A Privacy-Preserving, Data-Driven Personalization Framework for B2C Digital Sales Optimization Using Federated Learning and Customer 360 Integration

Sanjay Pradhan[1], Keshav Adhikari[2], and Ramesh Bhandari[3]

[1] *Tribhuvan University, Central Department of Management, Kirtipur, Kathmandu 44618, Nepal*

[2] *Kathmandu University School of Management, Department of Finance and Accounting, Dhulikhel, Kavrepalanchok 45200, Nepal*

[3] *Pokhara University, Faculty of Management Studies, Lekhnath 33700, Pokhara, Nepal*

Abstract

Business-to-consumer digital sales ecosystems have progressively shifted toward data-driven decisioning, with personalization engines determining exposure, pricing, and recommendations across web, mobile, and omnichannel surfaces. Concurrently, regulatory constraints, cross-jurisdictional data residency requirements, and increased sensitivity to surveillance practices have limited the feasibility of centralizing fine-grained behavioral data. Organizations operate multiple legacy stacks, fragmented identifiers, and inconsistent consent records that complicate unified modeling and raise the risk profile of conventional central warehouses. In this setting, there is interest in architectures that enable predictive personalization without concentrating raw identifiable data. This paper presents a privacy-preserving personalization framework that integrates a governed Customer 360 data model with federated learning and edge-resident policy execution for B2C digital sales optimization. The framework describes how identity resolution, feature derivation, consent-aware masking, and eligibility constraints can be combined with decentralized optimization protocols, secure aggregation, and calibrated noise. The focus is on compatibility with ranking, uplift, and budget-constrained exploration models under heterogeneous traffic, partial client participation, and non-stationary behavior. The analysis covers objective formulations, algorithmic components, and operational mechanisms for drift detection, fairness monitoring, and fail-safe fallbacks. Empirical evaluation on simulated and replayed multi-channel data illustrates the behavior of the proposed design under varying privacy parameters, participation levels, and consent churn, without overstating performance. The study aims to provide a technically detailed, implementation-oriented description of how Customer 360 integration and federated learning can be composed to support personalization that remains aligned with privacy, compliance, and engineering constraints in contemporary B2C digital sales environments.

1 Introduction

Personalization in business-to-consumer digital sales has evolved from heuristic segmentation toward continuous, model-driven decisioning embedded into every interaction surface.

Customer journeys now traverse responsive websites, native applications, search and social referral flows, in-store digitization, and post-purchase service channels. Each touchpoint emits structured and unstructured signals, including impression logs, browsing paths, cart events, transactional records, support tickets, and feedback traces. Organizations aim to translate these observations into calibrated decisions about which offers to display, which products to recommend, how frequently to engage, and through which channels to intervene. The underlying aspiration is to allocate sales and attention opportunities in a manner that is responsive to individual preferences and constraints, while remaining compatible with long-term relationship stability and regulatory scrutiny. [1]

Traditional approaches to this objective have relied heavily on centralized data platforms that consolidate fine-grained behavioral events and customer attributes into large-scale warehouses or lakes. In such architectures, model development teams extract training datasets, engineer features spanning multiple channels and time horizons, and train predictive or prescriptive models that are subsequently deployed into serving systems. Centralization facilitates cross-channel attribution, unified experimentation, and complex modeling techniques, since most relevant information is collected in one place. However, it also creates concentrated stores of sensitive data, including identifiers, inferred attributes, and detailed histories of customer behavior. These stores elevate the operational risk of unauthorized access, complicate compliance with evolving regulations, and introduce challenges for implementing timely consent revocation and purpose limitation. [2]

At the same time, the structural complexity of contemporary B2C ecosystems undermines the assumption that a fully unified, static central view is either attainable or desirable. Many organizations operate across jurisdictions with distinct data residency rules and sector-specific obligations. Historical systems coexist with newer platforms, each with their own identity schemes, logging conventions, and retention policies. Authentication is intermittent, devices are shared or replaced, and browser storage policies reduce the longevity of cookies. Identity resolution, which attempts to reconcile this fragmentation into coherent entities, is therefore probabilistic and dynamic. When central systems combine multiple identifiers too aggressively, they risk constructing linkages that exceed legitimate expectations or are difficult to unwind when new information arrives. When they combine them too conservatively, they fail to provide a consistent substrate for learning and measurement.

In response to these pressures, the notion of a Customer 360 model has emerged as a way to structure customer data into governed profiles with explicit lineage, linkage confidence, and consent semantics. Properly implemented, such a model is less a single database than a set of services responsible for identity resolution, feature computation, consent evaluation, and eligibility determination. These services attempt to make the conditions under which data may be used transparent and auditable [3]. They also define the vocabulary of features and masks through which downstream decision systems operate. However, many current personalization deployments still assume that these Customer 360 constructs will be materialized and processed centrally, even when the raw data feeding them originate at the network edge or in siloed environments.

Federated learning offers an alternative perspective by relocating parts of the training computation to where data is generated or regionally confined. Instead of uploading raw examples, participating devices or silos compute local model updates that are aggregated to produce a global model. When combined with secure aggregation and carefully designed clipping and noise mechanisms, this paradigm reduces the exposure of granular records while still enabling joint optimization [4]. The conceptual fit with B2C personalization is immediate: decisioning models depend on recent and sensitive interactions that can remain local, with only structured, masked updates crossing boundaries. Yet deploying federated learning in isolation does not resolve the conceptual and operational challenges of identity, consent, or channel constraints. Without alignment with a governed Customer 360 layer, local optimizations may use inconsistent feature definitions, disregard evolving eligibility rules, or implicitly reconstruct linkages that governance layers intend to control.

The introduction of federated learning into a B2C digital sales environment therefore raises several interdependent questions. First, how should the feature space be defined so that both edge and central components share a common representation grounded in Customer 360 semantics, while still honoring data minimization and consent scoping? Second, how should policies incorporate eligibility constraints, fairness conditions, pacing limits, and campaign-level goals in a setting where individual interaction logs are not centrally available for direct constraint enforcement? Third, how can the system maintain robustness under non-stationary conditions, such as promotions, seasonal effects, tracking changes, and consent churn, when observation of detailed distributions is restricted and learning is driven by decentralized, noisy updates? [5]

Aspect	**Description**	**Implication**
Personalization Evolution	Shift from heuristic segmentation to model-driven decisioning	Enables adaptive, individualized marketing actions
Interaction Channels	Web, app, in-store, service channels	Multi-source data integration required
Decision Objectives	Offers, recommendations, engagement frequency	Balancing personalization and compliance

Table 1. Evolution of B2C Personalization Context

Architecture	**Strengths**	**Weaknesses**
Centralized Data Platforms	Unified data access, powerful modeling	Privacy risk, compliance burden
Decentralized/Federated	Local computation, reduced exposure	Limited coordination, inconsistent features
Hybrid (Customer 360 + FL)	Governed identity, privacy alignment	Requires complex integration

Table 2. Architectural Trade-offs in Personalization Systems

Constraint Type	**Notation**	**Purpose**
Resource	$b(\theta) \leq \beta$	Enforces spend, pacing, exposure budgets
Fairness	$\Psi(\theta)$	Prevents biased action allocation
Privacy	Differential/Consent-based	Ensures lawful feature usage

Table 3. Constraint Taxonomy in the Optimization Problem

These questions point toward a need for an integrated framework rather than isolated techniques. A purely centralized Customer 360 approach without federated learning may

Component	Mathematical Role	Operational Meaning
Policy $\pi_\theta(a\|x)$	Decision function	Maps context to actions
Utility $u(y, a)$	Objective term	Captures reward or cost
Lagrange Multipliers (λ, γ)	Constraint weights	Balance between goals and limits

Table 4. Elements of the Constrained Optimization Formulation

Federated Step	Computation	Purpose
Local Gradient	$g_k(\theta)$ on D_k	Client-side learning under consent
Aggregation	$\bar{g}(\theta) = \sum w_k g_k$	Privacy-preserving update fusion
Global Update	$\theta \leftarrow \theta - \eta\bar{g}$	Model refinement without raw data sharing

Table 5. Federated Learning Process under Privacy Constraints

meet modeling requirements but can be misaligned with privacy constraints and operational risk management. A purely federated approach without explicit Customer 360 integration may achieve decentralization but lack clarity about which entities are being optimized, how consent and eligibility are enforced, and how to reconcile signals across channels. What is required is a composition in which governed identity resolution, feature engineering, and mask computation operate as preregistered transformations, and federated optimization consumes only the resulting permitted signals. In such a composition, personalization policies can be expressed in terms of abstracted feature vectors and eligibility masks instead of direct identifiers or raw events. [6]

From a modeling standpoint, B2C personalization introduces additional complexities beyond standard supervised learning. Policies must act sequentially, with feedback that is often partial or delayed. Logged data predominately reflects historical policies, creating confounding between actions and outcomes. Practical systems must handle a large and dynamic catalog of items or offers, sparse observations, and structural constraints such as inventory levels and contact frequency caps [7]. Federated learning in this setting must accommodate non-independent and non-identically distributed client data, as different devices or silos observe distinct user cohorts, channel mixes, and contextual distributions. It must also operate in resource-constrained environments that limit the number of local training steps, the complexity of models, and the size of updates [8].

To address these realities, the proposed framework views personalization as a constrained stochastic optimization problem defined over the policy space induced by Customer 360 features and eligible actions, optimized via federated coordination under explicit privacy mechanisms. Utility functions can encompass conversion and revenue, but also incorporate penalties for overexposure or misaligned allocations. Constraints express budgets, pacing, and fairness conditions encoded at the policy level rather than through direct per-event interventions [9]. The federated learning process approximates gradients of this objective using only local computations and privacy-preserving aggregates, ensuring that central systems never directly observe the raw contexts and outcomes that drive personalization decisions.

In operational terms, introducing such a framework into an existing B2C environment requires incremental, compatible design rather than disruptive replacement. Customer 360 components can be refactored to expose governed feature and mask services with clear contracts. Federated client libraries can be embedded into mobile applications, web edge infrastructure, or regional backends as an extension of existing telemetry and rec-

Customer 360 Element	Function	Governance Implication
Identity Resolution	Deterministic + probabilistic linkage	Confidence propagation, lineage tracking
Feature Computation	Aggregated behavioral vectors	Consent-filtered derivations
Action Masks	Eligibility constraints $m_{i,t}(a)$	Enforce legal and operational exclusions

Table 6. Core Components of the Customer 360 Governance Layer

ommendation capabilities. Early deployments can focus on constrained use cases, such as limited-scope recommendations or contact policy adjustments, allowing organizations to observe the behavior of decentralized learning under realistic conditions before expanding to broader decision surfaces [10]. Throughout this process, governance teams retain visibility into feature catalogs, mask rules, and privacy budgets, while operational teams examine calibration, pacing, and stability metrics derived from aggregates.

2 Theoretical Foundations and Problem Setting

The personalization problem in B2C digital sales can be viewed as a constrained stochastic decision problem with partially observed, heterogeneous, and evolving contexts. Let i index logical entities such as persons, households, or devices, depending on the governing linkage rules. Let t index decision opportunities, which may correspond to page views, app sessions, push eligibility windows, outbound communication slots, or in-session offers. Each decision opportunity is associated with a context vector $x_{i,t}$, constructed from permissible features provided by a Customer 360 pipeline under current consent state and identity resolution confidence.

For each opportunity, the policy selects an action $a_{i,t}$ from an action set $\mathcal{A}$ that may include ranked content sets, offer classes, message channels, or no-op choices. The outcome $y_{i,t}$ encodes a measurable response, such as click, conversion, incremental revenue, or a retention proxy. The system designer specifies a utility function $u(y_{i,t}, a_{i,t})$ that incorporates both direct outcomes and costs such as discount expenditure, inventory constraints, or communication fatigue.

The design objective is to identify a policy $\pi_\theta(a \mid x)$, parameterized by θ, that maximizes expected utility under privacy budgets, eligibility constraints, channel pacing limits, and fairness considerations [11]. Define a feasible policy family Π as the set of parameter vectors θ for which all applicable constraints are satisfied almost surely when actions are sampled from $\pi_\theta(a \mid x)$ given the realized context process. The core problem can be written as a constrained optimization over randomized policies.

$$\max_{\theta \in \Pi} \mathbb{E}[u(y, a)]$$

The expectation is taken with respect to the joint distribution of (x, a, y) induced by the environment and the randomized policy, acknowledging that y depends causally on a and may reflect delayed effects.

Operational constraints are encoded through additional functionals. Let $b(\theta)$ represent a vector of expected resource consumptions, such as average exposure rates, spend, or communication volume per channel and segment [12]. Let β denote allowable thresholds:

$$b(\theta) \leq \beta$$

Fairness conditions and brand safety rules can be represented by divergence constraints between allocation patterns across segments. Let $\rho_s(\theta)$ denote the long-run action rate for segment s and $\bar{\rho}_s$ a reference level, possibly derived from policy guidelines or legal requirements. A simple quadratic penalty captures deviations:

$$\Psi(\theta) = [13] \sum_s \left(\rho_s(\theta) - \bar{\rho}_s\right)^2$$

The overall learning problem can be expressed as minimizing a Lagrangian that balances utility, resource, and fairness terms:

$$\mathcal{J}(\theta, \lambda, \gamma) = -\mathbb{E}[u(y, a)] + \lambda^\top (b(\theta) - \beta) + \gamma\, \Psi(\theta)$$

Here λ and γ are non-negative multipliers. In a classical centralized setting, stochastic gradients of $\mathcal{J}$ could be approximated from logged data. However, in a privacy-preserving setting, access to raw samples is restricted [14]. Instead, gradient estimates must be constructed from local computations conducted where data resides, with only constrained aggregates revealed.

Federated learning offers a mechanism for such distributed optimization. Let clients (devices or regional services) be indexed by k, each holding a private dataset D_k of tuples $(x_{i,t}, a_{i,t}, y_{i,t})$ compatible with its consent and governance context. The global parameter vector θ is updated via rounds in which a subset of clients compute stochastic gradient estimates based on local data and then share privatized updates that are aggregated. Let $g_k(\theta)$ denote the clipped, privatized gradient computed by client k [15]. The central coordinator observes only an aggregate:

$$\bar{g}(\theta) = \sum_k w_k g_k(\theta)$$

where w_k are weights proportional to sample counts or participation probabilities. The coordinator then updates:

$$\theta \leftarrow [16]\theta - \eta\, \bar{g}(\theta)$$

This formulation embeds privacy guarantees via the properties of $g_k(\theta)$ and the secure aggregation protocol. The constraint functionals $b(\theta)$ and $\rho_s(\theta)$ are estimated from aggregated, privacy-limited statistics, ensuring that the enforcement of sales and policy constraints does not require central inspection of individual records.

Within this theoretical setting, the integration with a Customer 360 data model is not incidental. The definition of $x_{i,t}$, the mapping from identifiers to entities, the eligibility of actions per entity, and the reflection of consent are all managed upstream. The federated objective then operates over feature representations and eligibility masks that have already encoded governance constraints, allowing the mathematical optimization to focus on tradeoffs between utility, budgets, and fairness in a reduced space.

3 Customer 360 Data Model and Identity Resolution Mechanisms

The Customer 360 layer provides the representational substrate on which personalization policies are defined [17]. Its role is to fuse multi-channel signals into governed profiles without exposing unnecessary detail to the central learner. This layer can be treated as a set of services that, for each decision opportunity, return a context vector $x_{i,t}$, an identity confidence summary, and an action eligibility mask, all constructed according to lineage and consent rules.

Identifiers arise from login accounts, hashed emails, device IDs, browser cookies, in-store tokens, payment instruments, and support interactions. These identifiers are linked through both deterministic and probabilistic rules. Deterministic linkage relies on stable fields such as verified accounts or consistent login across channels. Probabilistic linkage addresses cases where evidence is partial, such as similar behavioral patterns or overlapping device characteristics [18]. Let $r_{p,q}$ denote the relational evidence between identifiers p and q, and let $z_{p,q}$ be a binary latent variable indicating whether they represent the same underlying entity. A logistic model approximates the posterior linkage probability.

$$\Pr(z_{p,q} = 1 \mid r_{p,q}) = \sigma\big(\alpha^\top \phi(r_{p,q})\big)$$

Here $\phi(r_{p,q})$ encodes similarity metrics and co-occurrence statistics, and σ is the logistic function. These linkage probabilities inform the formation of entity clusters. Instead of committing to a single hard clustering in all cases, the Customer 360 layer can propagate confidence scores to downstream consumers. Decision policies then use either deterministic clusters above a threshold or incorporate confidence into features, allowing more cautious treatment of uncertain linkages.

Feature construction within the Customer 360 model must respect several constraints [19]. Features should represent aggregated behavior over time windows that have clear semantics, such as recency-weighted counts, category affinities, response propensities, and channel interaction frequencies. They must be reproducible in both offline and online contexts, meaning that the transformation pipeline is shared or formally aligned. They must be segmented by consent, with sensitive categories either excluded or derived only in coarse, non-identifying forms. They should expose validity intervals or time-to-live metadata to prevent unbounded staleness in downstream use.

Consent is treated as a dynamic attribute influencing which raw events may contribute to features and which features can be activated for a given entity at serving time [20]. If consent is withdrawn or narrowed, the corresponding feature groups are recomputed or masked. The Customer 360 services encapsulate this logic so that personalization systems receive only feature vectors already filtered according to the current consent state. This approach reduces the risk that federated models inadvertently condition on attributes that should no longer be in scope.

Action eligibility masks express structural constraints such as geographic availability of products, regulatory exclusions, age gates, channel capabilities, and customer-specified preferences [21]. For each context, the mask $m_{i,t}(a)$ takes value 1 if action a is permitted and 0 otherwise. When a stochastic policy $\pi_\theta(a \mid x)$ is defined, actual sampling is performed from a renormalized distribution over allowed actions:

$$\pi_\theta^{(m)}(a \mid x) = \frac{m(a)\,\pi_\theta(a \mid x)}{\sum_{a'} m(a')\,\pi_\theta(a' \mid x)}$$

This masking ensures that optimization never implicitly assumes infeasible actions. From a data governance perspective, it is advantageous that governance constraints appear in masks and feature definitions before federated optimization occurs, keeping the roles of identity, consent, and modeling logically separated yet interoperable.

Stage	**Process**	**Objective**
Client Selection	Choose eligible devices or silos	Ensure participation diversity, reliability
Model Broadcast	Send $\theta^{(t)}$ to clients	Synchronize global and local states
Local Computation	Train on permitted local data	Compute clipped updates under consent
Secure Aggregation	Combine masked $\tilde{\Delta}_k^{(t)}$	Preserve privacy during update fusion
Central Update	Compute $\theta^{(t+1)}$	Refine model globally with weights $w_k^{(t)}$

Table 7. Federated Optimization Stages in B2C Personalization

Challenge	Mitigation Strategy	Outcome
Non-IID Data	Add proximal term to local loss	Stabilizes convergence, reduces bias
Heterogeneous Devices	Weighted participation, adaptive steps	Balanced contribution and efficiency
Intermittent Connectivity	Quorum-based aggregation	Robustness to dropouts

Table 8. Techniques for Robust Federated Optimization

Mechanism	Mathematical Representation	Guarantee
Secure Aggregation	Masked sum of Δ_k	Hides individual client updates
Differential Privacy	$\tilde{\Delta}_k = \Delta_k + \mathcal{N}(0, \sigma^2 C^2 I)$	Bounds single-client influence
Privacy Accounting	$\varepsilon = \Gamma(\sigma, q, T, \delta)$	Tracks cumulative privacy budget

Table 9. Privacy Mechanisms in Decentralized Learning

4 Federated Optimization Architecture for B2C Personalization

The federated optimization architecture coordinates edge-resident learners and a central aggregator under conditions typical of B2C ecosystems, including intermittent connectivity, uneven device capabilities, traffic bursts, and deployment across multiple jurisdictions. The architecture can be understood as a sequence of recurring stages: client selection, model broadcast, local computation, secure aggregation, central update, and policy deployment. [22]

A federation is defined over a set of clients, which may represent end-user devices, in-region data processors, or isolated data domains operated under separate legal entities. Each client is associated with a stream of decision and outcome data that can be used for local updates under consent and governance constraints. At each training round, the aggregator samples a subset of clients according to a participation scheme that may consider historical reliability, hardware profiles, and recent contribution levels. Selected clients receive the current global parameter vector $\theta^{(t)}$ and training configuration hyperparameters.

Each participating client constructs local examples compatible with the Customer 360 feature schemas and eligibility masks defined for its environment. Using these examples, it computes a local objective that may include empirical loss, regularization, and locally estimable parts of constraint penalties [23]. The client then performs a bounded number of optimization steps, such as stochastic gradient descent or local coordinate updates, subject to limits on compute time and data passes to avoid resource exhaustion and memorization.

To enforce privacy and robustness, each client clips its update $\Delta_k^{(t)}$ to a norm bound C and applies noise if required by the global privacy schedule. The update is then masked using cryptographic secure aggregation so that the aggregator only learns the sum or weighted average of updates once a minimum quorum of clients has contributed. The aggregator computes:

$$\theta^{(t+1)} = \theta^{(t)} + \eta^{(t)} \sum_k w_k^{(t)} \tilde{\Delta}_k^{(t)}$$

where $\tilde{\Delta}_k^{(t)}$ are privatized updates and $w_k^{(t)}$ are normalized weights. The learning rate

Compliance Aspect	Implementation	Impact
Consent Management	Feature and mask filtering in Customer 360	Prevents unauthorized data use
Data Retention	Time-bounded samples, expiring models	Aligns with regulatory timelines
Cross-Border Governance	Local training, aggregated updates only	Minimizes raw data transfers

Table 10. Compliance Controls Integrated in the Architecture

Model Component	Equation	Function
Ranking Policy	$\pi_\theta(a\|x) = \frac{e^{f_\theta(x,a)/\tau}}{\sum_{a'} e^{f_\theta(x,a')/\tau}}$	Scores and samples feasible actions
Uplift Estimation	$\Delta_\theta(x) = g_\theta(x) - h_\theta(x)$	Captures incremental treatment effect
Gradient Contribution	$G = (r - b)\nabla_\theta \log \pi_\theta(a\|x)$	Local policy gradient under privacy

Table 11. Core Formulas for Policy Learning in Federated Setting

$\eta^{(t)}$ and participation weights can adapt based on observed convergence behavior and participation statistics.

Because client data is typically non-IID, with some clients associated with specific geographies, segments, or interaction modes, naive federated averaging can converge slowly or to biased solutions. To counter this, local objectives may include proximal terms that keep local parameters near the global reference, promoting stability [24]. Additionally, the architecture can maintain multiple model heads or calibration layers that specialize to observed client clusters while sharing a global backbone representation. These variations are still coordinated via federated updates, with heads receiving updates only from relevant client subsets.

Deployment of the learned policy to serving systems reuses the same configuration as training to avoid skew. Edge services embed the current parameters and apply the same Customer 360 features and eligibility masks to compute action distributions. Online experimentation can be integrated by randomizing over policy variants within the federated framework, allowing robust evaluation of new modeling choices without degrading privacy guarantees. [25]

5 Privacy and Compliance Guarantees in Decentralized Learning

Formal privacy and compliance assurances are necessary to justify decentralization as more than an informal safeguard. The architecture employs a combination of secure aggregation, differential privacy, consent-aware data minimization, and auditable configuration management to bound information leakage and support regulatory alignment.

Secure aggregation ensures that the aggregator cannot inspect individual updates from clients. Each client k constructs a mask vector that, when combined with masks from other participants, cancels out in the aggregated sum but hides its own contribution. Only when a sufficient number of clients complete the protocol does the aggregator recover the summed update [26]. This mechanism reduces the risk of targeted inference about individual clients, provided the quorum thresholds are respected.

Constraint	Formulation	Purpose
Budget Limit	$\max_\theta \mathbb{E}[r] - \lambda(\mathbb{E}[c(a)] - C)$	Enforce spending caps during learning
Fairness Penalty	Segment disparity regularization	Prevent biased allocations
Exploration Balance	Temperature τ and score scaling	Maintain coverage and constraint adherence

Table 12. Budgeted and Fairness-Constrained Policy Optimization

Differential privacy further constrains what can be inferred from the aggregated results. By clipping updates to a maximum norm C and adding Gaussian noise with standard deviation proportional to C, the influence of any single client's data on the global model is bounded. A privatized client update can be expressed as:

$$\tilde{\Delta}_k = \Delta_k + \mathcal{N}(0, \sigma^2 C^2 I)$$

When many clients participate in a round, the added noise averages out at the aggregate level, allowing learning to proceed while maintaining individual protections [27]. The cumulative privacy loss over multiple rounds can be tracked via an appropriate accountant. If q denotes the sampling rate of clients per round and T the number of effective rounds, an approximate privacy budget ε at a given failure probability δ can be expressed as a function:

$$\varepsilon = \Gamma(\sigma, q, T, \delta)$$

where Γ is derived from the chosen analytical accountant [28]. This explicit relationship allows organizations to choose parameters σ, C, q, and T that align with internal and external expectations around acceptable privacy loss.

Compliance considerations extend beyond formal privacy metrics. The Customer 360 layer enforces purpose limitation and consent restrictions before data enters any learning process. Features representing sensitive categories can be omitted or transformed to non-identifying aggregates [29]. Data retention policies are implemented by limiting the temporal coverage of examples used for training and by expiring model snapshots according to documented schedules. Access controls and configuration logs establish who can modify thresholds, feature lists, or mask rules.

To address concerns about membership inference and reconstruction, the system avoids retaining raw update logs or sketches that could be combined with other signals to reverse engineer local contributions. Only aggregated, noise-perturbed updates are used to evolve models, and recovery procedures rely on checkpointed global states rather than historical per-client traces. These decisions narrow the potential avenues for misuse without relying on secrecy of algorithms. [30]

Finally, cross-border data transfers are mitigated when regional clients conduct training locally and only share aggregated, noise-added updates that do not contain direct identifiers. While legal interpretations vary, this structural reduction in raw data movement can be documented and evaluated, providing a concrete basis for governance discussions.

6 Policy Learning: Ranking, Uplift, and Budgeted Exploration Models

Within the federated and privacy-preserving setting, the policy learning component is responsible for mapping contexts to treatment probabilities in a way that balances relevance, incrementality, and resource management [31]. The policy is parameterized to support ranking of alternatives, estimation of incremental impact, and structured exploration that respects sales constraints.

Consider a parametric scoring function $f_\theta(x, a)$ that predicts the utility-associated signal for action a in context x [32]. A softmax transformation yields a stochastic policy, integrated with eligibility masks:

$$\pi_\theta(a \mid x) = \frac{\exp(f_\theta(x, a)/\tau)}{\sum_{a'} \exp(f_\theta(x, a')/\tau)}$$

and sampling is restricted to feasible actions via eligibility-aware renormalization. The temperature τ controls the spread of probabilities, influencing the degree of exploration.

Incremental impact, or uplift, is particularly relevant in B2C personalization, since observed responses confound baseline propensity and treatment effects [33]. A practical strategy approximates uplift using separate or shared models for treated and control conditions. Let $g_\theta(x)$ approximate the expected outcome under treatment and $h_\theta(x)$ under control. The uplift estimate is:

$$\Delta_\theta(x) = g_\theta(x) - h_\theta(x)$$

Policies can use $\Delta_\theta(x)$ as input to ranking or thresholding rules [34]. When budgets restrict the proportion of customers who can receive intensive treatments, a Lagrangian approach selects individuals with highest estimated uplift until the constraint is met, implemented through score-based sampling probabilities instead of hard thresholds to preserve exploration.

In federated training, clients update parameters of f_θ, g_θ, and h_θ based on locally observed outcomes, which may include bandit-style feedback where only the chosen action's outcome is known. Policy gradient or actor-critic methods can be adapted: each client computes an estimate of the gradient of expected utility with respect to θ from its logged interactions, clips and privatizes this estimate, and submits it for aggregation. For example, a simple local contribution for a selected action a with observed reward r and baseline b has the form:

$$G = (r - b)\nabla_\theta[35] \log \pi_\theta(a \mid x)$$

After clipping and noise addition, such terms are combined centrally to update the policy parameters. The baseline b can be approximated using local moving averages or simple predictive functions that remain on-device.

Budgeted exploration introduces global constraints, such as maximum impressions per period or per-campaign spending caps. Let $c(a)$ denote the cost of action a and let C denote an allowable average cost per decision [36]. The constrained optimization can be framed as maximizing expected reward minus a penalty on exceeding the budget:

$$\max_\theta \ \mathbb{E}[r] - \lambda(\mathbb{E}[c(a)] - C)$$

The multiplier λ is adjusted over time based on observed aggregate costs using a simple dual ascent rule that operates on privacy-limited cost aggregates. This adjustment modifies the effective policy by scaling scores associated with expensive actions, thereby aligning exploration with budget adherence without central inspection of individual-level data. [37]

Fairness-oriented constraints are implemented via penalties on disparities in action rates or outcomes across segments defined in the governed schema. These penalties are computed from aggregated summaries, such as privatized counts of actions per segment. Policy updates then incorporate gradients of these penalties, nudging parameters away from allocations that systematically over- or under-target certain groups, while preserving overall utility where possible.

7 Robustness, Drift Management, and Reliability Engineering

Robustness to distributional changes is a practical requirement in digital sales environments. Seasonal demand cycles, promotional events, macroeconomic shifts, interface changes, and regulatory updates can all alter the relationship between contexts, actions, and outcomes [38]. A federated, privacy-preserving system must detect and adapt to such shifts without relying on unrestricted access to raw logs.

The framework utilizes drift indicators constructed from aggregate feature and outcome summaries. Instead of collecting full distributions, clients compute compressed sketches over selected features and send noise-perturbed aggregates via secure aggregation. Let s_t denote a privatized sketch summarizing a feature dimension or model residuals at time t, and let s_0 represent a reference snapshot from a stable period. A simple drift score can be defined as: [39]

$$d_t = \|s_t - s_0\|$$

When d_t exceeds configured thresholds consistently, the system interprets this as evidence of distributional change. The coordinator may respond by increasing training frequency, adjusting learning rates, retraining calibration layers, or constraining exploration until new patterns are better understood.

Consent churn introduces another dimension of drift, as changes in policy or user choices alter which features and examples remain available. The Customer 360 services track the prevalence of consent conditions and the coverage of different feature groups [40]. Aggregated metrics on feature availability are used to adjust model architectures; if certain signals become sparse or prohibited, their corresponding parameters can be attenuated or removed. These adjustments mitigate the risk that policies depend on features with unstable legal status.

Reliability engineering encompasses the behaviors of the system under partial failures, connectivity issues, and anomalous contributions. The federated setup must tolerate stragglers, misconfigured clients, and transient spikes in noise. To address these, aggregation rounds can impose minimum and maximum participation bounds, discarding outlier contributions when necessary [41]. Models are checkpointed regularly, enabling rollback if inconsistent updates or implementation defects are detected.

Fallback policies are an important safeguard. When drift scores, calibration errors, or budget deviations exceed acceptable bands, serving systems may switch to conservative policies that rely on coarse segmentation or deterministic rules derived from stable aggregates. Because these policies are defined in terms of the same Customer 360 features and eligibility masks, the transition does not require re-engineering of integration points. Recovery to the learned policy follows once indicators stabilize over a verification window. [42]

Logging and observability are implemented with attention to privacy. Only aggregated indicators such as action shares, outcome rates, pacing measures, and drift scores are stored for long-term analysis. No individual decision traces or feature vectors are required for routine monitoring. These design choices support neutral and inspectable operations while reducing the exposure of sensitive details. [43]

8 Implementation Considerations and Deployment Patterns

Implementing a privacy-preserving, data-driven personalization framework that combines Customer 360 integration with federated learning in B2C digital sales requires attention to interactions between data engineering, model lifecycle management, security controls, and channel infrastructure. The conceptual architecture outlined earlier must be instantiated as a set of interoperable services with well-defined interfaces, consistent semantics, and bounded operational complexity. In practice, deployment occurs in environments where multiple legacy systems coexist, latency constraints are strict for certain decision paths, and organizational responsibilities are divided across analytics, security, marketing, and platform teams. This section discusses implementation choices in a neutral and detailed manner, examining deployment topologies, data flows, federation orchestration, observability, risk management, and incremental rollout, with the objective of establishing a coherent pattern that can be adapted to heterogeneous B2C contexts without relying on central accumulation of raw behavioral records.

A central consideration is the embedding of the Customer 360 layer as an authoritative provider of consent-aware profiles and features for both training and serving [44]. Rather than viewing Customer 360 as a monolithic warehouse, it is operationally effective to treat it as a collection of services: identity resolution, feature computation, consent and policy evaluation, and eligibility masking. Each service exposes deterministic interfaces that accept identifiers or tokens from channel endpoints and return feature vectors, confidence scores, and masks derived from governed logic. These services should be deployable in-region to conform with data residency expectations, with replication and configuration management that avoid divergence in definitions. To reduce training-serving skew, the same transformation code or declarative feature definitions can be compiled into both streaming feature pipelines and federated client libraries. When a mobile application or web edge node constructs local training examples, it uses the same feature derivations that power online policy evaluation, reducing discrepancies that might otherwise accumulate through bespoke implementations. [45]

Deployment of the federated learning layer can follow two principal patterns, depending on where sufficient computational capacity and data locality exist. In a device-centric pattern, end-user devices participate directly as clients, computing local updates from recent interactions stored transiently on the device under the platforms security model. In a silo-centric pattern, regional or domain-specific services act as clients, aggregating examples from a bounded set of sources within a legal or organizational boundary before performing local optimization. Both patterns are compatible with the same coordination protocol, provided the client library enforces clipping, noise addition where applicable, eligibility-aware sampling of examples, and secure aggregation participation. The choice between patterns is influenced by connectivity reliability, application footprint constraints, and feasible trust assumptions [46]. In many B2C deployments a hybrid arrangement is adopted, with certain channels contributing device-level updates while others are represented through regional silos, harmonized via shared schemas and model parameterization.

To formalize the scheduling behavior of federated rounds under operational constraints, one can define a selection function over the candidate client set. Let $\mathcal{K}_t$ denote the set of available clients at round t, each with a reliability score $q_k^{(t)}$ representing its historical completion rate and quality of updates. A probabilistic sampler defines inclusion probabilities $\pi_k^{(t)}$ that reflect both fairness across regions and the need to emphasize reliable contributors. An example rule, implemented at the coordinator, can be expressed as

$$\pi_k^{(t)} = \frac{q_k^{(t)}}{\sum_{j \in \mathcal{K}_t} q_j^{(t)}}$$

subject to minimum participation criteria for underrepresented silos enforced by lower bounds on aggregate inclusion over sliding windows. This expression illustrates how the selection mechanism remains simple yet tunable, without exposing individual-level information [47]. Reliability scores can themselves be derived from non-sensitive metrics such as timely completion rates and adherence to clipping constraints.

Practical serving patterns must reconcile global parameter propagation with channel-specific latency and availability requirements. A common design is to maintain a versioned model registry at the coordinator, from which signed, immutable artifacts are distributed to edge locations or clients according to update schedules. For low-latency web and app interactions, models are loaded into in-memory scoring services that rely on cached features and masks from the Customer 360 layer. To avoid excessive chatter, model refreshes occur at controlled intervals, and each artifact carries metadata specifying compatible feature definitions, expected privacy budgets, and applicable channels [48]. When a policy evaluation is requested, the serving stack retrieves a profile from the local or regional Customer 360 service, applies eligibility masks, computes scores using the current model, and samples actions according to the configured stochastic rule. If any dependency fails, a deterministic fallback policy defined over coarse aggregates is invoked, minimizing disruption while respecting constraints.

A key requirement for reliable deployment is consistent semantics of consent and eligibility between training and serving. Since federated learning relies on local construction of examples, the client library must embed the same consent evaluation rules as the runtime profile services. This can be achieved by shipping versioned policy configurations that encode which categories of data are admissible for training at a given time and in a given jurisdiction [49]. When a consent state changes, devices or regional services adjust their local data retention and feature inclusion accordingly, and future updates reflect the narrowed scope without requiring central intervention. To maintain traceability, each training round can record the configuration identifiers used by participating clients, allowing auditors to verify that updates were produced under approved policies.

Resource efficiency considerations influence the design of update computation and communication patterns. B2C environments may include millions of potential clients, only a fraction of which can be engaged per round without incurring prohibitive network and battery costs [50]. Compression techniques reduce the size of updates while respecting the imposed clipping bounds. For example, clients may quantize gradient coordinates to a small number of discrete levels or transmit only top entries above a magnitude threshold, with shared randomization ensuring that expectations remain unbiased. Let g_k denote a clipped gradient vector on client k. A simple sparsified estimator $\hat{g}_k$ can retain a subset of coordinates indexed by $\mathcal{I}_k$ and rescale them appropriately. In conceptual form,

$$\hat{g}_{k,j} = \begin{cases} \frac{g_{k,j}}{p_{k,j}} & j \in \mathcal{I}_k \\ 0 & \text{otherwise} \end{cases}$$

where $p_{k,j}$ is the probability that coordinate j is selected. Such mechanisms are configured to preserve statistical properties of updates while reducing communication volume, and their parameters are recorded to support interpretability of convergence characteristics. [51]

Observability and monitoring in this framework must operate under privacy constraints and operational noise while still providing sufficient signals for safe operation. Rather than capturing raw examples, the system can track aggregated counters, distributional sketches, and calibration metrics computed from privatized summaries. For instance, the discrepancy between predicted and observed outcome rates in each segment over a window can be approximated from counts of predictions and events, with random noise added where necessary to enforce aggregation thresholds. These indicators support detection of miscalibration, model drift, or unintended biases. Alerting thresholds are framed in terms of stable ratios or error bands, chosen to account for the variance introduced by noise and partial participation [52]. Engineering teams rely on these indicators to decide when to trigger retraining, rolling back a model, or adjusting hyperparameters, without requiring invasive inspection of individual-level trajectories.

Change management and rollout strategies are particularly important in B2C settings where even small perturbations in personalization logic can have wide impact. A disciplined pattern introduces new models through shadow modes and incremental traffic allocation. In shadow mode, a candidate policy runs in parallel with the current policy, consuming the same contextual inputs but not influencing actions, enabling the collection of aggregate performance statistics. Once stability is observed, a small share of traffic is allocated to the new policy, and results are analyzed using conservative estimators compatible with bandit-style logging [53]. During this process, federated training continues under both parameterizations, but secure aggregation prevents the coordinator from tying differences back to any specific client. Rollback remains straightforward because model artifacts are versioned and dependencies on Customer 360 features and policies are explicitly declared.

Failure modes deserve explicit attention. Misconfigurations in feature definitions, masking rules, or privacy parameters can lead to degraded performance or non-compliant behavior if undetected. To reduce such risks, a layered validation approach verifies at build time that candidate configurations satisfy structural constraints, such as the absence of disallowed attributes, consistency of feature namespaces, and compatibility between model signatures and Customer 360 schemas [54]. At runtime, synthetic test profiles, constructed without personal data, exercise the same code paths as production requests, ensuring that eligibility masks, pacing logic, and fallback rules behave as expected. When anomalies are detected in production metrics, the system prefers safe degradation to aggressive self-correction, switching to simpler policies that are easier to reason about until root causes are assessed.

The interaction between privacy budgets and operational objectives must be handled as a continuous tuning process rather than a one-time choice. Organizations select noise scales, clipping bounds, and sampling rates that correspond to acceptable aggregate privacy guarantees over defined horizons, subject to internal standards and applicable law. These choices affect learning efficiency; therefore, they are best evaluated through controlled experiments using synthetic or de-identified replay data, followed by cautious deployment with explicit documentation [55]. Governance bodies can review the mapping between configuration parameters and resulting privacy budgets expressed through an accountant function, assessing whether tradeoffs are balanced. Because federated learning is integrated with a Customer 360 layer that already performs data minimization and masking, the effective exposure per individual remains structured, and additional adjustments can be made iteratively as operational experience accumulates.

Finally, the implementation pattern benefits from a modularization that aligns with organizational boundaries. The Customer 360 services, federated orchestration layer, and policy serving components can be developed and maintained by distinct teams, provided their interfaces are concrete and stable. Security and compliance teams focus on configuration baselines, audit trails, and independent verification of feature and mask rules [56]. Product and marketing teams interact mainly with policy-level abstractions such as objectives, eligible actions, and pacing constraints, without direct involvement in raw data handling. This separation reduces coupling and supports incremental refinement of the overall system. Rather than framing the resulting deployment as a singular solution, it is more appropriate to regard it as a disciplined composition of components, each with limited responsibility, which collectively enable personalization in B2C digital sales while respecting privacy, governance, and reliability requirements.

9 Experimental Design, Results, and Operational Insights

Evaluating a privacy-preserving, federated Customer 360 framework requires methods that approximate production conditions without compromising the very protections the system is designed to uphold. Synthetic datasets and replay-based simulations offer one approach: they emulate cross-channel interactions, identity linkages, and consent dynamics based on observed structural patterns, without using raw identifiers or sensitive attributes. [57]

In a representative evaluation setting, a population of simulated entities is constructed with latent preferences, price sensitivities, and propensities to respond across channels. Each entity produces events across devices and sessions, sometimes authenticated and sometimes anonymous. An identity resolution module, parameterized independently of the policy, generates deterministic links for high-certainty cases and probabilistic scores for ambiguous ones, mirroring real-world uncertainty. A Customer 360 process then derives features such as recency, frequency, category engagement, and channel mix, constrained by synthetic consent flags that can change over time. [58]

The action set includes promotional offers, content variants, and frequency-control decisions, each with associated constraints and costs. Outcomes are generated from structural models that reflect both baseline behavior and incremental response. Logging encompasses bandit feedback: only outcomes for chosen actions are recorded. The federated learning system is deployed over many logical clients, each holding a shard of the interaction history. Training proceeds with secure aggregation, update clipping, and varying levels of Gaussian noise to emulate privacy budgets. [59]

Results are analyzed along multiple dimensions. First, convergence behavior of the federated models is compared to a centralized benchmark trained on the same synthetic but fully accessible data. Across a broad range of participation rates and clipping norms, the federated system tracks the centralized objective within modest gaps, particularly when the number of rounds is sufficient to offset noise. Second, uplift-based policies are assessed using replay estimators grounded in inverse-propensity weighting and doubly robust corrections. Under realistic parameterizations, the policies learned in a federated manner achieve consistent incremental gains relative to non-personalized baselines, without requiring direct sharing of raw data across clients. [60]

Third, the impact of privacy parameters is examined. As the noise scale increases to represent stricter privacy guarantees, performance shows gradual degradation, but remains within practically acceptable ranges for a significant interval. Beyond this region, learning becomes unstable, highlighting that privacy budgets should be chosen with awareness of participation levels and objective complexity. Fourth, drift scenarios are simulated by modifying latent preferences, seasonality, or consent policies mid-run. Drift indicators based on privatized sketches respond to these changes, and adaptive retraining reduces miscalibration and budget misallocation. [61]

Operational insights emerging from such experiments include the observation that stable feature definitions and eligibility masks, anchored in the Customer 360 layer, are critical to reducing training-serving skew, regardless of centralization choices. Federated learning benefits from modest over-provisioning of participating clients per round, as this accelerates effective averaging of noise and variance. Update clipping bounds require calibration to local gradient distributions; excessively tight bounds slow adaptation, while excessively loose bounds amplify noise impact without yielding proportional information gains.

From a process perspective, integrating governance approval into feature and mask configuration, and documenting privacy accountant settings, supports more predictable deployments. The experiments suggest that a neutral, well-specified composition of Customer 360 integration, federated optimization, and constrained policy learning can function reliably under a range of conditions typical for B2C digital sales, without relying on central pools of raw event-level data. [62]

10 Conclusion

This paper has presented a detailed formulation and architectural description of a privacy-preserving, data-driven personalization framework for B2C digital sales optimization that combines Customer 360 integration with federated learning and constrained policy execution. The discussion began with a problem setting in which personalization must operate under fragmented identifiers, heterogeneous channels, dynamic consent, and regulatory constraints that discourage unrestricted data centralization. To address this, the framework relies on a governed Customer 360 layer to perform identity resolution, construct reproducible and consent-aware features, and expose action eligibility masks that encode structural and policy constraints upstream from model training.

On top of this representational substrate, federated optimization protocols coordinate decentralized learning, using secure aggregation, clipping, and differential privacy mechanisms to limit the influence and visibility of individual contributions. The policy learning component supports ranking, uplift estimation, and budgeted exploration within a mathematically explicit constrained optimization view, integrating fairness and pacing terms derived from aggregated indicators [63]. Robustness is treated through drift detection using privatized sketches, adaptive retraining strategies, and conservative fallback policies that provide predictable behavior under uncertainty. Experimental considerations, based on synthetic and replayed scenarios, were used to characterize how participation levels, privacy parameters, and configuration choices affect stability and utility, without attributing undue significance to any single metric.

The resulting picture is one of a modular system in which Customer 360 services, federated learning, and operational guardrails interact through clearly defined interfaces and objectives. The framework does not remove the need for careful engineering, monitoring, and governance, but it describes how organizations can structure personalization workflows so that sensitive data remains closer to its source while central components operate on aggregated, privacy-limited signals. Future refinements may adjust modeling techniques, privacy accountants, or deployment practices; the core principles of governed feature construction, decentralized optimization, constrained policy learning, and auditable behavior provide a basis for incremental evolution aligned with both commercial requirements and privacy expectations in B2C digital sales environments. [64]

References

[1] R. Wu, “Research on order tracking system for mto,” *DEStech Transactions on Computer Science and Engineering*, no. iccis, Nov. 13, 2019.

[2] M. Wodnicka and D. Skurpel, “Growth global market of e-commerce cross border: The case of poland,” *EUROPEAN RESEARCH STUDIES JOURNAL*, vol. XXIV, no. Special Issue 1, pp. 1121–1135, Mar. 1, 2021.

[3] N. Susanti, “Analisis implikasi kepuasan pelanggan terhadap perilaku pasca pembelian melalui testimoni dalam situs pemasaran internet,” *Jurnal Manajemen Teori dan Terapan| Journal of Theory and Applied Management*, vol. 2, no. 1, Apr. 22, 2009.

[4] H. Etemad, “The emergence of online global market place and the multilayered view of international entrepreneurship,” *Journal of International Entrepreneurship*, vol. 15, no. 4, pp. 353–365, Dec. 19, 2017.

[5] S. R. Clinton, “Importance of technology investments in the logistics service providers: A case study of ups and its use of online tools,” *Journal of Applied Business Research (JABR)*, vol. 24, no. 2, Jan. 14, 2011.

[6] S. V. Pascalau, “Application of b2c digital marketing,” *AGORA INTERNATIONAL JOURNAL OF ECONOMICAL SCIENCES*, vol. 15, pp. 13–16, Feb. 14, 2022.

[7] D. Choi, C. Y. Chung, and J. Young, “Sustainable online shopping logistics for customer satisfaction and repeat purchasing behavior: Evidence from china,” *Sustainability*, vol. 11, no. 20, pp. 5626–, Oct. 12, 2019.

[8] S. H. Kukkuhalli, “Enabling customer 360 view and customer touchpoint tracking across digital and non-digital channels,” *Journal of Marketing & Supply Chain Management*, vol. 1, no. 3, 2022.

[9] K. Matsuoka, “Exploring the interface between management accounting and marketing: A literature review of customer accounting,” *Journal of Management Control*, vol. 31, no. 3, pp. 157–208, Apr. 17, 2020.

[10] N. A. Morgan, H. Feng, and K. A. Whitler, “Marketing capabilities in international marketing,” *Journal of International Marketing*, vol. 26, no. 1, pp. 61–95, Mar. 1, 2018.

[11] B.-C. Su, H. Lin, and Y.-M. Wang, “The business model of digital platforms for the sharing economy: Intensive case study methodology for rover.com pet boarding platform,” *Sustainability*, vol. 14, no. 23, pp. 16 256–16 256, Dec. 6, 2022.

[12] W. You, M. Xia, L. Liu, and D. Liu, “Customer knowledge discovery from online reviews,” *Electronic Markets*, vol. 22, no. 3, pp. 131–142, Jun. 7, 2012.

[13] S. Rungsrisawat and T. Chankoson, “Understanding social media effects across different parties interactions,” *Journal of Security and Sustainability Issues*, vol. 9, no. 4, pp. 1363–1377, Jun. 30, 2020.

[14] M. Zavyalova, N. Skrynko, and Z. Helevachuk, “E-commerce: Global and domestic development trends,” *Scientific Journal of Polonia University*, vol. 35, no. 4, pp. 40–51, Jun. 28, 2019.

[15] S. P. Sazonov, I. A. Ezangina, A. Polianskaia, and A. I. Chunakov, “Digital transformation and its role in the reproduction of innovative development of the modern banking institution of russia,” *SHS Web of Conferences*, vol. 114, pp. 01 009–, Jul. 13, 2021.

[16] A. Houcheimi, “The key e-tail opportunities and challenges in the lebanese e-commerce market,” *Journal of Information System and Technology Management*, vol. 7, no. 26, pp. 13–31, Jun. 10, 2022.

[17] H. N. Utami, D. T. Alamanda, and R. M. Ramdani, “Factors determining buyer-seller relationships: Empirical results from an agribusiness perspective,” *Sosiohumaniora*, vol. 24, no. 1, pp. 140–140, Mar. 2, 2022.

[18] C. Lorenzo-Romero, M.-Á. Gómez-Borja, null Alej, and ro Molla-Descals, “Effects of utilitarian and hedonic atmospheric dimensions on consumer responses in an online shopping environment,” *African Journal of Business Management*, vol. 5, no. 21, pp. 8649–8667, Sep. 23, 2011.

[19] O. F. Bustinza, G. Parry, and F. Vendrell-Herrero, "Supply and demand chain management: The effect of adding services to product offerings," *Supply Chain Management: An International Journal*, vol. 18, no. 6, pp. 618–629, Sep. 23, 2013.

[20] F. D. Puspadewa and H. Suhaimi, "Developing personal selling sop, improving website and improving company profile of msme of pt. selula dwiphaloka teknologi," *IPTEK Journal of Proceedings Series*, vol. 0, no. 5, pp. 8–, Dec. 25, 2019.

[21] G. Salvietti, C. Ziliani, C. Teller, M. Ieva, and S. Ranfagni, "Omnichannel retailing and post-pandemic recovery: Buildingăaăresearch agenda," *International Journal of Retail & Distribution Management*, vol. 50, no. 8/9, pp. 1156–1181, Feb. 28, 2022.

[22] H. Webb, S. Liu, and M.-R. Yan, "Evaluation of m-payment technology and sectoral system innovationa comparative study of uk and indian models," *Electronics*, vol. 8, no. 11, pp. 1282–, Nov. 4, 2019.

[23] I. Samanta, "Investigating the buyer-seller relationships in the economic recession: A qualitative approach," *Independent Journal of Management & Production*, vol. 7, no. 2, pp. 340–366, Jun. 1, 2016.

[24] K. ervená and M. Sabayová, "Sharing economy in the slovak republic (selected aspects)," *Financial Law Review*, no. 24 (4), Dec. 30, 2021.

[25] C. Suteja, "Development stages and online marketing analysis of open trip e-tourism websites in indonesia," *International Journal of Advanced Trends in Computer Science and Engineering*, vol. 8, no. 5, pp. 1988–1995, Oct. 15, 2019.

[26] G. R. Sanden, "Language strategies in multinational corporations: A cross-sector study of financial service companies and manufacturing companies," *AIB Insights*, vol. 16, no. 3, pp. 18–21, Aug. 1, 2016.

[27] N. Vila and I. Küster, "The role of usability on stimulating smes on line buying intention: An experiment based on a ficticius web site design," *Quality & Quantity*, vol. 46, no. 1, pp. 117–136, Apr. 9, 2010.

[28] T. Hughes, M. Stone, E. Aravopoulou, L. T. Wright, and L. Machtynger, "Academic research into marketing: Many publications, but little impact?" *Cogent Business & Management*, vol. 5, no. 1, pp. 1 516 108–1 516 108, Jan. 1, 2018.

[29] R. Gruescu, R. Nanu, and A. Tanasie, "Human resources development and ict contribution to the tourist destination competitiveness," *EUROPEAN RESEARCH STUDIES JOURNAL*, vol. XII, no. 4, pp. 87–100, Nov. 1, 2009.

[30] M. Kotarba, "Measuring digitalization key metrics," *Foundations of Management*, vol. 9, no. 1, pp. 123–138, Feb. 23, 2017.

[31] S. H. Kukkuhalli, "Increasing digital sales revenue through 1:1 hyper-personalization with the use of machine learning for b2c enterprises," *Journal of Artificial Intelligence, Machine Learning and Data Science*, vol. 1, no. 1, 2023.

[32] C. Guo, Z. Du, and X. Kou, "Products ranking through aspect-based sentiment analysis of online heterogeneous reviews," *Journal of Systems Science and Systems Engineering*, vol. 27, no. 5, pp. 542–558, Sep. 29, 2018.

[33] P. Rogers, "Rented but mine! application of psychological ownership theory to access-based consumption and the circular economy," *Circular Economy and Sustainability*, vol. 1, no. 2, pp. 719–744, Sep. 7, 2021.

[34] M. N. Iskandar and F. Aprilianty, “Building an effective branding strategy: A study case of raiment,” *Malaysian Journal of Social Sciences and Humanities (MJSSH)*, vol. 6, no. 9, pp. 566–567, Sep. 10, 2021.

[35] S. H. Kukkuhalli, “Improving digital sales through reducing friction points in the customer digital journey using data engineering and machine learning,” *International Journal of Innovative Research in Engineering & Multidisciplinary Physical Sciences*, vol. 10, no. 3, 2022.

[36] E. Y. Saatci and I. Ç. Erdönmez, “Corporate social responsibility within social media-stakeholder relations strategies: Case of turkish stock exchange companies,” *Business and Management Studies*, vol. 2, no. 1, pp. 1–11, Sep. 18, 2015.

[37] L. Wang, “Game playing in the platform society: A cultural-political economy analysis of the live streaming industry in china,” *Journal of Transcultural Communication*, vol. 2, no. 1, pp. 24–50, Sep. 1, 2022.

[38] W. Krings, R. Palmer, M. J. Harrison, and A. Inversini, “Digital media as a game-changer in b2b buyer-vendor relationships,” *Journal of Sustainable Business and Economics*, vol. 5, no. 3, pp. 27–43, Aug. 23, 2022.

[39] N. Chornopyska and L. Bolibrukh, “The influence of the covid-19 crisis on the formation of logistics quality,” *Electronic Scientific Journal Intellectualization of Logistics and Supply Chain Management 1 2020*, vol. 1, no. 2, pp. 88–98, Aug. 4, 2020.

[40] R. P. Bringula, “Taxonomy of factors influencing non-use of online shopping of students,” *Modern Applied Science*, vol. 10, no. 4, pp. 119–, Feb. 2, 2016.

[41] R. Varadarajan, “Innovating for sustainability: A framework for sustainable innovations and a model of sustainable innovations orientation,” *Journal of the Academy of Marketing Science*, vol. 45, no. 1, pp. 14–36, Aug. 18, 2015.

[42] R. S. M. Aime, G. C. Premananto, and S. Rakotoarisoa, “Marketing mix, customers’ attitude, and purchasing intention in social commerce with internet access as a moderating variable,” *Jurnal Manajemen Teori dan Terapan | Journal of Theory and Applied Management*, vol. 15, no. 1, pp. 62–76, Apr. 29, 2022.

[43] Y. Wu, Y. Zhu, and J. Zhao, “The propulsion path of synergy and linkage based on artificial intelligence and digital economy.,” *Frontiers in psychology*, vol. 13, pp. 854 542–, May 17, 2022.

[44] J. Majerova and A. Kubjatkova, “Brand value building and management on b2b markets,” *Technology transfer: innovative solutions in Social Sciences and Humanities*, vol. 3, pp. 45–48, Apr. 30, 2020.

[45] C. Homburg and D. M. Wielgos, “The value relevance of digital marketing capabilities to firm performance.,” *Journal of the Academy of Marketing Science*, vol. 50, no. 4, pp. 666–688, Apr. 20, 2022.

[46] V. Chlebovsky, “Customer solutions management (csm) empirical model based on european machine building sector experience,” *Engineering Economics*, vol. 27, no. 5, pp. 586–593, Dec. 22, 2016.

[47] R. Agnihotri, A. Kalra, H. Chen, and P. J. Daugherty, “Utilizing social media in a supply chain b2b setting: A knowledge perspective,” *Journal of Business Logistics*, vol. 43, no. 2, pp. 189–208, Aug. 24, 2021.

[48] “Article abstracts,” *Interactive Marketing*, vol. 5, no. 1, pp. 77–92, Jul. 1, 2003.

[49] B. Castillo-Abdul, M. Bonilla-del-Río, and E. Núñez-Barriopedro, "Influence and relationship between branded content and the social media consumer interactions of the luxury fashion brand manolo blahnik," *Publications*, vol. 9, no. 1, pp. 10–, Mar. 1, 2021.

[50] S.-L. Huang and Y.-J. Lee, "Diagnosing service success and failure incidents in the consumer-to-business sharing economy: A case of logistics sharing," *Journal of Global Information Management*, vol. 30, no. 2, pp. 1–16, Sep. 15, 2021.

[51] M. Butarbutar, A. Sudirman, A. P. Windarto, E. Chandra, and O. S. Sinaga, "Business strategy training for "yuni phea" sewing business housewives group in south siantar district, pematang siantar city," *Jurnal Pengabdian Masyarakat*, vol. 3, no. 2, pp. 310–327, Nov. 7, 2022.

[52] A. Surianggo, null Amelia, and null Ronald, "Analysis of the effect of online convenience dimensions on customer satisfaction and behavioral intention of tokopedia customers in surabaya," *International Journal of Research Publications*, vol. 69, no. 1, pp. 18–18, Feb. 1, 2021.

[53] E. Marzai, "Bancassurance in a digital era," *Proceedings of the International Conference on Business Excellence*, vol. 12, no. 1, pp. 601–611, May 1, 2018.

[54] H. Luo, "Study of influence on e-commerce b2c marketing strategy on sichuan consumers' intention to purchase local specialties," *Learning & Education*, vol. 10, no. 8, pp. 117–117, Jun. 20, 2022.

[55] S. Kathuria, A. Grover, V. M. E. Perego, A. Mattoo, and P. Banerjee, "Unleashing e-commerce for south asian integration," pp. 1–97, Nov. 26, 2019.

[56] S.-K. Jung, H.-i. Choe, and J.-W. Byun, "Internet marketing strategy of a wholesale tour agency in korea: Case of hana tour," *Journal of Service Science*, vol. 1, no. 1, pp. 83–104, Jun. 25, 2009.

[57] L. F. Moraes and G. Campos, "Problemas corriqueiros no e-commerce sob a percepção dos consumidores," *Marketing & Tourism Review*, vol. 5, no. 2, Apr. 20, 2021.

[58] null Vijayan*, V. Mareeswari, C. Navaneethan, S. Prasanna, and K. Yaswanth, "Calculating effective product marketing on e-commerce applications based on customer rating using big data," *International Journal of Innovative Technology and Exploring Engineering*, vol. 8, no. 12, pp. 5130–5136, Oct. 30, 2019.

[59] B. Kennedy, "Customer experience management: Industry trends an interview with bob kennedy of tealeaf technologies," *Journal of Digital Asset Management*, vol. 5, no. 3, pp. 135–147, Jun. 10, 2009.

[60] Y. Zong and M. He, "The impact imposed by brand elements of enterprises on the purchase intention of consumers-with experience value taken as the intermediary variable.," *Frontiers in psychology*, vol. 13, pp. 873 041–, Jun. 9, 2022.

[61] S. Mukherjee, "Internet based business model in india: Its challenges and opportunities," *International Journal for Research in Applied Science and Engineering Technology*, vol. 6, no. 4, pp. 1175–1181, Apr. 30, 2018.

[62] C. S. G. and I. Singh, "Optimized recommendation system for e commerce on product features and user behavior," *International Journal of Recent Technology and Engineering (IJRTE)*, vol. 8, no. 2, pp. 748–749, Jul. 30, 2019.

[63] N. G. Furtado, R. S. Canto, S. L. I. de Oliveira, and S. L. do Amaral Moretti, "Perceptions in the use of technology for payments: A study of customer behavior in food and beverage sector," *Ágora : revista de divulgação científica*, vol. 22, no. 2, pp. 4–23, Dec. 19, 2017.

[64] S. Q. A.-K. Al-Maliki, "Increasing non-oil revenue potentiality through digital commerce: The case study in ksa," *Journal of Money and Business*, vol. 1, no. 2, pp. 65–83, Nov. 2, 2021.