falsifiable assumption; deployment evidence indicating that some dimensions dominate should replace it.

The ratings below are structured expert judgments by the author, informed by the public deployments cited in this paper; they are offered as falsifiable claims that new deployment evidence should revise. Three ratings illustrate the evidence sensitivity: convenience-store deployment economics are rated M rather than L because Żabka Nano's operation of Europe's largest autonomous chain suggests viable unit economics in dense small formats, though the figures are operator-reported [20] , while the January 2026 closure of Amazon's first-party Go chain [31] pulls in the opposite direction, underscoring that convenience economics are operator- and location-dependent; corporate-campus payment favorability is rated M–H rather than H because closed-loop and payroll integration reduces scheme complexity while introducing integration complexity of its own; and corporate-campus trust favorability is rated M–H rather than H in view of employee privacy expectations in monitored workplaces.

Table 3. Scalability Matrix: Dimension Favorability Ratings (Structured Expert Judgment)

Format	1 Urgency	2 Basket	3 Identity	4 Payment	5 Ops control	6 Economics	7 Trust	Mean	Overall fit
Stadiums and arenas	H	H	M	M	H	H	H	2.71	Very high
Airports	H	H	M	M	H	M–H	M–H	2.57	High

Corporate campuses	M–H	H	H	M–H	H	M	M–H	2.64	High
University campuses	M	H	H	M	M–H	M	M–H	2.43	High
Hospitals	M–H	H	M	M	M	M	M	2.21	Medium–high
Convenience stores	M	H	L–M	M	M	M	M	2.07	Medium–high
Hotels	M	H	M–H	M	M–H	M	M	2.29	Medium–high
Large grocery	L	L	M	L	M	L	L	1.29	Medium–low
General merchandise	L	L–M	L–M	M	M	L–M	L–M	1.57	Medium–low

Sensitivity of verdicts to the scoring assumptions. Because equal weighting and the band thresholds are analyst choices, the verdicts were stress-tested computationally against two families of perturbation. Doubling the weight of any single dimension (seven variants, 63 format-verdict evaluations) changes 5 of 63 verdicts; removing any single dimension entirely (a harsher perturbation) changes 11 of 63. Every change is a shift of exactly one band, concentrated at band boundaries (University campuses at mean 2.43 against the 2.40 High threshold; Stadiums at 2.71 against the 2.70 Very-high threshold); no format ever moves more than one band, and the rank ordering of formats is essentially preserved. The verdicts are therefore robust to moderate weighting disagreement, while boundary cases , University campuses in particular , should be treated as High with low confidence.

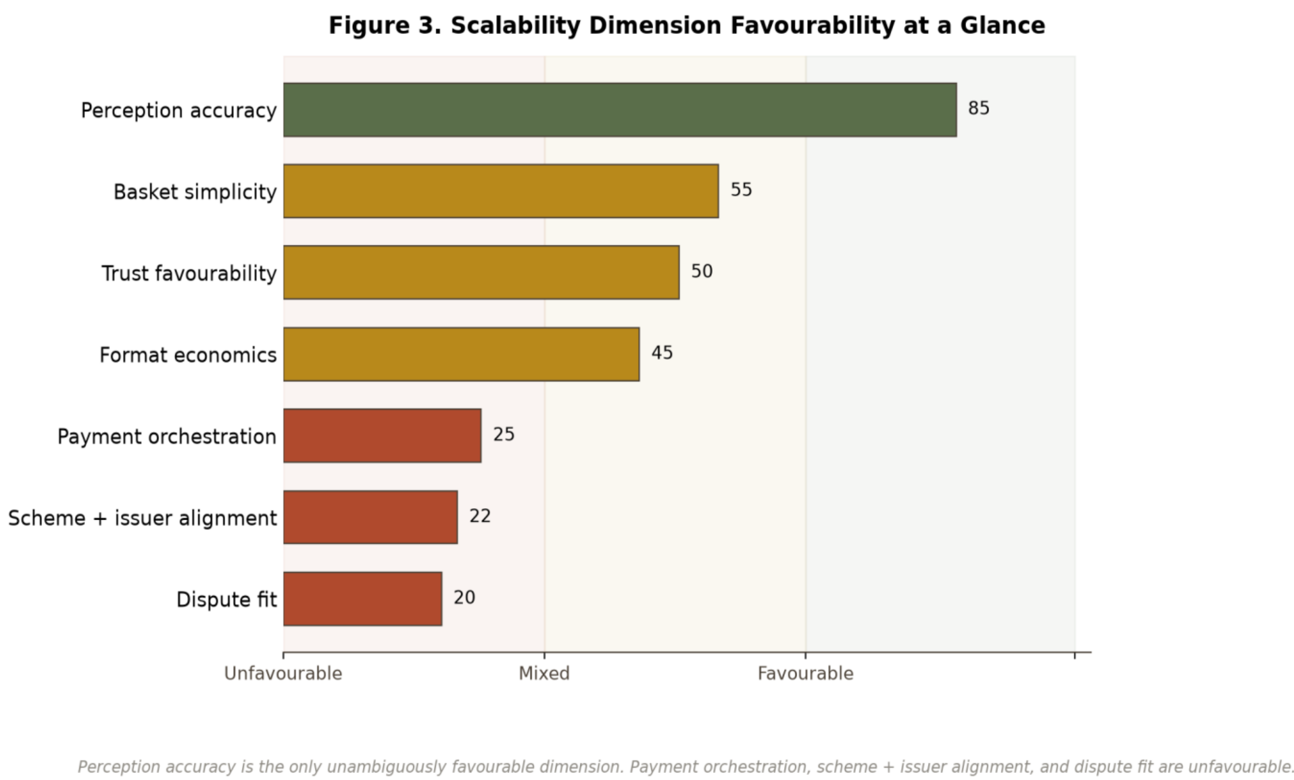

Figure 3. Scalability dimension favorability at a glance. Perception accuracy is the only unambiguously favorable dimension; payment orchestration, scheme + issuer alignment, and dispute fit are unfavorable.

Two verdicts deserve annotation. Large grocery's Medium–low rating is a statement about the present: the dimensions that suppress it (basket simplicity, payment complexity, trust favorability) are the ones hybrid models, smart carts, scan-and-go , are designed to relieve. General merchandise, previously the hardest format to rate, scores Medium–low on average, but RFID tagging can materially raise the basket and identity scores for selected categories such as apparel; the format-level verdict conceals high subcategory variance.

Table 4. Scalability matrix: format rationale and constraints

Format	Why it can scale	Primary constraints
Stadiums and arenas	Peak demand, short breaks, high revenue upside; vendor-reported Lumen Field figures cited 85% more transactions and 112% higher sales per game in a JWO location [5]; Zippin reports checkout-free stores across a large share of U.S. professional venues [21].	Alcohol controls, crowding, payment/network reliability, surge staffing and support.
Airports	Time-constrained shoppers, small baskets, premium convenience value, long operating hours.	International cards, traveler unfamiliarity, tenant/POS complexity, receipt timing.
Corporate campuses	Known users, badge identity, predictable demand, limited break windows, 24/7 amenity potential.	Subsidies, payroll/closed-loop integration, employee privacy expectations.
University campuses	Campus card/app identity, late-night food access, controlled population, small baskets.	Meal plans, age restrictions, accessibility, student trust and support.
Hospitals	Need for unattended food access for staff, visitors, and patients; labor constraints.	Mixed user identities, compliance-sensitive environment, accessibility and support needs.
Convenience stores	Small baskets and frequent trips; strong fit when labor is scarce or hours are extended; Żabka Nano demonstrates the format at chain scale in Europe [20].	Capex sensitivity, shrink, local payment mix, cash/prepaid inclusion.
Hotels	24/7 guest convenience and unmanned retail; hotel folio or card-on-file opportunities.	Guest folio integration, refunds, staff support, transient users.
Large grocery	High volume and strong friction pain point.	Basket size, produce, coupons, substitutions, family shopping, real-time spend visibility [3].
General merchandise	RFID and merchandise tagging may help selected categories.	SKU diversity, returns, theft, item tagging cost.

The matrix supports a central conclusion: autonomous checkout is most scalable where queue elimination directly monetizes limited time. It is least straightforward where shoppers require ongoing price visibility, promotion validation, budget management, and visible control.

B. Retrodictive Check Against Held-Out Cases

The deployments that informed the ratings above (Lumen Field, Żabka Nano, Amazon Fresh) cannot also validate them. As a demonstration step, the matrix is therefore checked against public cases that were not used in its construction. For the check to be capable of failing, each band must first be given an operational meaning, stated before the cases are examined.

Very high predicts sustained multi-operator expansion of the format over a three-to-five-year horizon, with failures that are site-specific rather than format-level. High predicts sustained operation with net location growth by at least one operator over a comparable horizon, without format-level retreat. Medium–high predicts conditional viability: deployments persist where the format's stated enabling conditions hold (for convenience: labor scarcity, extended unattended hours, dense traffic) and retreat where they do not , this band predicts a pattern of outcomes, not a single outcome. Medium–low predicts no sustained at-scale fully autonomous deployment without hybrid remediation of the suppressing dimensions. Low predicts no sustained deployment. Accordingly: a High format exhibiting format-level retreat within the horizon is a miss; a Medium–high format failing where its enabling conditions held is a miss; a Medium–low format scaling without hybrid remediation is a miss.

Three held-out cases were scored against these criteria. No candidate held-out case was examined and then excluded.

Amazon Go, first-party urban convenience; format rated Medium–high. The chain closed eight stores in Seattle, San Francisco, and New York in March 2023 [32], shrank by roughly half from its ~30-store peak, and was shut down entirely in January 2026 alongside the Amazon Fresh brand [31]. Scored against the conditional prediction: Go's first-party urban model largely lacked the band's enabling conditions, stores were staffed, concentrated in office districts whose weekday traffic collapsed after 2020, and did not exploit extended unattended hours. Retreat under absent enabling conditions is a pass, though the weakest kind: the conditional structure of the Medium–high band makes it the hardest band to falsify, and the same band must accommodate Żabka Nano's growth where the enabling conditions do hold. The band would have been missed had Go failed while those conditions held.

Standard AI, autonomous checkout vendor in convenience settings; format rated Medium–high. The vendor pivoted from autonomous checkout to a camera-based vision-analytics product in 2024, publicly attributing the decision to slower-than-anticipated shopper adoption and low returns on investment [33]. Pass, and the stated reason is itself informative: the constraint the vendor named is adoption and ROI, not sensor accuracy, which is the pattern the matrix's trust and economics dimensions predict.

Hudson Nonstop, checkout-free airport format; rated High. Expanded from its 2021 launch to fifteen North American airport locations by mid-2024 [34], with no format-level retreat. Pass against the High criterion of sustained net growth by at least one operator.

The January 2026 announcement [31] also bears on the framework as a whole. Amazon exited first-party autonomous retail while continuing to sell Just Walk Out to third parties. Under the matrix, this divergence is the format-specificity signature rather than a contradiction: the formats Amazon exited, first-party grocery and urban convenience, are among the matrix's weakest fits, while the formats where third-party sales continue , stadiums, airports, campuses , are its strongest. If first-party economics had been the binding constraint everywhere, third-party venue expansion should have stalled as well; it did not. The honest caveat is that operator strategy and format fit are confounded in any single company's exit, which is why multi-operator evidence (Żabka, Zippin, Hudson) carries more weight than Amazon's alone.

The check remains deliberately modest: three cases, press-level outcome attributions, author-selected for public-data availability, and scored by the same analyst who constructed the matrix. One further limitation deserves explicit statement: although the band criteria are presented before the cases in the text, they were necessarily authored with knowledge of the case outcomes, retrodiction cannot escape this. The criteria should therefore be read as retrospectively constructed benchmarks whose value lies in binding future scorings; prospective application to deployments whose outcomes are not yet known is the genuine test. The check is offered as a falsifiability demonstration and as the demonstration step of the design-science structure described in Section II. Multi-expert scoring remains future work (Section XI).

9. FRICTIONLESS COMMERCE PAYMENTS MATURITY MODEL

Table 5. Frictionless commerce payments maturity model

Level	Payment model	Capabilities	Autonomous retail implication
1	Traditional POS	Card-present terminal, cashier flow, visible receipt.	Baseline; high trust but high checkout friction.
2	Self-checkout payment	Customer-operated terminal, kiosk controls, interventions.	Lower labor but higher customer work and shrink risk.
3	App or scan-and-go payment	Credential-on-file, wallet, loyalty, self-scanning.	Lower capex but requires scanning and app adoption.
4	Autonomous payment orchestration	Entry binding, pre-auth/final auth, tokenization, reversals, itemized receipt, dispute flows.	Core target state for stadiums, airports, campuses, hospitals, hotels, and convenience.
5	Predictive and ecosystem-integrated commerce	Dynamic routing, risk-based authorization, identity federation, loyalty/tax/receipt standards, automated reconciliation.	Future state; requires standards and cross-ecosystem support.

The distinguishing feature of this model is that it does not measure payment maturity by acceptance breadth alone. It measures the ability of payments to disappear from the shopper's foreground while becoming more sophisticated in the background. Consistent with maturity-model construction guidance [29], the levels are defined as cumulative capability increments rather than technology brands, the model's application domain (payment orchestration for physical retail) is explicit, and the model is offered for comparison and iterative refinement rather than as a finished measurement instrument.

Figure 4. Frictionless Commerce Payments Maturity Model

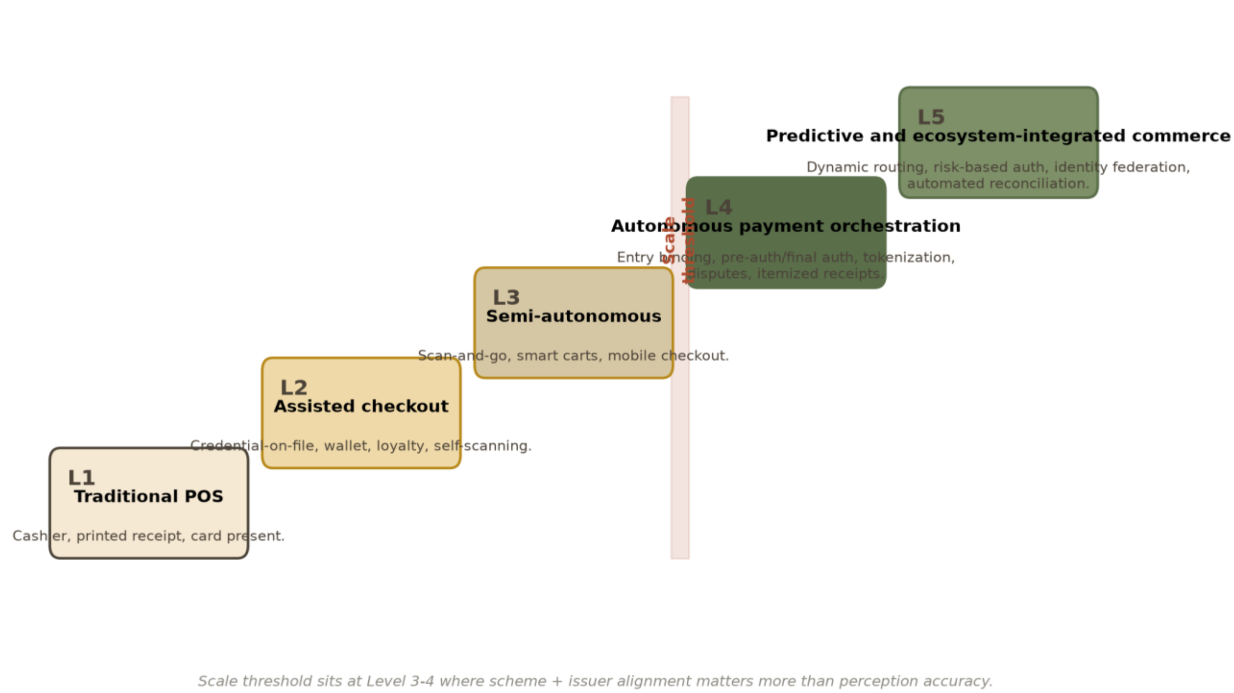

Figure 4. Frictionless commerce payments maturity model. Scale threshold sits at Level 3-4 where scheme + issuer alignment matters more than perception accuracy.

Assessment indicators. Measurable indicators can locate an organization on the model. The Level 2–3 boundary is marked by credential-on-file share and self-scan adoption rates. Level 4 is marked by the share of transactions processed through variable-amount authorization flows, receipt delivery latency, hold-release latency after final capture, dispute rate and dispute-overturn rate, and the reconciliation auto-match rate between item-level, payment-level, and settlement-level records. Level 5 would be evidenced by cross-provider dynamic routing and standardized item-level receipt exchange. Format archetypes illustrate the levels without naming organizations: a conventional supermarket lane operates at Level 1–2, a scan-and-go app retailer at Level 3, a checkout-free stadium or airport store at Level 4; Level 5 remains aspirational , no public deployment presently exhibits its full capability set.

10. DISCUSSION: WHY IT CAN SCALE AND WHY IT MAY NOT

A. Why Autonomous Retail Can Scale

Queue elimination creates measurable value in surge environments. Vendor-reported figures for JWO at Lumen Field cited strong transaction and sales uplift [5], and multi-vendor deployment growth in stadiums and European convenience formats shows the pattern is not vendor-specific [20], [21].

Hardware, AI models, and RFID approaches are maturing, enabling more format-specific deployments [6], [7], [15].

Tokenized payments, wallets, and account-based commerce make it easier to bind identity and payment at entry [10].

Labor constraints and demand for 24/7 access create use cases in hospitals, hotels, campuses, and corporate environments [2].

B. Why Autonomous Retail May Not Scale Universally

Large-format grocery requires real-time basket visibility, promotion transparency, coupon handling, variable-weight item support, and family-shopping clarity [3].

Payment ecosystems are inconsistent across card products, regions, schemes, acquirers, and issuers, especially for incremental authorization, partial approval, and prepaid handling [8], [9], [23].

Privacy and surveillance concerns can limit adoption if customers do not understand what data are collected and how charges are calculated; the 2024 human-review controversy shows how quickly opacity becomes a trust liability [24].

Integration cost can overwhelm ROI for low-traffic stores unless unattended hours or labor savings create incremental value.

Disputes and delayed receipts can erode trust quickly because the checkout ritual is invisible.

11. ALTERNATIVE EXPLANATIONS, AND WHY THEY REINFORCE THE THESIS

A critical reader may object that the most prominent retreat in the field, Amazon’s Just Walk Out's removal from most U.S. Amazon Fresh stores, has been publicly attributed to sensor economics, operating cost, and shopper experience rather than to payments [3], [4]. The objection deserves a direct answer, for two reasons.

First, this paper does not claim payments are the sole binding constraint; it claims they are a co-equal necessary condition that the perception-centric narrative omits. Sensor cost explains why a deployment is expensive. It does not explain which formats succeed at identical sensor cost, that variation tracks the payment and trust dimensions in the Scalability Matrix.

Second, the features shoppers reportedly missed in large-format grocery, running totals, visible savings, immediate receipts, confidence about what they will be charged [3] , fall largely within the payment trust layer as bounded in Section V. The claim must be stated carefully: a running total on a smart-cart screen is jointly produced, with item identification belonging to perception and the priced, promotion-adjusted expected charge belonging to the payment layer. What the Fresh case shows is that charge transparency, the customer's ability to anticipate and verify what they will pay , was binding for that format, and charge transparency is predominantly a payment-layer output under the Section V boundary. The strongest public counterexample to a payment-led reading is therefore, on inspection, substantially consistent with it, though not exclusively so.

What would falsify the thesis. The claim that payment maturity is a co-equal necessary condition is empirical, and the following evidence would count against it: (a) sustained large-scale autonomous deployments operating with only terminal-style final-amount authorization, with no variable-amount orchestration, delayed capture, or reversal handling; (b) adoption evidence showing customer acceptance is insensitive to receipt latency, authorization holds, and dispute experience; or (c) format outcomes that track sensor economics while cutting across the payment and trust dimensions of the matrix , for example, systematic failure in stadium-like formats despite favorable payment conditions, or grocery success without any charge-transparency remediation. The retrodictive cases in Section VII.B are the beginning of such a test, not its conclusion.

12. IMPLICATIONS FOR MERCHANTS, PAYMENT PROVIDERS, AND SCHEMES

Table 6. Stakeholder implications

Stakeholder	Implication
Merchants	Do not start with the technology. Start with the format: peak demand, basket complexity, labor constraints, identity model, and payment risk.
Acquirers/processors	Develop autonomous-retail transaction playbooks for pre-auth, incremental auth, reversal, delayed capture, refunds, chargebacks, and settlement.
Issuers	Improve approval logic for legitimate autonomous retail flows while protecting against prepaid abuse and unusual basket-finalization timing.
Schemes	Clarify rules and indicators for autonomous checkout transaction types, including card-present entry with delayed amount finalization.
Wallets and token providers	Support domain-controlled token use, lifecycle management, credential updates, and privacy-preserving identity binding.
Customers	Expect convenience only if receipts, corrections, and dispute flows remain transparent and easy.

13. LIMITATIONS AND RESEARCH AGENDA

This paper has six principal limitations. First, the Scalability Matrix ratings are single-analyst structured judgments; the retrodictive check in Section VII.B is a demonstration on three held-out cases, not validation, and multi-expert scoring with reported inter-rater agreement (for example, a Delphi procedure) remains the appropriate next step. Second, the public evidence base is weighted toward the most heavily reported deployments, partially offset by the additional multi-provider cases cited [20], [21], [33]. Third, several deployment figures originate in vendor or operator communications and should be treated as self-reported. Fourth, the author's professional proximity to the subject, disclosed in Section II, is both the paper's evidentiary strength and a potential source of bias. Fifth, the transaction-flow analysis is grounded primarily in the international card networks' U.S.-published rules; Section VI.G identifies, but does not resolve, jurisdictional variation in strong customer authentication, merchant-initiated transaction treatment, and domestic debit scheme support. Sixth, no quantitative performance data , approval rates, dispute rates, hold-release latencies, or deployment economics , were collected or analyzed; every quantitative figure in the paper is quoted from a cited public source, and the sensitivity analysis in Section VII.A perturbs the author's own ratings, not field measurements.

14. THESE LIMITATIONS DEFINE A RESEARCH AGENDA:

• Under what format conditions does autonomous retail generate positive value compared with self-checkout, scan-and-go, or smart carts?

• How should autonomous retail transaction flows be classified and processed when entry is physical but basket finalization is delayed?

• Which combination of estimated authorization, incremental authorization, partial reversal, delayed capture, and tokenization best balances approval rate, fraud control, prepaid inclusion, and customer trust?

• How do receipt timing, real-time basket visibility, and self-service disputes affect customer adoption of invisible checkout?

• Can the Autonomous Retail Scalability Matrix predict which retail formats scale and which require hybrid models , and can multi-expert scoring (a Delphi panel with reported inter-rater agreement) confirm or revise its ratings?

• Can the matrix survive prospective application, scoring announced deployments with pre-registered band predictions before their outcomes are known, which is a strictly stronger test than the retrodiction of Section VII.B?

• What standards are needed across identity, payment tokens, receipts, item-level logs, and settlement reconciliation to reduce merchant integration burden?

15. CONCLUSION

Autonomous retail is best understood as a transformation of the checkout boundary. The customer no longer experiences payment as a final action; instead, payment is distributed across entry, identity, basket inference, risk scoring, authorization, receipting, dispute handling, and settlement. This creates a powerful convenience proposition but also raises the standard for payment reliability and transparency.

The paper's contribution is the claim that autonomous retail's future will be determined by payment-centric scalability as much as by perception accuracy. Stadiums, arenas, airports, campuses, hospitals, convenience stores, and hotels are strong near-term fits because speed, access, and throughput matter more than continuous basket control. Large grocery and general merchandise will likely adopt hybrid models until real-time visibility, economics, and payment certainty improve. Autonomous retail can scale, but only where the invisible payment layer is mature enough to make frictionless commerce trustworthy.

16. DISCLOSURE STATEMENT

The author has held payments and product roles in the commerce and retail-technology sector. This paper uses only publicly available information about autonomous retail; it discloses no proprietary architecture, internal metrics, partner terms, or non-public implementation details. The views expressed are the author's own.

17. ACKNOWLEDGMENT

The author thanks the anonymous reviewers for their constructive feedback that improved this manuscript.

REFERENCES

1. Amazon Web Services, "Just Walk Out Technology," public product page, accessed July 1, 2026.
2. Just Walk Out technology by Amazon, "What is Just Walk Out technology, and how does it work?" Feb. 26, 2026.
3. Associated Press, "Amazon is removing Just Walk Out technology from its Fresh grocery stores in the US," Apr. 2, 2024.
4. Reuters, "Amazon to push cashierless shopping tech into more third-party stores, while backing off itself," Apr. 17, 2024.
5. About Amazon, "Amazon's Just Walk Out tech rolls out to more stadiums and universities," Sept. 3, 2024.
6. Just Walk Out technology by Amazon, "Computer vision in retail: the technology behind frictionless checkout," Feb. 26, 2026.
7. Amazon, "Just Walk Out's new portable lanes bring checkout-free shopping to more event formats," Jan. 13, 2026.
8. Visa, "Authorization and Reversal Processing: Best Practices for Merchants," 2024.
9. Mastercard, "Transaction Processing Rules," Dec. 9, 2025.

10. EMVCo, "EMV Payment Tokenisation," public specification overview, accessed July 1, 2026.
11. PCI Security Standards Council, "PCI Data Security Standard," public standards overview, accessed July 1, 2026.
12. Drexel University LeBow College of Business, "Does Self-Checkout Impact Grocery Store Loyalty?" Jan. 4, 2024.
13. ECR Retail Loss Group, "Self-checkout loss in retail: Global use, impact, and control," Mar. 8, 2022.
14. National Retail Federation, "The Impact of Retail Theft & Violence 2024," Dec. 17, 2024.
15. S. Rukundo et al., "A Survey of Challenges and Sensing Technologies in Autonomous Retail Systems," arXiv:2503.07997, 2025.
16. X. Liu et al., "Grab: Fast and Accurate Sensor Processing for Cashier-Free Shopping," arXiv:2001.01033, 2020.
17. D. Basin, R. Sasse, and J. Toro-Pozo, "The EMV Standard: Break, Fix, Verify," arXiv:2006.08249, 2020.
18. C. D. Schultz, "Customer benefits and self-service satisfaction in cashierless shopping systems," Journal of Retailing and Consumer Services, 2025.
19. CNBC, "Retail lobby group's retraction of key crime claim shows how hard it is to track theft," Dec. 7, 2023.
20. Żabka Group, "Żabka Nano autonomous stores powered by AiFi outpace industry with 25 live locations," Jan. 2022.
21. Stadium Tech Report, "Zippin adds five checkout-free stores at MLB stadiums," 2025, accessed July 2026.
22. About Amazon UK, "How to shop at Amazon Fresh stores and why you should try it," accessed July 2026.
23. Visa, "Expanded Eligibility for Estimates and Incremental Authorization Requests," Article ID AI09108, 2019.
24. Retail Technology Innovation Hub, "Amazon hits out at report that Just Walk Out technology relies on human reviewers watching from afar," Apr. 17, 2024.
25. PayPal Braintree, "All Spring 2023 Network Updates: Visa, Expansion of Authorization Framework," developer documentation, 2023, accessed July 2026.
26. Visa, "Visa Core Rules and Visa Product and Service Rules," published public edition, Apr. 2026.
27. A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design Science in Information Systems Research," MIS Quarterly, vol. 28, no. 1, pp. 75–105, 2004.
28. K. Peffers, T. Tuunanen, M. A. Rothenberger, and S. Chatterjee, "A Design Science Research Methodology for Information Systems Research," Journal of Management Information Systems, vol. 24, no. 3, pp. 45–77, 2007.
29. J. Becker, R. Knackstedt, and J. Pöppelbuß, "Developing Maturity Models for IT Management: A Procedure Model and its Application," Business & Information Systems Engineering, vol. 1, no. 3, pp. 213–222, 2009.
30. European Commission, "Commission Delegated Regulation (EU) 2018/389 with regard to regulatory technical standards for strong customer authentication and common and secure open standards of communication," Official Journal of the European Union, 2018.
31. B. Dworski, "Amazon Go is shutting down," C-Store Dive, Jan. 27, 2026.
32. GeekWire, "Amazon closing eight Amazon Go convenience stores in latest tightening of its physical retail," Mar. 2023.
33. CSP Daily News, "Standard AI Moves Away From Autonomous Shopping and Toward AI Cameras," 2024.
34. Retail Technology Innovation Hub, "Amazon Just Walk Out tech powered Hudson Nonstop store hits Washington Dulles International Airport," Jun. 18, 2024.
35. Verisave, "April 2023 Card Brand Rate and Rule Changes," 2023.
36. F. D. Davis, "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology," MIS Quarterly, vol. 13, no. 3, pp. 319–340, 1989.
37. E. M. Rogers, Diffusion of Innovations, 5th ed. New York: Free Press, 2003.
38. J. D. Lee and K. A. See, "Trust in Automation: Designing for Appropriate Reliance," Human Factors, vol. 46, no. 1, pp. 50–80, 2004.
39. P. C. Verhoef, P. K. Kannan, and J. J. Inman, "From Multi-Channel Retailing to Omni-Channel Retailing: Introduction to the Special Issue on Multi-Channel Retailing," Journal of Retailing, vol. 91, no. 2, pp. 174–181, 2015.
40. S. L. Vargo and R. F. Lusch, "Evolving to a New Dominant Logic for Marketing," Journal of Marketing, vol. 68, no. 1, pp. 1–17, 2004. [1]The views expressed in this work are those of the author and do not necessarily reflect the views of any current or former employers.