JOURNAL OF GLOBAL TRENDS IN SOCIAL SCIENCE

J G T S S

Review article

https://doi.org/10.70731/1178nq90

Privacy-Preserving and Trustworthy AI Infrastructures for Digital Commerce: Federated Learning, Cross-Channel Measurement, and Social Advertising

Ahmad Zulkifli Bin Idris [1], **Weiling Tan** [2,*], **Kavitha Rajendran** [3]

Received 19 January 2026

Accepted 23 March 2026

Published 31 March 2026

Abstract: Digital commerce is shifting from unconstrained data accumulation to a context in which privacy regulation, platform restrictions, and public distrust shape how artificial intelligence is designed and deployed. This review argues that privacy-preserving and trustworthy commerce AI should be understood as an infrastructure problem rather than a set of isolated model-level improvements. It synthesizes research on privacy-preserving machine learning, digital advertising measurement, recommender systems, social commerce, creator monetization, platform architecture, and AI governance, and proposes an analytical framework structured around four layers: data topology, learning protocol, measurement logic, and governance architecture. At the data layer, consumer traces are fragmented across merchants, platforms, creators, and devices, making centralized modeling increasingly costly, risky, and legally fragile. At the learning layer, federated learning, differential privacy, secure aggregation, and related techniques enable distributed training and protected analytics, but introduce trade-offs in accuracy, communication, personalization, and security. At the measurement layer, the erosion of third-party identifiers makes privacy-preserving attribution and audience matching central to campaign optimization. At the governance layer, multi-tenant architecture, API standardization, auditability, and risk management determine whether such systems can operate at scale, especially for SMBs. The review concludes that the future of commerce AI will depend less on prediction alone than on the ability to institutionalize privacy, transparency, and accountability within commercially viable infrastructures.

Keywords: Digital commerce; Federated learning; Differential privacy; Secure aggregation; Social advertising; Creator economy; Recommender systems; AI governance

ISSN 2759-7830 (Online)
ISSN 2760-2508 (Print)

Introduction

The technical and economic logic of digital commerce has changed. This transition is also part of a broader restructuring of digital and social media marketing research, where data access, platform governance, and AI-mediated personalization have become central rather than peripheral concerns (Dwivedi et al., 2021). For roughly two decades, the dominant paradigm assumed that more data centralization would reliably produce better personalization, more precise advertising, stronger attribution, and lower customer acquisition cost. Recommendation engines, bidding systems, and marketing analytics were therefore designed around the continuous extraction and fusion of behavioral traces across sites, devices, and application contexts. That model is now under pressure from multiple directions. Privacy regulation has expanded the compliance burden surrounding personal data processing; platform policies have restricted tracking and identifier sharing; consumers have become more aware of surveillance-oriented advertising; and firms increasingly face the reputational and operational costs of building AI systems on data practices that are difficult to justify or audit (Boerman et al., 2017; Dinev et al., 2013; Morimoto, 2021; Truong et al., 2021). As a result, digital commerce has moved into a stage in which the central challenge is no longer only how to optimize

[1] University of Malaya, Kuala Lumpur 50603, Malaya; [2] Monash University Malaysia, Selangor 47500, Malaysia; [3] University of Technology Malaysia, Johor 81310, Malaysia.
* Corresponding author. Email: weiling.tan@monash.edu.

AI models, but how to construct AI infrastructures that remain effective under privacy, governance, and interoperability constraints.

This shift is especially consequential because digital commerce is structurally heterogeneous. Data relevant to a single commercial decision can be distributed across retailers, marketplaces, social platforms, creators, payment intermediaries, analytics providers, and cloud services. A recommendation model may depend on on-site browsing, purchase history, inventory states, content metadata, and social signals. An advertising system may combine audience estimation, creative selection, attribution, fraud detection, and budget allocation. A creator monetization workflow may require audience matching, engagement prediction, content moderation, revenue allocation, and contractual auditability. In each case, the technical object is not a single model but a socio-technical pipeline embedded in a wider platform ecosystem (Hein et al., 2020; Mukhopadhyay & Bouwman, 2019). This is why privacy-preserving AI in commerce cannot be understood merely as the addition of a privacy mechanism to an otherwise unchanged system. It entails redesigning the infrastructure through which data, incentives, models, and accountability are coordinated.

Research on trustworthy AI has reinforced this broader view. Surveys and policy-oriented syntheses increasingly argue that trustworthy AI is not exhausted by fairness or explainability at the model level; it also includes robustness, traceability, responsibility allocation, and risk management across the system life cycle (Corrêa et al., 2023; Díaz-Rodríguez et al., 2023; Laux et al., 2024; Lewis & Moorkens, 2020; Tabassi, 2023). In digital commerce, these issues are intensified by commercial imperatives. Firms want more granular measurement, faster experimentation, and scalable personalization, while regulators and users demand minimization, transparency, and control. The resulting tension cannot be resolved by normative statements alone. It requires architectures capable of combining privacy preservation with commercially usable outputs.

Federated learning has become a prominent candidate because it allows multiple parties or devices to train shared models without directly pooling raw data (Cheng et al., 2020; Kairouz et al., 2021). Yet federated learning by itself does not solve the underlying commerce problem. Distributed training still leaves questions about leakage from gradients, participant heterogeneity, personalization quality, verification, content governance, and downstream measurement. Differential privacy can limit leakage, but often at the cost of utility (Abadi et al., 2016; Wei et al., 2020). Secure aggregation can reduce exposure of participant updates, but it complicates orchestration and operational resilience (Bonawitz et al., 2017). Zero-knowledge proofs and related cryptographic mechanisms can improve verification and accountability, but they introduce complexity and latency that many commerce systems are not designed to absorb (Sun et al., 2021). In short, privacy-preserving AI for commerce is modular, not monolithic.

The importance of this problem is amplified by two business trends. First, commerce journeys are now inherently cross-channel. Consumers move between search, social feeds, creator content, marketplaces, merchant websites, loyalty applications, and physical channels. Measurement and attribution have therefore become more difficult at exactly the moment when identifiers are less stable. The classical attribution problem, already difficult in pay-per-conversion settings, has become more severe in privacy-constrained environments (Jordan et al., 2011). Second, small and medium-sized businesses increasingly rely on external AI infrastructures rather than internal data science teams. This expands the relevance of multi-tenant architectures, standardized recommendation and measurement APIs, and platform-level governance mechanisms that can lower adoption barriers while embedding compliance and auditability into default workflows (Arnold et al., 2022; Gieß & Hutterer, 2025; Xue et al., 2019).

Against this background, the present review treats privacy-preserving and trustworthy commerce AI as an integrated infrastructure domain. It does not approach recommendation, advertising, attribution, creator monetization, and SMB enablement as separate literatures that only happen to share some methods. Instead, it examines how these functions become interdependent once privacy constraints limit frictionless data pooling. The review makes three contributions. First, it synthesizes the technical stack of federated learning, differential privacy, secure aggregation, and zero-knowledge verification in terms of their relevance for commerce rather than in purely generic machine learning terms. Second, it links privacy-preserving learning to cross-channel measurement and social advertising, two domains where commercial value depends on the controlled recovery of signal from partially observable environments. Third, it shows why trustworthy AI in commerce must be operationalized at the platform architecture and governance level, especially for SMB-focused ecosystems.

The rest of the paper follows this logic. Section 2 defines the main commerce scenarios and explains why they generate distinct data topologies and privacy risks. Section 3 reviews the core privacy-preserving technical stack. Section 4 addresses cross-channel measurement and attribution in a post-cookie environment. Section 5 examines social commerce advertising and the creator economy. Section 6 analyzes trustworthy AI platforms for SMBs. Section 7 discusses governance, standards, and accountability. Section 8 concludes with a research agenda centered on the balance between privacy, utility, and commercial implementability.

Scenario Decomposition

Digital commerce is frequently described as a unified domain, but the infrastructural requirements of its main AI tasks differ sharply. A useful review must therefore begin by distinguishing scenarios according to their decision objects, data topology, latency requirements, and regulatory exposure. Five scenarios are especially central: recommendation, advertising delivery and targeting, cross-channel measurement and attribution, creator monetization, and platformized AI services for SMBs.

Recommendation remains the most mature commerce AI application. It includes product ranking, bundle suggestion,

personalized search, media selection, and retention-oriented recommendation. The modern literature emphasizes increasingly complex representations, including graph neural networks, cross-domain transfer, and self-supervised learning (Chen et al., 2025; Gao et al., 2023; Sharma et al., 2024). However, the privacy problem in recommendation is unusually acute because consumer preference data are both economically valuable and behaviorally intimate. Historical work on privacy-preserving recommender systems focused on anonymization, obfuscation, or cryptographic computation (Checco et al., 2017; Sweeney, 2002; Wang et al., 2018). More recent work shifts toward federated recommendation, where interaction data remain local to users, enterprises, or domains while models are coordinated centrally or semi-centrally (Asad et al., 2023; Kalloori & Klingler, 2021; Luo et al., 2024; Qin et al., 2021). The central trade-off is that preserving locality often worsens sparsity, non-IID data problems, and personalization difficulty.

Advertising delivery and targeting form a second scenario. Here the key task is not simply ranking items but matching audiences, contexts, and creatives under auction-like or budget-constrained conditions. Behavioral advertising research has long shown that performance depends on the availability of reliable behavioral features and identity linkage, but these same features trigger privacy concern, persuasion resistance, and trust deterioration when consumers perceive tracking as intrusive (Boerman et al., 2017; Carlson et al., 2022; Jiang et al., 2013). In social media settings, personalization also interacts with platform-specific motivations and trust perceptions, which means that advertising effectiveness depends not only on the accuracy of targeting but on the legitimacy of the targeting process itself (Carlson et al., 2022; Morimoto, 2021). Privacy-preserving advertising therefore requires both a technical solution to audience estimation and an institutional solution to perceived manipulation.

Cross-channel measurement and attribution constitute a third scenario and should not be reduced to a reporting problem. In digital commerce, firms optimize campaigns, recommendations, and incentives on the basis of measured downstream outcomes. When user journeys span multiple channels, devices, and actors, the absence of stable identifiers creates measurement loss: some conversions cannot be linked, some touchpoints cannot be sequenced, and some causal contributions cannot be isolated. The attribution literature already identified the difficulty of assigning value across multiple ad exposures and interactions even before current privacy restrictions (Jordan et al., 2011). The contemporary challenge is more severe because signal loss is now built into the environment. Privacy-preserving conversion measurement, secure aggregation of event counts, and differentially private reporting are increasingly central to campaign optimization rather than supplementary safeguards (Delaney et al., 2024; Farahat et al., 2009; Zhong et al., 2022).

Creator monetization forms a fourth scenario, especially within social commerce and influencer-led retail. In this setting, the economic unit is often neither the platform alone nor the merchant alone, but a triangular relationship among creators, audiences, and commercial sponsors. The relevant AI tasks include audience matching, engagement prediction, conversion estimation, content recommendation, fraud or manipulation detection, and revenue allocation. Research on influencer and live-streaming commerce shows that consumer response is mediated by credibility, parasocial relationships, platform affordances, and the perceived authenticity of content (Bargoni et al., 2023; Bi & Zhang, 2023; Libai et al., 2025; Teng et al., 2022; Xue & Liu, 2023). This complicates privacy-preserving design because the system must often infer high-value audience segments from interactional signals that are socially embedded and partly creator-specific. The result is a need for architectures that can support monetization without forcing creators to surrender comprehensive audience-level data to centralized intermediaries.

The fifth scenario is AI-as-infrastructure for SMBs. Most small firms cannot build proprietary recommender systems, privacy engineering pipelines, or governance functions. They rely on platforms that expose standardized APIs for recommendation, campaign management, analytics, catalog enrichment, or content moderation. From a systems perspective, SMB adoption depends on whether these services can be offered in a multi-tenant form that is sufficiently modular, interoperable, and compliant. Research on digital platform ecosystems shows that value creation depends on orchestrating complementors through shared interfaces and governance rules rather than merely providing software functionality (Hein et al., 2020; Mukhopadhyay & Bouwman, 2019; Xue et al., 2019). Recent platform research also highlights the importance of data architecture, modularity, and governance in shaping who can participate and under what conditions (Arnold et al., 2022; Gieß & Hutterer, 2025). Trustworthy AI for SMBs is therefore not only a matter of offering "responsible" models; it is a matter of packaging compliance, auditability, and privacy guarantees into accessible infrastructure.

These scenarios share some common technological primitives, but their data topologies differ. Recommendation often involves horizontally partitioned user-interaction data or cross-domain preference transfer. Advertising may combine platform-side audience estimates with merchant-side conversion signals. Attribution connects event logs generated across multiple systems. Creator monetization adds relational data between creators and audiences, as well as platform-specific content signals. SMB services require multi-tenant separation, tenant-specific policy controls, and standard interfaces that can bridge heterogeneous tools. Vertical federated learning is relevant when distinct organizations hold different features about overlapping user sets, while horizontal federated learning is more natural when participants hold similar features over disjoint populations (Khan et al., 2025; Kalloori & Klingler, 2021). Cross-silo and cross-device settings therefore have different security and orchestration implications.

A scenario-based view also clarifies why the privacy question cannot be resolved by abstract "compliance" language. The same privacy mechanism can be appropriate in one scenario and inadequate in another. For example, local differential privacy may be tolerable for aggregate measurement but too destructive for personalized recommendation. Secure aggregation may protect participant updates in cross-silo learn-

Table 1 | Core digital commerce scenarios and their infrastructural implications

Scenario	Main AI task	Data topology	Primary privacy risk	Main infrastructure implication
Recommendation	Ranking, retrieval, personalization	User- or domain-local interaction data	Preference leakage, re-identification, profiling	Federated recommendation, privacy-aware personalization
Social advertising	Audience matching, bidding, creative optimization	Platform-side behavioral traces plus merchant outcomes	Opaque targeting, cross-context tracking	Privacy-preserving targeting and reporting
Cross-channel measurement	Attribution, lift analysis, conversion counting	Event logs across channels and devices	Linkage risk, unverifiable conversion claims	Secure aggregation, DP reporting, auditable measurement
Creator monetization	Matching sponsors, audiences, and content	Creator-specific audiences plus platform interaction graphs	Centralized exposure of audience value	Verifiable revenue sharing and protected audience analytics
SMB platform services	Model hosting, recommendation APIs, content governance	Multi-tenant enterprise data and policy settings	Tenant leakage, inconsistent compliance	Trustworthy multi-tenant architecture and standardized APIs

ing but still leave unresolved questions about malicious clients or biased participation. Zero-knowledge proofs may be particularly valuable where revenue sharing or creator compensation requires verifiable accounting, but less necessary in low-stakes catalog ranking. A useful commerce architecture must therefore map privacy tools to scenario-specific functional requirements. **Table 1** summarizes the scenario logic used in this review.

The remainder of the review builds on this scenario differentiation by examining how privacy-preserving technologies can be assembled into commerce-specific infrastructures.

Privacy-Preserving Technical Stack

The privacy-preserving technology stack relevant to digital commerce is best understood as layered. Federated learning determines where model training occurs and how updates are coordinated. Differential privacy constrains what can be inferred from outputs or intermediate updates. Secure aggregation protects the visibility of participant-level updates during coordination. Zero-knowledge proofs and related cryptographic techniques support verification and accountability. Each layer addresses a different failure mode, and none is sufficient on its own.

Federated learning as an infrastructural rather than purely algorithmic choice

Federated learning emerged as a response to the concentration of data in centralized machine learning pipelines. Instead of moving raw data to a single repository, participants compute local updates that are then aggregated into a shared model (Cheng et al., 2020; Kairouz et al., 2021). This basic idea is attractive for commerce because merchants, platforms, creators, and user devices often have incentives not to share raw data. Federated learning can preserve local data residency while still enabling collective model improvement.

Yet the relevance of federated learning to commerce depends on deployment form. Cross-device federated learning is suitable when the primary participants are consumer devices, for example in on-device ranking or personalized content selection. Cross-silo federated learning fits settings where participants are institutions, such as merchants collaborating with an advertising network or multiple brands joining a retail media platform. Horizontal and vertical federated learning further distinguish whether parties hold similar feature spaces over different users or different feature spaces over overlapping users (Kalloori & Klingler, 2021; Khan et al., 2025). These distinctions matter because the operational problems differ. Cross-device settings face high churn, weak trust assumptions, and device heterogeneity. Cross-silo settings face stronger governance and contractual issues, but usually enjoy more stable participation and richer local computation.

For recommender systems, federated learning is appealing because interaction histories are especially privacy-sensitive. Surveys show rapid growth in federated recommendation, including matrix factorization, neural recommendation, graph-based models, and personalized federated strategies (Asad et al., 2023; Kalloori & Klingler, 2021; Luo et al., 2024; Qin et al., 2021). At the same time, several technical obstacles remain persistent. First, user behavior data in commerce are highly non-IID. Consumers differ across language, category preference, price sensitivity, and channel use, which makes global models unstable. Second, sparse and long-tail item spaces create difficulties when local clients only observe tiny slices of the inventory. Third, recommendation quality often depends on rich side information, some of which may be distributed across merchants, platforms, and content systems rather than concentrated on a single device. These problems mean that federated learning must often be combined with personalization layers, cross-domain transfer, or graph-aware representations (Chen et al., 2025; Gao et al., 2023; Sharma et al., 2024).

Federated learning also creates new attack surfaces. Even when raw data stay local, model updates can leak information, and malicious participants can poison training or infer sensitive attributes. Surveys consistently identify gradient leakage, inference attacks, poisoning, backdoors, and free-riding as major risks (Gosselin et al., 2022; Jere et al., 2021; Mothukuri et al., 2021; Papadopoulos et al., 2021). For commerce applications, these risks are not only technical. An adversarial merchant could manipulate shared recommendation training; a platform participant could attempt to infer the strategic value of another tenant's audience; or a creator-side

application could introduce distorted engagement signals. Consequently, federated learning should be seen not as a privacy guarantee but as a controlled redistribution of where risk appears.

Differential privacy and the pricing of privacy loss

Differential privacy provides a mathematically explicit way to limit the impact of any single record on released outputs or model parameters (Dwork, 2006). In practice, it is relevant to commerce for three broad reasons. First, it can reduce the likelihood that recommendation, advertising, or measurement outputs expose identifiable behavior. Second, it provides a language of privacy budgets that can be documented and governed. Third, it allows organizations to communicate formal guarantees in settings where informal assurances about "anonymization" have become unreliable.

The introduction of differential privacy into deep learning made these ideas operational for modern models, although often with substantial utility trade-offs (Abadi et al., 2016). In federated settings, the utility-privacy trade-off is further complicated by client heterogeneity, communication limits, and the fact that noise may be inserted at local or aggregate stages (Wei et al., 2020). For commerce, the key question is not whether differential privacy can be added, but where in the pipeline it is most valuable. Local differential privacy offers stronger participant-side protection but often destroys fine-grained signal needed for personalization. Central differential privacy, applied after secure aggregation, can preserve more utility but requires stronger trust in the aggregation server or protocol. Aggregate reporting tasks such as campaign measurement may tolerate higher noise than item-level ranking or dynamic pricing decisions.

The privacy budget perspective is particularly useful for commerce governance. Recommendation, targeting, and measurement are not one-time computations; they are recurring processes that consume privacy budget across repeated releases, experiments, and reporting intervals. A firm that runs continuous campaign attribution, creator analytics, and user segmentation needs an accounting framework for cumulative privacy loss. This links privacy engineering to platform governance: privacy budgets are not only mathematical objects but resource-allocation decisions shaped by business priorities. Surveys on privacy-preserving federated learning show that practical deployments must decide which outputs are worth preserving at higher fidelity and which can absorb more noise (Gu et al., 2023; Truong et al., 2021). In that sense, differential privacy is also a managerial instrument.

Secure aggregation and the hidden middle layer of trust

Secure aggregation protects the confidentiality of individual participant updates during federated coordination. The canonical design goal is that the server should learn only the aggregate sum of client updates, not the update of any single client (Bonawitz et al., 2017). In commerce, this is critical when multiple organizations jointly train models or contribute outcome events. Without secure aggregation, federated learning may offer only superficial privacy because the coordinating party can inspect participant-level gradients.

Secure aggregation is especially relevant to cross-channel measurement. If advertisers, platforms, merchants, and analytics providers each contribute event signals, a naïve aggregation system can expose partner-level conversion patterns or strategic performance information. Privacy-preserving event reporting and frequency management were already recognized in earlier online advertising work (Farahat et al., 2009). The current environment strengthens the case for secure aggregation because direct identifier-level matching is both harder and more controversial. Measurement systems increasingly need to recover aggregate signal while hiding participant-specific detail.

However, secure aggregation is often treated as a neutral technical layer when it is actually a design choice with governance implications. The protocol determines fault tolerance, dropout handling, communication burden, and the point at which trust is centralized. In cross-silo commerce systems, these operational details affect participation. Small merchants or creators may not tolerate protocols that are too costly or brittle. Thus, secure aggregation should be evaluated not only in cryptographic terms but in infrastructural terms: who can realistically join, who controls orchestration, and what evidence of correct execution is available.

Zero-knowledge proof, verification, and accountable commerce AI

Zero-knowledge proofs (ZKPs) are usually discussed in blockchain contexts, but their relevance to commerce AI lies more broadly in verifiable computation and accountability. A zero-knowledge protocol allows one party to prove that a statement is true without revealing the underlying secret information (Sun et al., 2021). In commerce infrastructures, this can support claims such as: a conversion count was computed according to an agreed protocol; a revenue share was allocated using an approved formula; a participant satisfied eligibility conditions without disclosing raw data; or a model-serving process complied with a defined policy rule.

This is particularly salient in creator monetization and federated advertising. Creators often depend on platform-provided analytics to assess sponsorship value and payout fairness. Merchants depend on platforms to report audience quality and conversions. When raw logs cannot be disclosed for privacy or proprietary reasons, verification becomes difficult. ZKPs do not solve all of these problems, but they provide a route toward auditable claims without full data exposure. Their main limitation is practical: proof generation and verification impose computational costs, and integration into real-time or near-real-time systems is still complex.

Architectural composition and trade-offs

The main analytical point is that commerce systems should not select privacy tools one by one in isolation. The functional unit is a composed stack. Federated learning addresses data locality; differential privacy constrains inferential exposure; secure aggregation protects intermediate coordination; ZKPs support verification; and API-level governance

Table 2 | Privacy-preserving building blocks for commerce AI

Building block	Primary function	Commerce use cases	Main strengths	Main limitations
Federated learning	Distributed model training	Recommendation, audience modeling, shared fraud detection	Preserves data locality; enables multi-party learning	Non-IID data, communication overhead, attack surface
Differential privacy	Formal privacy guarantee through controlled noise	Reporting, analytics, recommendation training, measurement	Quantifiable privacy budget; useful for governance	Utility loss; difficult budget allocation across repeated tasks
Secure aggregation	Hides participant-level updates during coordination	Federated training, conversion counting, partner reporting	Reduces visibility of local updates	Protocol complexity; orchestration burden; dropout handling
Zero-knowledge proof	Verifiable computation without revealing raw data	Payout verification, conversion claims, compliance checks	Supports auditability under data minimization	Computation cost; integration complexity
Standardized APIs and policy layers	Controlled exposure of model and analytics functions	SMB enablement, platform orchestration, tenant governance	Scalability, interoperability, embedded compliance	Requires mature governance and version control

structures determine how these components are exposed to tenants and partners. Research on privacy-preserving recommendation-as-a-service and privacy-aware commercial AI points toward this compositional logic, even when implementation details vary (Wang et al., 2018; Papadopoulos et al., 2021). **Table 2** summarizes the main roles and trade-offs of these primitives.

The literature suggests that utility-preserving privacy in commerce will not come from maximizing any single technique. Instead, it will come from designing a stack whose components are matched to the information requirements and trust assumptions of specific scenarios. This point becomes clearer once we turn from training infrastructure to measurement infrastructure.

Cross-Channel Measurement and Attribution in A Post-Cookie Environment

Cross-channel measurement is the point at which privacy preservation most visibly collides with commercial decision making. Recommendation and targeting models are valuable because they influence outcomes, but firms cannot justify continued investment without measurement. Once identifiers become unstable and event-level linkage becomes constrained, campaign optimization, budget allocation, and channel evaluation all degrade. The core question is therefore not simply how to measure less invasively, but how to reconstruct enough signal for action while respecting data minimization.

The attribution literature identified the difficulty of assigning conversion credit long before today's privacy restrictions. In pay-per-conversion advertising, multiple exposures may contribute jointly to a conversion, making straightforward last-touch accounting strategically misleading (Jordan et al., 2011). The contemporary setting adds two new complications. First, observable user paths are incomplete because tracking is fragmented across browsers, applications, platforms, and walled gardens. Second, even when data exist somewhere in the system, they may not be legally or contractually combinable. This transforms attribution from an econometric challenge into an infrastructural one.

One response is to shift attention from user-level traceability to privacy-preserving aggregate measurement. Recent work on ad conversion measurement develops differentially private mechanisms that release aggregate campaign statistics while bounding privacy leakage (Delaney et al., 2024). Similar efforts in advertising security propose privacy-preserving conversion tracking and bidding protocols that allow parties to optimize campaigns without fully exposing conversion logs (Zhong et al., 2022). Historically, even comparatively narrow tasks such as frequency capping generated privacy concerns because they required maintaining exposure histories (Farahat et al., 2009). The current research frontier generalizes this insight: almost every useful advertising metric depends on linking behavior over time, which means privacy-preserving measurement must selectively reconstruct just enough linkage to support action.

This requirement produces measurement loss, a concept that is useful even when different papers use different terminology. Measurement loss refers to the gap between true causal or transactional activity and what the system can legitimately and technically observe. In commerce, this loss has four dimensions. First, identity loss occurs when the system cannot reliably connect touchpoints to the same user. Second, path loss occurs when intermediate interactions are not visible. Third, outcome loss occurs when conversions happen in contexts not observable to the measurement partner. Fourth, semantic loss occurs when privacy-preserving aggregation removes detail needed to interpret heterogeneous conversions. Differential privacy can protect against leakage, but it also enlarges semantic loss if the released outputs are too coarse. Secure aggregation can hide local contributions, but it does not guarantee causal interpretability. This is why privacy-preserving measurement must be paired with careful metric design.

A second response is to redesign incentives around partial observability. If the system cannot perfectly observe every conversion path, then contracts and optimization rules must be robust to incomplete measurement. This insight has practical relevance in creator commerce and affiliate-like settings, where payout formulas may depend on noisy or delayed attribution. Platform ecosystems can reduce conflict by using

standardized reporting protocols and clearly documented confidence intervals or eligibility rules, rather than pretending that privacy-preserving measurement is exact. The problem is not only technical uncertainty but contestability. When merchants, creators, and platforms do not share raw logs, they need rules for how approximate measurement enters billing, budgeting, and compensation.

A third response is to exploit structured forms of distributed learning and matching. Vertical federated learning is potentially relevant when different organizations hold complementary features about overlapping user sets, such as platform-side engagement features and merchant-side purchase outcomes (Khan et al., 2025). Federated cross-domain recommendation and federated rating prediction also illustrate how signal can be transferred across domains without unrestricted data pooling (Wu et al., 2022). These approaches are not measurement systems in the narrow sense, but they show how cross-channel information can be leveraged under partitioned data conditions. The limitation is that overlap resolution and identity correspondence remain difficult, especially when privacy constraints forbid explicit linkage.

The measurement problem is therefore best understood as a trade space among fidelity, privacy, and verifiability. High-fidelity user-level attribution maximizes optimization value but creates the greatest exposure. Strong privacy with heavy noise or coarse aggregation minimizes exposure but may be too weak for commercial action. Verifiable protocols improve trust but can slow deployment. A privacy-preserving commerce system must choose a point in this trade space based on decision purpose. Strategic budget allocation across channels may tolerate aggregate reporting with uncertainty bounds. Real-time bidding and creative optimization may require more granular proxies. Creator compensation may require auditable but delayed settlement rather than instantaneous perfect attribution.

Behavioral research further indicates that measurement design shapes trust. Consumers' privacy concerns are not driven solely by formal data collection volume, but by perceived loss of control, opacity, and manipulative intent (Chaudhuri et al., 2023; Dinev et al., 2013; Jiang et al., 2013). Thus, privacy-preserving measurement can have indirect commercial benefits if it reduces the perception that advertising relies on hidden surveillance. At the same time, if reporting becomes too opaque or technically obscure, merchants and creators may distrust the platform instead. The design task is therefore dual-facing: measurement must be privacy-legible to users and accountability-legible to business participants.

The literature suggests several research priorities. First, more work is needed on the interaction between differential privacy parameters and business decision quality in realistic campaign settings. Second, attribution models should explicitly incorporate observability constraints rather than treating missingness as a nuisance. Third, privacy-preserving measurement systems need stronger audit layers, potentially including cryptographic proofs of correct aggregation or rule execution. Fourth, incentive-compatible contracts are needed for settings where measurement is approximate by design. Privacy-preserving measurement is therefore not a residual technical adjustment to post-cookie advertising; it is becoming the central coordination mechanism through which digital commerce decides what counts as performance.

A recent preprint by Yi (2026a) is illustrative here. It proposes a federated and differentially private framework for cross-channel measurement that explicitly integrates Topics and Protected Audience on the web side with Attribution Reporting and SKAdNetwork on the app side, while linking measurement to incentive allocation for SMB advertisers. Because the work is currently a preprint, it should be interpreted cautiously. Even so, it is directly relevant to this review because it treats post-cookie measurement as a joint problem of privacy budgeting, channel harmonization, and decision support rather than as a narrow reporting task.

Social Commerce Advertising and The Creator Economy

Social commerce and the creator economy have transformed the structure of advertising by embedding persuasion, discovery, and transaction inside relationship-rich media environments. In conventional display advertising, targeting and measurement are often discussed as relatively separate layers. In creator commerce, they are intertwined. The value of an impression depends not only on who sees it but on who delivers it, how the content is framed, what platform norms govern disclosure, and how conversion is attributed across social and transactional touchpoints. This makes privacy-preserving design especially difficult.

Research in marketing and interactive media shows that creator influence depends on credibility, perceived similarity, parasocial interaction, and the alignment between message form and platform culture (Bi & Zhang, 2023; Carlson et al., 2022; Teng et al., 2022). Consumers often respond to creators not as interchangeable ad inventory but as trusted or quasi-relational intermediaries. In live-streaming and ecosystem-based analyses of social selling, value creation flows through interactions among platforms, anchors, brands, and audiences rather than through one-directional promotion alone (Xue & Liu, 2023). Recent synthesis work on the creator economy similarly emphasizes that value chains are multi-actor systems involving platforms, creators, advertisers, agencies, and analytics providers (Libai et al., 2025). This means that the data generated in creator commerce are relational, contextual, and partially co-produced.

From a privacy perspective, this relationality matters in two ways. First, creator audiences are strategic assets. Platforms and brands want to infer which audiences are likely to convert, but creators may resist architectures that fully expose their audience data because those data underpin bargaining power. Second, consumers may accept creator recommendations partly because they perceive them as socially situated rather than purely algorithmic. Excessively invasive targeting can undermine that perception and erode trust. The problem is therefore not simply how to protect user privacy,

but how to preserve the autonomy and informational position of intermediaries within the commerce ecosystem.

Federated and privacy-enhancing methods are relevant because they can decouple shared model improvement from unrestricted data access. Audience modeling or recommendation for creator-brand matching could, in principle, be trained across creators, merchants, or platforms without centralizing all raw interaction data. Surveys of federated recommendation and privacy enhancement suggest that this is technically plausible, particularly where participants have partially aligned objectives but do not want to share raw histories (Asad et al., 2023; Gosselin et al., 2022; Papadopoulos et al., 2021). Yet social commerce introduces additional complications. Engagement signals can be noisy, strategic, and vulnerable to manipulation. Creator-side optimization may also encourage gaming behaviors if payout formulas are visible but measurement is imperfect. Thus, privacy-preserving social advertising requires both learning protection and integrity control.

The creator economy also raises the question of revenue sharing and compensation verification. Influencer campaigns, affiliate arrangements, and platform-mediated creator funds all depend on claims about reach, engagement quality, conversions, or downstream sales. Bargoni et al. (2023) show that endorsement services can affect campaign outcomes, but commercial value is contingent on how those services are operationalized. If the creator, platform, and sponsor do not share raw data, disputes can arise over whether audience delivery and conversions were measured correctly. Here zero-knowledge proofs or other verifiable reporting methods may become useful because they allow a party to demonstrate compliance with an agreed formula without disclosing raw logs. Even when such mechanisms are not fully implemented, the design logic is clear: creator monetization requires privacy-preserving visibility rather than either total secrecy or total transparency.

Another relevant strand of literature concerns virtual influencers and platform-mediated identity. The emergence of virtual or AI-driven influencers intensifies privacy tensions because the distinction between content generation, user data exploitation, and synthetic persuasion becomes blurred. Recent work conceptualizes this through a multi-privacy paradox in which consumers may disclose or accept more than they normatively endorse under conditions of convenience, entertainment, or social immersion (Liyanaarachchi et al., 2024). This observation has broader relevance for creator commerce. A privacy-preserving infrastructure cannot assume that disclosed preference is a reliable measure of informed consent. Trustworthy design must therefore include procedural safeguards and governance constraints, not only predictive optimization.

There is also an asymmetry between large platforms and small creators or merchants. Major platforms can absorb privacy engineering costs and may internalize large-scale behavioral data regardless of external restrictions. Smaller actors depend on whatever analytics and APIs the platform exposes. This suggests that trustworthy creator commerce cannot be evaluated solely at the firm level; it must be assessed at the ecosystem level. Standardized, privacy-preserving reporting interfaces can improve access for smaller participants, but only if the platform's governance rules also address auditability, explainability of payout logic, and content moderation consistency. Otherwise, privacy discourse may simply mask further asymmetry in informational control.

A useful design principle is to treat audience matching, conversion estimation, and revenue allocation as distinct yet connected layers. Audience matching can often tolerate partial decentralization and privacy-preserving representation learning. Conversion estimation may rely on differentially private or aggregated outcome reporting. Revenue allocation may require explicit verification and contract rules. Collapsing all three into one opaque platform metric produces efficiency in the short term but undermines long-term trust. Social advertising under privacy constraints therefore benefits from modularity.

The literature also points toward several open questions. First, little is known about how privacy-preserving advertising architectures affect creator bargaining power and market concentration. Second, more research is needed on fairness in creator recommendation and monetization when audience data are unevenly observable. Third, current work still underspecifies how content governance and privacy protection interact. A platform may privacy-protect user data while still recommending harmful or misleading commercial content. Finally, the boundary between recommendation and advertising becomes increasingly blurred in creator ecosystems, suggesting that governance categories inherited from earlier digital advertising may be analytically insufficient. Trustworthy commerce AI in this domain must therefore integrate privacy, content accountability, and economic verification rather than addressing them sequentially.

A closely related 2026 article by Yi (2026c) extends this line of thinking to social e-commerce advertising and creator monetization. Its proposed combination of federated learning and zero-knowledge verification is useful for this review not because it resolves the empirical question of platform fairness, but because it makes explicit a crucial architectural distinction: audience modeling, ad interaction verification, and creator payout accountability can be separated and then recomposed. That formulation reinforces the argument advanced here that trustworthy creator commerce depends on privacy-preserving visibility rather than either unrestricted surveillance or opaque platform reporting.

Trustworthy AI platforms for SMBs

The commercial relevance of privacy-preserving AI depends heavily on whether it can be operationalized for small and medium-sized businesses. Large technology firms can build proprietary data infrastructure, privacy engineering teams, and internal audit capabilities. Most SMBs cannot. They adopt AI through platforms, software-as-a-service vendors, marketplaces, and ecosystem intermediaries. The critical question is therefore how trustworthy AI can be packaged into multi-tenant infrastructures that are technically scalable and organizationally usable.

Table 3 | Design blueprint for trustworthy SMB-oriented commerce AI infrastructure

Layer	Design objective	Key components
Data layer	Protect tenant and user data while enabling useful learning	Local storage controls, data minimization, retention policies, tenant isolation
Learning layer	Support shared improvement without unrestricted pooling	Federated learning, secure aggregation, personalization modules, attack monitoring
Measurement layer	Provide actionable analytics under privacy constraints	Differentially private reporting, aggregate conversion measurement, confidence disclosure
Governance layer	Make the system auditable and controllable	Risk management workflows, logging, incident response, human oversight
Interface layer	Lower adoption barriers for SMBs	Standardized APIs, service cards, policy-aware configuration, interoperable documentation

Multi-tenant architecture is central because SMB-facing systems must serve many organizations with limited customization cost. Yet multi-tenancy creates its own privacy and governance problems. Tenant data must be logically or cryptographically separated; model improvements may need to be shared without leaking competitive information; policy controls must vary across sectors and jurisdictions; and audit logs must remain intelligible to clients who are not AI specialists. Architectural work on platforms and industrial data infrastructures shows that modularity, orchestration logic, and interface design strongly shape adoption and control (Arnold et al., 2022; Gieß & Hutterer, 2025; Hein et al., 2020). In commerce AI, the same principle applies. A platform that merely offers an API endpoint for "recommendation" without governance metadata, logging, and privacy options is not providing trustworthy infrastructure; it is externalizing governance burdens to the least capable actors.

Standardized APIs occupy a strategic position in this architecture. They translate complex learning, ranking, and measurement processes into callable services for merchants, creators, and app developers. Research on APIs and complementor innovation shows that interface standardization can stimulate external innovation, but it can also reproduce dependency and copying risks when governance is weak (Xue et al., 2019). For privacy-preserving commerce AI, API design should do more than expose functionality. It should define the permissible scope of data use, the granularity of outputs, the retention and deletion logic, and the audit trails associated with each invocation. Put differently, trustworthy AI for SMBs requires policy-aware APIs rather than bare prediction endpoints.

Recommendation-as-a-service provides a concrete example. Traditional centralized services ask clients to upload user and item data to the provider. Privacy-preserving recommendation-as-a-service aims to protect data while still enabling shared computation, often through distributed learning or cryptographic protocols (Wang et al., 2018). For SMBs, the attraction is clear: they can access advanced recommendation without building internal pipelines. But unless the service also supports tenant-level governance, documentation of privacy guarantees, and clear liability boundaries, adoption may be superficial. The platform may be "privacy-enhancing" in a technical sense while remaining opaque in contractual or operational terms.

Content governance is another essential component. Digital commerce increasingly depends on AI systems that rank or generate product descriptions, moderate user-generated reviews, screen promotional content, and route creator materials. Trustworthy infrastructure therefore requires alignment between privacy design and content governance. A platform that protects training data but fails to control harmful or noncompliant commercial content is not trustworthy in the commerce sense. Conversely, aggressive moderation without due process or auditability can damage SMBs that depend on the platform for visibility. The governance challenge is to embed policy constraints into workflow design rather than bolting them onto downstream review.

This is where risk management frameworks become important. NIST's AI RMF frames trustworthy AI through functions such as govern, map, measure, and manage, emphasizing life-cycle controls rather than one-off ethical commitments (Tabassi, 2023). The European trustworthy AI guidance and associated assessment tools similarly move from principles toward operational checklists and self-assessment structures (High-Level Expert Group on Artificial Intelligence [HLEG], 2019, 2020). For SMB-oriented commerce platforms, these frameworks suggest that providers should supply governance scaffolding as part of the service. Examples include configurable privacy budgets, model cards or service cards, incident reporting procedures, escalation channels, and human-review triggers for sensitive content or decisions.

A trustworthy SMB platform must also reconcile global shared models with tenant-specific context. A generic recommender or targeting model may not reflect local catalog structure, legal obligations, or sector-specific sensitivities. Federated approaches offer one route by enabling shared learning across tenants while keeping raw data local, but this is only part of the solution. Platforms also need tenant-aware policy layers and configurable governance modules. The most promising direction is therefore not "one model for all" but a layered system in which shared representation learning, tenant-specific fine-tuning, privacy accounting, and rule enforcement are separated but interoperable. **Table 3** summarizes a design blueprint for trustworthy SMB AI platforms.

The SMB context highlights a broader point: trustworthy AI is not simply a quality of algorithms but a service design principle. It concerns how capabilities are exposed, constrained, documented, and monitored. Privacy-preserving infrastructures become economically meaningful only when

these properties are operationalized in ways that non-expert organizations can actually use.

On the provider side, Yi (2026b) offers a complementary architecture paper on trusted AI commercialization infrastructure for SMBs. The paper emphasizes multi-tenant governance controls, standardized recommendation APIs, incentive systems, and content-governance workflows. Its importance for the present review lies in showing that privacy preservation alone is insufficient for SMB adoption unless it is packaged together with auditability, reusable interfaces, and policy-aware service design. As with other recent architecture papers, the contribution is best read as a design blueprint rather than as settled evidence of market-wide effectiveness.

Governance and Standards

The governance debate surrounding commerce AI has matured beyond general ethical exhortation, but operational gaps remain significant. Trustworthy AI is now commonly associated with transparency, accountability, human oversight, robustness, fairness, and privacy. The difficulty is that these principles are often stated at a high level while commerce systems require concrete allocations of decision rights, reporting duties, and technical responsibilities across multiple actors.

Large-scale reviews of AI governance guidelines show convergence around a limited set of principles, especially transparency, justice or fairness, non-maleficence, responsibility, and privacy (Corrêa et al., 2023). However, the same reviews also show fragmentation in interpretation and implementation. Commerce systems intensify this problem because many relevant decisions are distributed across platforms, advertisers, creators, vendors, and analytics intermediaries. A social advertising platform may control ranking and moderation; a merchant may control pricing and campaign objectives; a creator may control message framing; and a measurement vendor may control attribution outputs. Governance therefore has to assign responsibility across a chain rather than inside a single organization.

Trustworthy AI scholarship increasingly argues that this chain perspective must be made explicit. Díaz-Rodríguez et al. (2023) connect AI principles to system requirements and regulation, emphasizing that trustworthy AI involves the translation of abstract norms into design controls, governance processes, and compliance mechanisms. Laux et al. (2024) further caution that trustworthiness should not be conflated with mere acceptability of risk. This distinction matters in commerce because a platform may technically comply with risk categories while still cultivating opaque incentive structures or asymmetrical information control. Governance should therefore ask not only whether risk is bounded, but also who defines acceptable risk and who bears the residual cost when the system fails.

Transparency is often the first principle invoked, yet it is especially difficult to operationalize in commerce. Consumers do not need the same disclosure as merchants, regulators, or creators. Overly general notices can satisfy formal disclosure requirements while communicating little about actual system behavior. Conversely, highly technical descriptions may be unusable for most stakeholders. The practical lesson is that transparency must be role-specific. Users need intelligible explanations of data use and targeting logic at an appropriate level. Business participants need service-level information about measurement error, payout rules, and moderation triggers. Regulators and auditors need traceable logs, documentation, and evidence of policy enforcement. Trustworthy commerce AI therefore requires layered transparency rather than a single disclosure artifact.

Auditability is the next key requirement. AI audits in commerce should extend beyond model performance to include data lineage, privacy accounting, policy conformance, and dispute resolution. Existing governance work and emerging audit discussions indicate that standards boards, risk frameworks, and formalized assessment procedures are increasingly necessary if organizations are to move from principle statements to repeatable oversight (Tabassi, 2023). In digital commerce, audits should be able to address questions such as: Were conversion counts generated according to the documented privacy-preserving protocol? Were creators paid according to the disclosed revenue-sharing rule? Were moderation decisions applied consistently across comparable commercial content? Was model retraining triggered by valid and logged inputs? These are infrastructural audit questions, not merely algorithmic ones.

Responsibility allocation is also changing because privacy-preserving architectures can blur agency. When a recommendation outcome results from federated updates contributed by many participants, or when attribution depends on securely aggregated event streams, it can become harder to determine who is accountable for a harmful or misleading result. Governance must therefore define accountability despite distributed computation. One approach is to separate responsibility by function: participants remain responsible for data lawfulness at source; orchestrators are responsible for protocol integrity and model governance; service providers are responsible for API behavior and documentation; and deploying firms remain responsible for downstream business use. While such separation is imperfect, it is more realistic than treating "the AI system" as a singular responsible entity.

A rights-based perspective strengthens this analysis. Research on trustworthy AI in social media argues that rights language can clarify what is at stake when commercial platforms use AI to shape expression, exposure, and behavior (Lewis & Moorkens, 2020). In commerce, rights-based thinking helps connect privacy to due process, contestability, and freedom from manipulative opacity. This is especially important in creator ecosystems and SMB platform dependence, where weaker actors may be constrained by platform-defined rules they cannot meaningfully negotiate. Governance should thus include appeal mechanisms, explanation pathways, and contractual clarity for participants, not only consumer-facing privacy controls.

The practical standards landscape is becoming denser. NIST's AI RMF provides a flexible but structured reference for organizational risk management (Tabassi, 2023). The EU's trustworthy AI guidelines and assessment tools provide lifecycle-oriented self-assessment logic (HLEG, 2019, 2020). Reg-

ulatory developments such as the EU AI Act intensify the pressure to specify risk classification, technical documentation, and accountability arrangements, even if debate continues regarding the relation between trustworthiness and legal acceptability (Laux et al., 2024). For commerce platforms, the implication is clear: governance cannot remain an informal compliance appendix. It has to be embedded in architecture, documentation, vendor management, and interface design.

Several unresolved issues deserve emphasis. First, privacy and transparency can conflict. Strong privacy-preserving computation may reduce the amount of raw evidence available for audit or explanation. This increases the importance of protocol-level verification and carefully designed logs. Second, governance frameworks still under-address platform power asymmetries. A dominant platform may comply procedurally while imposing opaque economic dependencies on merchants and creators. Third, the connection between AI governance and content governance remains insufficiently theorized in commerce research. Commercial content ranking, sponsored recommendations, and creator monetization all blur the line between economic infrastructure and informational governance. Finally, most current governance frameworks are organization-centric, whereas digital commerce is ecosystemic. Future governance work must therefore address interoperability of audit artifacts, shared incident taxonomies, and cross-organizational accountability standards.

In summary, governance is not an external constraint on privacy-preserving commerce AI. It is the mechanism that determines whether privacy-preserving methods become trustworthy infrastructures or merely technical shields around opaque systems.

Conclusion

This review has argued that privacy-preserving and trustworthy AI in digital commerce should be understood as an infrastructure problem. Recommendation, advertising, cross-channel measurement, creator monetization, and SMB enablement all depend on the controlled circulation of data, model updates, metrics, and governance signals across organizational boundaries. Privacy constraints do not eliminate the need for these flows; they force them to be redesigned.

The literature shows that federated learning, differential privacy, secure aggregation, and zero-knowledge verification offer a meaningful technical toolkit, but their value depends on how they are composed. Federated learning without governance remains vulnerable. Differential privacy without scenario-sensitive metric design can produce unusable outputs. Secure aggregation without verifiability may protect data while undermining trust among business participants. Zero-knowledge methods without operational integration remain aspirational. The most promising direction is therefore modular composition anchored in scenario-specific requirements.

Three conclusions follow. First, cross-channel measurement has become the strategic bottleneck of privacy-preserving commerce AI. Without credible and privacy-aware measurement, firms cannot allocate budgets, evaluate creators, or refine recommendation and advertising systems. Second, social commerce and creator monetization reveal that trustworthiness is both consumer-facing and partner-facing. Privacy protection must coexist with auditable revenue logic, content governance, and role-specific transparency. Third, SMB adoption depends on whether trustworthy AI is embedded in multi-tenant platforms, standardized APIs, and governance-by-design rather than offered as a collection of advanced but inaccessible technical options.

Future research should therefore move in four directions. One direction is quantitative evaluation of privacy-utility trade-offs in realistic commerce workflows rather than in abstract benchmarks alone. A second is the development of privacy-preserving measurement systems that include uncertainty modeling, audit support, and incentive compatibility. A third is the study of ecosystem power under privacy-preserving architectures, especially whether such architectures decentralize control or simply repackage centralization. A fourth is the design of interoperable governance artifacts for platform ecosystems, including service cards, privacy ledgers, and standardized audit evidence. The next generation of digital commerce AI will be judged not only by whether it predicts well, but by whether it can produce commercially useful intelligence without undermining privacy, contestability, and accountability.

References

1. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC conference on computer and communications security (pp. 308–318). https://doi.org/10.1145/2976749.2978318
2. Arnold, L., Jöhnk, J., Vogt, F., & Urbach, N. (2022). IIoT platforms' architectural features: A taxonomy and five prevalent archetypes. Electronic Markets, 32(2), 927–944. https://doi.org/10.1007/s12525-021-00520-0
3. Asad, M., Shaukat, S., Javanmardi, E., Nakazato, J., & Tsukada, M. (2023). A comprehensive survey on privacy-preserving techniques in federated recommendation systems. Applied Sciences, 13(10), Article 6201. https://doi.org/10.3390/app13106201
4. Bargoni, A., Giachino, C., Battisti, E., & Iaia, L. (2023). The effects of influencer endorsement services on crowdfunding campaigns. Journal of Services Marketing, 37(1), 40–52. https://doi.org/10.1108/JSM-12-2021-0444
5. Bi, N. C., & Zhang, R. (2023). "I will buy what my 'friend' recommends": The effects of parasocial relationships, influencer credibility and self-esteem on purchase intentions. Journal of Research in Interactive Marketing, 17(2), 157–175. https://doi.org/10.1108/JRIM-08-2021-0214
6. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. In Proceedings of the 2017 ACM SIGSAC conference on computer and communications security (pp. 1175–1191). https://doi.org/10.1145/3133956.3133982
7. Boerman, S. C., Kruikemeier, S., & Zuiderveen Borgesius, F. J. (2017). Online behavioral advertising: A literature review and research agenda. Journal of Advertising, 46(3), 363–376. https://doi.org/10.1080/00913367.2017.1339368
8. Carlson, J. R., Hanson, S., Pancras, J., Ross, W. T., & Rousseau-Anderson, J. (2022). Social media advertising: How online motivations and congruency influence perceptions of trust. Jour-

nal of Consumer Behaviour, 21(2), 197–213. https://doi.org/10.1002/cb.1989

9. Chaudhuri, R., Chatterjee, S., & Vrontis, D. (2023). Antecedents of privacy concerns and online information disclosure: Moderating role of government regulation. EuroMed Journal of Business, 18(3), 467–486. https://doi.org/10.1108/EMJB-11-2021-0181
10. Chen, Z., Gan, W., Wu, J., Hu, K., & Lin, H. (2025). Data scarcity in recommendation systems: A survey. ACM Transactions on Recommender Systems, 3(3), Article 27. https://doi.org/10.1145/3639063
11. Cheng, Y., Liu, Y., Chen, T., & Yang, Q. (2020). Federated learning for privacy-preserving AI. Communications of the ACM, 63(12), 33–36. https://doi.org/10.1145/3387107
12. Checco, A., Bianchi, G., & Leith, D. J. (2017). BLC: Private matrix factorization recommenders via automatic group learning. ACM Transactions on Privacy and Security, 20(2), Article 4. https://doi.org/10.1145/3041760
13. Corrêa, N. K., Galvão, C., Santos, J. W., Del Pino, C., Pontes Pinto, E., Barbosa, C., Massmann, D., Mambrini, R., Galvão, L., Terem, E., & de Oliveira, N. (2023). Worldwide AI ethics: A review of 200 guidelines and recommendations for AI governance. Patterns, 4(10), Article 100857. https://doi.org/10.1016/j.patter.2023.100857
14. Delaney, J., Ghazi, B., Harrison, C., Ilvento, C., Kumar, R., Manurangsi, P., Pál, M., Prabhakar, K., & Raykova, M. (2024). Differentially private ad conversion measurement. Proceedings on Privacy Enhancing Technologies, 2024(2), 124–140. https://doi.org/10.56553/popets-2024-0044
15. Díaz-Rodríguez, N., Del Ser, J., Coeckelbergh, M., López de Prado, M., Herrera-Viedma, E., & Herrera, F. (2023). Connecting the dots in trustworthy artificial intelligence: From AI principles, ethics, and key requirements to responsible AI systems and regulation. Information Fusion, 99, Article 101896. https://doi.org/10.1016/j.inffus.2023.101896
16. Dinev, T., Xu, H., Smith, J. H., & Hart, P. (2013). Information privacy and correlates: An empirical attempt to bridge and distinguish privacy-related concepts. European Journal of Information Systems, 22(3), 295–316. https://doi.org/10.1057/ejis.2012.23
17. Dwivedi, Y. K., Ismagilova, E., Hughes, L., Carlson, J., Filieri, R., Jacobson, J., Jain, V., Karjaluoto, H., Kefi, H., Krishen, A. S., Kumar, V., Rahman, M. M., Raman, R., Rauschnabel, P. A., Rowley, J., Salo, J., Tran, G. A., & Wang, Y. (2021). Setting the future of digital and social media marketing research: Perspectives and research propositions. International Journal of Information Management, 59, Article 102168. https://doi.org/10.1016/j.ijinfomgt.2020.102168
18. Dwork, C. (2006). Differential privacy. In M. Bugliesi, B. Preneel, V. Sassone, & I. Wegener (Eds.), Automata, languages and programming (pp. 1–12). Springer. https://doi.org/10.1007/11787006_1
19. Farahat, A., Duda, C., Doshi, A., & Muthukrishnan, S. (2009). Privacy preserving frequency capping in internet banner advertising. In Proceedings of the 18th international conference on World Wide Web (pp. 1147–1148). https://doi.org/10.1145/1526709.1526900
20. Gao, C., Zheng, Y., Li, N., Li, Y., Qin, Y., Piao, J., Quan, Y., Chang, J., Jin, D., He, X., & Li, Y. (2023). A survey of graph neural networks for recommender systems: Challenges, methods, and directions. ACM Transactions on Recommender Systems, 1(1), 1–51. https://doi.org/10.1145/3568022
21. Gieß, A., & Hutterer, A. (2025). The future of data management: A delimitation of data platforms, data spaces, data meshes, and data fabrics. Information Systems and e-Business Management, 23(4), 971–997. https://doi.org/10.1007/s10257-025-00707-4
22. Gosselin, R., Vieu, L., Loukil, F., & Benoit, A. (2022). Privacy and security in federated learning: A survey. Applied Sciences, 12(19), Article 9901. https://doi.org/10.3390/app12199901
23. Gu, X., Sabrina, F., Fan, Z., & Sohail, S. (2023). A review of privacy enhancement methods for federated learning in healthcare systems. International Journal of Environmental Research and Public Health, 20(15), Article 6539. https://doi.org/10.3390/ijerph20156539
24. Hein, A., Schreieck, M., Riasanow, T., Setzke, D. S., Wiesche, M., Böhm, M., & Krcmar, H. (2020). Digital platform ecosystems. Electronic Markets, 30(1), 87–98. https://doi.org/10.1007/s12525-019-00377-4
25. High-Level Expert Group on Artificial Intelligence. (2019). Ethics guidelines for trustworthy AI. Publications Office of the European Union. https://doi.org/10.2759/346720
26. High-Level Expert Group on Artificial Intelligence. (2020). Assessment list for trustworthy artificial intelligence (ALTAI) for self-assessment. Publications Office of the European Union. https://doi.org/10.2759/002360
27. Jiang, Z., Heng, C. S., & Choi, B. C. F. (2013). Research note—Privacy concerns and privacy-protective behavior in synchronous online social interactions. Information Systems Research, 24(3), 579–595. https://doi.org/10.1287/isre.1120.0441
28. Jere, M., Farnan, T., & Koushanfar, F. (2021). A taxonomy of attacks on federated learning. IEEE Security & Privacy, 19(2), 20–28. https://doi.org/10.1109/MSEC.2020.3039941
29. Jordan, P., Mahdian, M., Vassilvitskii, S., & Vee, E. (2011). The multiple attribution problem in pay-per-conversion advertising. In M. Mavronicolas & V. V. Vazirani (Eds.), Algorithmic game theory (pp. 31–43). Springer. https://doi.org/10.1007/978-3-642-24829-0_5
30. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Nitin Bhagoji, A., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., … Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210. https://doi.org/10.1561/2200000083
31. Kalloori, S., & Klingler, S. (2021). Horizontal cross-silo federated recommender systems. In Proceedings of the 15th ACM conference on recommender systems (pp. 680–684). https://doi.org/10.1145/3460231.3478863
32. Khan, A., ten Thij, M., & Wilbik, A. (2025). Vertical federated learning: A structured literature review. Knowledge and Information Systems, 67, 3205–3243. https://doi.org/10.1007/s10115-025-02356-y
33. Laux, J., Wachter, S., & Mittelstadt, B. (2024). Trustworthy artificial intelligence and the European Union AI act: On the conflation of trustworthiness and acceptability of risk. Regulation & Governance, 18(1), 3–32. https://doi.org/10.1111/rego.12512
34. Lewis, D., & Moorkens, J. (2020). A rights-based approach to trustworthy AI in social media. Social Media + Society, 6(3). https://doi.org/10.1177/2056305120954672
35. Libai, B., Rosario, A. B., Beichert, M., Donkers, B., Haenlein, M., Hofstetter, R., van der Lans, R., Lanz, A., Li, H. A., Mayzlin, D., Muller, E., Shapira, D., Yang, J., & Zhang, L. (2025). Influencer marketing unlocked: Understanding the value chains driving the creator economy. Journal of the Academy of Marketing Science, 53, 4–28. https://doi.org/10.1007/s11747-024-01073-2
36. Liyanaarachchi, G. P., Mifsud, M., & Viglia, G. (2024). Virtual influencers and data privacy: Introducing the multi-privacy paradox. Journal of Business Research, 176, Article 114584. https://doi.org/10.1016/j.jbusres.2024.114584
37. Luo, S., Xiao, Y., Zhang, X., Liu, Y., Ding, W., & Song, L. (2024). PerFedRec++: Enhancing personalized federated recommendation with self-supervised pre-training. ACM Transactions on

Intelligent Systems and Technology, 15(5), Article 98. https://doi.org/10.1145/3664927

38. Morimoto, M. (2021). Privacy concerns about personalized advertising across multiple social media platforms in Japan: The relationship with information control and persuasion knowledge. International Journal of Advertising, 40(3), 431–451. https://doi.org/10.1080/02650487.2020.1796322
39. Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. Future Generation Computer Systems, 115, 619–640. https://doi.org/10.1016/j.future.2020.10.007
40. Mukhopadhyay, S., & Bouwman, H. (2019). Orchestration and governance in digital platform ecosystems: A literature review and trends. Digital Policy, Regulation and Governance, 21(4), 329–351. https://doi.org/10.1108/DPRG-11-2018-0067
41. Papadopoulos, P., Abramson, W., Hall, A. J., Pitropakis, N., & Buchanan, W. J. (2021). Privacy and trust redefined in federated machine learning. Machine Learning and Knowledge Extraction, 3(2), 333–356. https://doi.org/10.3390/make3020017
42. Qin, J., Liu, B., & Qian, J. (2021). A novel privacy-preserved recommender system framework based on federated learning. In Proceedings of the 4th international conference on simulation, modeling and intelligent computing systems (pp. 82–88). https://doi.org/10.1145/3451471.3451485
43. Sharma, K., Lee, Y.-C., Nambi, S., Salian, A., Shah, S., Kim, S.-W., & Kumar, S. (2024). A survey of graph neural networks for social recommender systems. ACM Computing Surveys, 56(10), Article 265. https://doi.org/10.1145/3661821
44. Sun, X., Yu, F. R., Zhang, P., Sun, Z., Xie, W., & Peng, X. (2021). A survey on zero-knowledge proof in blockchain. IEEE Network, 35(4), 198–205. https://doi.org/10.1109/MNET.011.2000473
45. Sweeney, L. (2002). k-anonymity: A model for protecting privacy. International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, 10(5), 557–570. https://doi.org/10.1142/S0218488502001648
46. Tabassi, E. (Ed.). (2023). Artificial intelligence risk management framework (AI RMF 1.0). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1
47. Teng, T., Li, H., Fang, Y., & Shen, L. (2022). Understanding the differential effectiveness of marketer versus user-generated advertisements in closed social networking sites: An empirical study of WeChat. Internet Research, 32(6), 1910–1929. https://doi.org/10.1108/INTR-04-2021-0268
48. Truong, N., Sun, K., Wang, S., Guitton, F., & Guo, Y. (2021). Privacy preservation in federated learning: An insightful survey from the GDPR perspective. Computers & Security, 110, Article 102402. https://doi.org/10.1016/j.cose.2021.102402
49. Wang, J., Arriaga, A., Tang, Q., & Ryan, P. Y. A. (2018). Facilitating privacy-preserving recommendation-as-a-service with machine learning and cryptographic techniques. In Proceedings of the 2018 ACM SIGSAC conference on computer and communications security (pp. 1900–1917). https://doi.org/10.1145/3243734.3278504
50. Wei, K., Li, J., Ding, M., Ma, C., Yang, H. H., Farokhi, F., Jin, S., Quek, T. Q. S., & Poor, H. V. (2020). Federated learning with differential privacy: Algorithms and performance analysis. IEEE Transactions on Information Forensics and Security, 15, 3454–3469. https://doi.org/10.1109/TIFS.2020.2988575
51. Wu, M., Li, L., Tao, C., Rigall, E., Wang, X., & Xu, C.-Z. (2022). FedCDR: Federated cross-domain recommendation for privacy-preserving rating prediction. In Proceedings of the 31st ACM international conference on information & knowledge management (pp. 2179–2188). https://doi.org/10.1145/3511808.3557320
52. Xue, J., & Liu, M. T. (2023). Investigating the live streaming sales from the perspective of the ecosystem: The structures, processes and value flow. Asia Pacific Journal of Marketing and Logistics, 35(5), 1157–1186. https://doi.org/10.1108/APJML-11-2021-0822
53. Xue, L., Song, P., Rai, A., Zhang, C., & Zhao, X. (2019). Implications of application programming interfaces for third-party new app development and copycatting. Production and Operations Management, 28(8), 1887–1902. https://doi.org/10.1111/poms.13021
54. Zhong, K., Ma, Y., & Angel, S. (2022). Ibex: Privacy-preserving ad conversion tracking and bidding. In Proceedings of the 2022 ACM SIGSAC conference on computer and communications security (pp. 3223–3237). https://doi.org/10.1145/3548606.3560651
55. Yi, X. (2026a). A federated and differentially private incentive–marketing framework for privacy-preserving cross-channel measurement in AI-powered digital commerce. Preprints, 202602.1929. https://doi.org/10.20944/preprints202602.1929.v1
56. Yi, X. (2026b). Trusted AI commercialization infrastructure for SMBs: A unified multi-tenant architecture integrating incentive systems, content governance, and standardized recommendation APIs. Preprints, 202602.1885. https://doi.org/10.20944/preprints202602.1885.v1
57. Yi, X. (2026c). Privacy-enhanced ad targeting for social e-commerce: A federated learning framework with zero-knowledge verification for creator monetization. Frontiers in Business and Finance, 3(1), 102–113. https://doi.org/10.71465/fbf653