Learning to Reduce Cart Abandonment: A Data-Driven Personalization Framework Based on Customer 360 and Journey Friction Analytics

Owenpress:
OPEN ACCESS JOURNALS
1–14
©owenpress:
OPEN ACCESS JOURNALS
Article reuse guidelines:
https://owenpress.com/

OwenPress

Arvin Dela Cruz[1] Jericho Manalansan[2]

Abstract
Online retail environments expose shoppers to a wide variety of products, promotions, and interface patterns, yet a large fraction of initiated baskets never reach checkout completion. Cart abandonment arises from a combination of intent volatility, perceived risk, cognitive load, and friction in the interaction flow. As digital commerce platforms evolve toward real-time decisioning, there is growing attention on learning-based personalization strategies that adapt the interface and messaging to each visitor. At the same time, many organizations are assembling Customer 360 assets that aggregate multi-source data into unified profiles spanning browsing, transactions, service interactions, and marketing responses. These assets remain underexploited for fine-grained modeling of abandonment dynamics. Parallel advances in instrumentation provide detailed telemetry on micro-frictions along the journey, including latency spikes, validation errors, and micro-patterns of hesitancy. This paper develops a technical framework that combines Customer 360 representation learning with journey friction analytics and policy optimization for intervention selection. The framework casts cart abandonment reduction as a sequential decision problem over states that encode both long-term customer attributes and short-term friction signals. It supports offline learning from large historical logs under practical constraints such as delayed outcomes, partial observability, and action-dependent exposure. The presentation emphasizes formal problem definitions, model structures, and training objectives rather than any specific deployment context. The resulting formulation is intended to guide the design of data pipelines, feature learning modules, and policy optimization layers in personalization systems targeting reductions in cart abandonment.

Introduction

Cart abandonment is a persistent phenomenon in digital commerce, with a substantial share of initiated sessions terminating before payment confirmation (*1*). While some proportion reflects rational reconsideration or exploratory behavior, a considerable fraction is associated with avoidable friction in the customer journey. Such friction manifests as delayed page loads, confusing layout transitions, rigid form validation, opaque shipping or tax estimation, and other obstacles that increase cognitive effort or perceived risk. When these factors interact with heterogeneous customer preferences and constraints, they shape the probability that a given session proceeds to conversion. Understanding and influencing this probability in a principled manner is central to personalization strategies that attempt to adapt experience elements such as recommendations, discounts, nudges, and navigation shortcuts.

In parallel, many organizations have invested in Customer 360 initiatives that seek to unify disparate sources of customer data into consolidated profiles. These profiles typically aggregate information from historical transactions, browsing behavior on web and mobile channels, customer support interactions, marketing engagement, device characteristics, and sometimes third party enrichment. However, the raw existence of a Customer 360 platform does not immediately translate into actionable personalization. The challenge lies in constructing representations and models that exploit this high dimensional, sparse, and temporally structured data to predict abandonment risk and to evaluate alternative interventions. Moreover, it is necessary to connect long horizon signals captured in Customer 360 with the fine-grained, session level friction events that emerge during an ongoing visit (*2*).

Journey friction analytics attempts to quantify such micro-level impediments by transforming event logs, latency measurements, error messages, and cursor or scroll patterns into interpretable features. These features can characterize

[1]San Felipe College of Business Studies, Department of Financial Management, Bonifacio Drive, San Fernando City 2000, Philippines
[2]Luzviminda Institute of Commerce, Department of Business Administration, Mabini Street 112, Legazpi City 4500, Philippines

both systematic design issues and idiosyncratic responses of particular users. While friction analytics is often used descriptively to diagnose problems in funnels, its integration into predictive modeling and sequential decision frameworks is less mature. Yet, interventions triggered without regard to friction context may either fail to reach users at moments of high abandonment risk or be applied redundantly in already smooth journeys.

The central idea in this work is to formulate cart abandonment reduction as a learning problem over a state space that blends Customer 360 embeddings with online friction indicators. The state captures a compressed representation of who the visitor is, as inferred from historical data, and what they are experiencing in the current journey. The action space comprises candidate interventions such as showing alternative shipping options, offering contextual help, adjusting the ordering of steps, or deploying incentives. The system must choose actions based on the state, using models that are trained from historical logs where actions were selected under previous policies. The objective is to minimize the probability of abandonment subject to constraints such as limited incentive budget, interaction fatigue, and latency overhead (*3*).

This formulation raises several technical questions. First, one must define a representation learning pipeline that maps high dimensional heterogeneous Customer 360 data into compact vectors suitable for downstream policy learning. Second, one must design friction signals that are sensitive to subtle changes in behavior but robust to noise and missing data. Third, one must select modeling assumptions that link state and action to abandonment probabilities or long term value and that are estimable from logged data subject to confounding and selection bias. Finally, one must address evaluation, both offline via counterfactual estimators and online via controlled experiments, under metrics that reflect not only conversion but also revenue, satisfaction, and operational constraints.

This paper addresses these questions through a technical framework that combines representation learning, friction analytics, and policy optimization. The presentation focuses on the internal structure of the models and the learning objectives rather than empirical outcomes in any specific deployment. By grounding the discussion in explicit mathematical formulations, the framework aims to clarify design choices and trade offs when constructing systems that learn to reduce cart abandonment from rich historical data.

Background and Problem Formulation

Consider a population of users indexed by an integer variable i. Each user may generate multiple sessions, indexed by j, and within each session there is a sequence of time ordered interaction events indexed by t. The raw digital trace for session j of user i can be described at an abstract level by a sequence

$$e_{i,j,t},$$

where each event encodes attributes such as page type, action type, timestamps, device identifiers, and observed response fields. Additionally, each event may be associated with system side measurements such as latency, server side errors, and recommendation exposures. From the perspective of cart abandonment, a session is typically associated with a binary outcome

$$y_{i,j} \in \{0,1\},$$

where $y_{i,j} = 1$ indicates that the session culminates in an order placement and $y_{i,j} = 0$ indicates abandonment before payment completion.

Beyond session logs, Customer 360 infrastructure provides a high dimensional profile vector for each user, aggregating signals over prior sessions and other channels. Let

$$c_i \in \mathbb{R}^{d_c}$$

denote a raw feature vector constructed from such sources. Components of c_i may represent frequency of past purchases, recency measurements, average basket composition, engagement with marketing campaigns, historical returns behavior, payment preferences, and other attributes. The dimensionality d_c can be large, and many components may be sparse or missing for new users.

The objective in this framework is to learn a policy that maps the evolving state of a session to interventions that reduce the probability of abandonment. The state at discrete decision step t within session j of user i is denoted by (*4*)

$$s_{i,j,t}.$$

This state must summarize both the Customer 360 information and the interaction context up to step t. Interventions are elements of a finite action set, written as

$$a_{i,j,t} \in \mathcal{A},$$

where $\mathcal{A}$ may include no action as well as variations of messaging, navigation aids, and selective promotions. The system must choose actions according to a policy

$$\pi(a \mid s),$$

which defines a probability distribution over actions conditioned on the state.

The underlying dynamics of the session can be viewed through the lens of a partially observed Markov decision process. There is an unobserved latent state that captures the true intent, budget constraints, and satisfaction of the user, while the observed state $s_{i,j,t}$ is a function of this latent state and the historical events. The evolution of the session depends on both the user decisions and the system actions. From the logs, one observes sequences of the form

$$\left(s_{i,j,0}, a_{i,j,0}, s_{i,j,1}, a_{i,j,1}, \ldots, s_{i,j,T_{i,j}}, y_{i,j}\right),$$

where $T_{i,j}$ is the last decision step prior to termination of the session.

When the primary objective is reduction of cart abandonment, a natural reward formulation assigns a terminal reward of unity for conversion and zero for abandonment. For each step, the instantaneous reward can be written as

$$r_{i,j,t} = 0 \quad \text{for} \quad t < T_{i,j},$$

and

$$r_{i,j,T_{i,j}} = y_{i,j}.$$

Alternative reward structures may incorporate basket value, predicted return probability, or estimates of satisfaction, but the binary formulation provides a concise baseline (5). The cumulative reward for a session is then

$$R_{i,j} = \sum_{t=0}^{T_{i,j}} \gamma^t r_{i,j,t},$$

where $\gamma \in (0, 1]$ is a discount factor that controls the relative emphasis on immediate versus delayed outcomes.

The learning problem is complicated by several factors. First, actions in the historical logs were chosen according to a logging policy that may depend on the state. This induces selection bias in the observed action distribution and affects the estimation of counterfactual outcomes under alternative policies (6). Second, state features derived from Customer 360 profiles and friction analytics may be high dimensional and noisy, requiring representation learning techniques to produce compact embeddings. Third, the outcome of interest is sparse, as only a fraction of sessions convert, and the temporal credit assignment from intermediate actions to final outcomes is nontrivial.

To formalize the policy optimization objective, one can define the expected return under a candidate policy π as

$$J(\pi) = \mathbb{E}[R_{i,j} \mid \pi],$$

where the expectation is taken over the distribution of users, sessions, and stochastic transitions induced by the policy. The goal is to find a policy π^* that maximizes this quantity,

$$\pi^* = \arg\max_{\pi} J(\pi).$$

In practice, the maximization takes place over a parameterized class of policies with parameters θ, and one seeks an approximate maximizer θ^* using finite logged data.

Throughout the remainder of the paper, this general formulation serves as the anchor for more specific model constructions (7) (8). First, the raw Customer 360 features c_i are transformed into embeddings that can be incorporated into session states. Second, friction analytics are used to construct additional state components that capture the evolving difficulty of the journey. Third, parametric models link these states and actions to estimated rewards, enabling policy learning under constraints imposed by offline data.

Customer 360 Representation Learning

The Customer 360 feature space is often heterogeneous, combining continuous metrics, categorical attributes, count statistics, and temporal summaries. Directly feeding these raw features into policy models can lead to instability, overfitting, and difficulties in handling missingness. A more robust approach is to learn a lower dimensional embedding of the Customer 360 vector that captures salient structure while discarding noise. Let

$$z_i \in \mathbb{R}^{d_z}$$

denote such an embedding, with d_z considerably smaller than d_c. The mapping from raw features to embedding is defined by a parameterized function

$$z_i = f_\theta(c_i),$$

where f_θ may be implemented as a deep network or another flexible parametric function.

The learning objective for f_θ should encourage embeddings that are predictive of behaviors relevant to abandonment. One straightforward strategy is to use supervised learning with labeled outcomes derived from historical sessions. For each user i, one can define an aggregated binary label

$$\tilde{y}_i \in \{0, 1\},$$

indicating whether the user converted in a specified historical window (9). A simple logistic regression style objective can then be written as

$$L(\theta, w) = \sum_i \ell\big(\tilde{y}_i, \sigma(w^\top z_i)\big),$$

where w is a parameter vector, σ is the logistic function, and ℓ is the cross entropy loss. The optimization of L encourages embeddings that carry information about conversion propensity. However, such a coarse label may obscure richer structure.

To address this limitation, it is natural to consider sequence level supervision (10). Each session j of user i can be associated with features summarizing pre session Customer 360 data and an outcome $y_{i,j}$. One may then define a representation learning objective that aggregates over sessions,

$$L(\theta, v) = \sum_{i,j} \ell\big(y_{i,j}, g_v(z_i, u_{i,j})\big),$$

where $u_{i,j}$ captures session specific descriptors such as entry channel or initial intent and g_v is a prediction network with parameters v. In this construction, the embedding z_i must support discrimination across a variety of session contexts.

Beyond purely supervised objectives, one can employ self supervised or contrastive learning methods to exploit unlabeled structure in Customer 360 data. For example, different views of a user profile can be constructed by subsampling attributes or by aggregating features over disjoint time windows. The embedding function f_θ can be trained to maximize similarity between embeddings of views from the same user while minimizing similarity across users. A simple contrastive objective for a batch of pairs $(c_i^{(1)}, c_i^{(2)})$ with positive pairings can be written as

$$L(\theta) = \sum_i \Phi\Big(z_i^{(1)}, z_i^{(2)}, \{z_k^{(2)}\}_{k \neq i}\Big),$$

where $z_i^{(1)}$ and $z_i^{(2)}$ are embeddings of the two views of user i and Φ is a loss function that encourages high similarity for positive pairs relative to negatives. The exact functional form of Φ can be chosen to control the temperature and margin properties of the contrast (*11*).

Handling cold start users is an important aspect of the representation design. For new visitors with little or no historical data, the Customer 360 vector may be largely empty or populated only with coarse attributes such as location or device type. The function f_θ should map such profiles into embeddings that reflect uncertainty rather than extreme beliefs. One way to achieve this is through variational encoders that associate each user with a latent random vector. For each user i, one can posit a distribution

$$z_i \sim \mathcal{N}\big(\mu_i, \Sigma_i\big),$$

with mean and covariance given by learned functions of c_i. In downstream models, either the mean embedding μ_i or samples drawn from this distribution can be used as the state component. The variance structure Σ_i can capture the degree of uncertainty associated with sparse profiles.

A further extension arises when Customer 360 data incorporates graph structure, for example when users are linked via household relationships or shared devices. In such cases, graph neural networks can be used to propagate information across connected entities before computing embeddings. Let G denote a graph with nodes corresponding to users and edges representing relationships (*12*). A generic message passing update can be written as

$$h_i^{(l+1)} = \psi\left(h_i^{(l)}, \sum_{k \in \mathcal{N}(i)} \phi\Big(h_i^{(l)}, h_k^{(l)}\Big)\right),$$

where $h_i^{(l)}$ is the representation of node i at layer l, $\mathcal{N}(i)$ is the neighborhood of i, and ψ and ϕ are learnable functions. After a finite number of layers, the final embedding z_i is taken as $h_i^{(L)}$. This mechanism enables embeddings to reflect shared behavioral patterns across connected users in addition to individual level features.

The resulting Customer 360 embeddings become components of the session states $s_{i,j,t}$. At a given step t, the state can be written as a concatenation

$$s_{i,j,t} = \big(z_i, x_{i,j,t}\big),$$

where $x_{i,j,t}$ summarizes the intra session context up to step t. In later sections, the construction of $x_{i,j,t}$ will be augmented with journey friction features to produce a combined representation suitable for policy learning (*13*).

Journey Friction Analytics and Feature Construction

While Customer 360 embeddings capture long term behavioral and preference structure, cart abandonment is strongly influenced by real time friction within a session. Friction can be conceptualized as any factor that increases the cognitive or temporal cost of progressing to checkout. To operationalize this concept, one leverages detailed event logs and telemetry signals recorded at each interaction step. Let $d_{i,j,t}$ denote a vector of raw friction indicators associated with event $e_{i,j,t}$. Examples include client side latency measurements, counts of error messages, occurrences of validation failures, scroll depth metrics, and indicators of repeated navigation between pages without progression.

These raw signals often require aggregation and smoothing to form stable features. For each step t, one may construct cumulative measures such as

$$\Delta_{i,j,t} = \sum_{\tau=0}^{t} \delta_{i,j,\tau},$$

where $\delta_{i,j,\tau}$ is a scalar or vector representing a specific friction type at step τ. For instance, $\delta_{i,j,\tau}$ could be the time spent on a page segment that does not permit progression, or an indicator of a failed payment attempt. The cumulative measure $\Delta_{i,j,t}$ then reflects the accumulated friction experienced up to step t.

A more dynamic viewpoint models the hazard of abandonment as a function of both Customer 360 embedding and friction features. Let T denote the step at which the session terminates, and define a discrete time hazard function

$$\lambda_t(s) = \Pr\big(T = t \mid T \geq t, s_t = s\big),$$

where s_t is the state at step t. A simple parametric form for the hazard can be written as

$$\lambda_t(s) = \sigma\big(\alpha_t + u^\top z + v^\top \phi_t\big),$$

where z is the Customer 360 embedding, ϕ_t is a vector of friction features at step t, and α_t, u, and v are parameters. This structure allows friction features to influence the instantaneous risk of termination, modulated by user level characteristics.

From the hazard function, one can derive the survival probability of a session through step t as

$$S_t(s) = \prod_{\tau=0}^{t} \big(1 - \lambda_\tau(s_\tau)\big), (14)$$

where s_τ is the state at step τ. The probability that a session terminates in abandonment before conversion under a given policy depends on the sequence of hazards along the path. Friction features that increase the hazard at particular steps thereby elevate the overall abandonment risk (*15*).

In practice, friction signals can be numerous and correlated. Dimensionality reduction or representation learning applied specifically to friction data can produce compact descriptors. For each step t, consider an encoder

$$\phi_{i,j,t} = h_\eta(d_{i,j,t}, \Delta_{i,j,t}) \,,$$

where h_η is a neural network with parameters η. The combined representation $\phi_{i,j,t}$ can capture both instantaneous and cumulative friction. Training objectives for h_η may include reconstruction losses, contrastive learning where steps leading to abandonment are contrasted against those leading to conversion, or joint training with hazard or policy models.

Another useful construct is a scalar friction index that summarizes the severity of friction at each step. Define a function

$$F_{i,j,t} = q^\top \phi_{i,j,t},$$

where q is a parameter vector. The index $F_{i,j,t}$ can be calibrated so that higher values correspond to higher estimated hazard of abandonment, for example by regressing step level labels indicating whether a session terminates shortly after the step. A smoothed version can be obtained by exponential averaging, (*16*)

$$\bar{F}_{i,j,t} = \beta \bar{F}_{i,j,t-1} + (1-\beta) F_{i,j,t},$$

with $\beta \in [0, 1)$ controlling the memory of the smoothing. The smoothed index serves as a compact state component that reflects both current and recent friction.

There is also value in distinguishing structural friction, which affects many users equally, from idiosyncratic friction that interacts with Customer 360 attributes. Structural friction may arise from consistent latency spikes on specific pages or design choices that are confusing to a broad audience. Idiosyncratic friction may involve mismatches between content and user preferences, such as irrelevant recommendations or unsuitable payment methods. To model this distinction, one can decompose the friction representation into a shared component and an interaction component. For step t, write

$$\phi_{i,j,t} = \phi_t^{(0)} + \phi_{i,j,t}^{(1)},$$

where $\phi_t^{(0)}$ is a function of the page or template independent of user and $\phi_{i,j,t}^{(1)}$ captures user specific deviations. The shared component can be estimated by averaging friction signals across many users, while the interaction component can be modeled as a function of both Customer 360 embedding and event context (*17*).

The journey friction descriptors enter the state vector as

$$x_{i,j,t} = \big(g_{i,j,t}, \phi_{i,j,t}, \bar{F}_{i,j,t}\big),$$

where $g_{i,j,t}$ collects other intra session features such as number of items in the cart, presence of discounts, and time since session start. The full state becomes

$$s_{i,j,t} = \big(z_i, x_{i,j,t}\big),$$

as in the previous section. Policy models can then condition on this enriched state to choose interventions that are sensitive to both long term user attributes and instantaneous friction conditions.

Personalization Policy Learning and Optimization

Given the state representation combining Customer 360 embeddings and friction analytics, the personalization engine must learn a policy that maps states to actions (*18*). Depending on the degree of sequential structure in the interventions, one may choose between a contextual bandit formulation and a full reinforcement learning formulation. In a contextual bandit setting, each session is treated as a single decision point with context

$$x_{i,j} = \psi\big(s_{i,j,0}, \ldots, s_{i,j,T_{i,j}}\big) \,,$$

where ψ is a summary function, and the action is selected once per session. In a reinforcement learning setting, the policy selects actions at multiple steps within the session, and the cumulative reward depends on the sequence.

Consider first a parameterized stochastic policy in the reinforcement learning setting. Let $\pi_\theta(a \mid s)$ denote the probability of selecting action a in state s under parameters θ. A common choice for discrete actions is the softmax parameterization, (*19*)

$$\pi_\theta(a \mid s) = \frac{\exp\big(f_\theta(s,a)\big)}{\sum_{a' \in \mathcal{A}} \exp\big(f_\theta(s,a')\big)},$$

where $f_\theta(s,a)$ is a score function. To optimize θ, one can employ policy gradient methods that estimate the gradient of the expected return $J(\pi_\theta)$ with respect to θ.

In an online reinforcement learning setting with on policy data cite41, the policy gradient can be estimated using trajectories generated by the current policy. For offline training with logged data generated by a different policy μ, importance sampling or other off policy estimators are required. The importance ratio at step t is

$$w_{i,j,t} = \frac{\pi_\theta(a_{i,j,t} \mid s_{i,j,t})}{\mu(a_{i,j,t} \mid s_{i,j,t})},$$

assuming the logging policy probabilities are known or estimable. The cumulative importance weight for a trajectory

Friction Source	Description	Impact on Abandonment
Delayed Page Loads	High latency during key steps	Increases cognitive cost and exit likelihood
Confusing Layouts	Disorientation during navigation	Interrupts progression toward checkout
Rigid Validation	Strict or unclear form checks	Blocks user completion paths
Opaque Cost Info	Missing shipping/tax previews	Raises perceived transaction risk

Customer 360 Component	Data Type	Behavioral Insight
Purchase History	Transaction logs	Reveals spending patterns
Marketing Engagement	Email/app interactions	Indicates responsiveness to stimuli
Device and Channel Signals	Web/mobile metadata	Reflects preferred access modes
Service Interactions	Support cases	Captures past friction episodes

State Component	Origin	Modeling Role
Customer Embedding z_i	Historical multi-source data	Encodes long-horizon behavior
Session Context $x_{i,j,t}$	Real-time telemetry	Tracks evolving user intent
Friction Features $\phi_{i,j,t}$	Event logs and errors	Measures moment-level impediments

Modeling Challenge	Source	Reason for Difficulty
High Dimensionality	Customer 360 vectors	Sparse, heterogeneous attributes
Sequential Credit Assignment	Multi-step sessions	Rewards occur only at termination
Policy Confounding	Historical action bias	Logging policy influences observed data
Friction Noise	Telemetry signals	Latency and errors fluctuate rapidly

up to step $T_{i,j}$ is

$$W_{i,j} = \prod_{t=0}^{T_{i,j}} w_{i,j,t}.$$

A naive off policy estimator for the policy gradient would involve products of such weights, which can have high variance. To mitigate this, various techniques such as per decision importance sampling, clipping, or doubly robust estimators can be employed.

An alternative class of methods is based on value function approximation. Let $Q_\theta(s,a)$ denote an estimate of the expected return starting from state s taking action a and following policy π_θ thereafter. The policy can be derived by greedily or softly selecting actions with respect to Q_θ. For example, one may use

$$\pi_\theta(a \mid s) = \frac{\exp\big(Q_\theta(s,a)/\tau\big)}{\sum_{a'} \exp\big((20)Q_\theta(s,a')/\tau\big)},$$

with temperature parameter τ. Learning Q_θ from offline data requires algorithms designed for batch reinforcement learning, which control extrapolation error when evaluating actions that are infrequently observed in the logs.

When the sequential aspect of interventions is limited, a contextual bandit formulation can reduce complexity. Each decision instance is characterized by a context vector x, an action a, and an observed reward r. Logged data consist of tuples (x_n, a_n, r_n), where actions were chosen according to a logging policy μ. The objective is to find a new policy π that maximizes expected reward

$$J(\pi) = \mathbb{E}[r \mid a \sim \pi(\cdot \mid x)].$$

Off policy evaluation can be conducted via inverse propensity scoring, with estimator (*21*)

$$\hat{J}_{\text{IPS}}(\pi) = \frac{1}{N}\sum_{n=1}^{N} r_n \frac{\pi(a_n \mid x_n)}{\mu(a_n \mid x_n)}.$$

Doubly robust estimators augment this with outcome models, improving stability.

In the cart abandonment context, actions may include interventions that carry direct cost, such as financial incentives. It is therefore natural to formulate constrained optimization problems where the policy must satisfy resource budget constraints. For instance, suppose each action a has an associated expected cost $c(a)$, and the expected cost under policy π must not exceed a budget B. The constrained optimization problem is

$$\max_{\pi} J(\pi) \quad \text{subject to} \quad C(\pi) \leq B,$$

where

$$C(\pi) = \mathbb{E}[c(a) \mid a \sim \pi(\cdot \mid x)] \,.$$

One can handle this via Lagrangian relaxation, introducing a multiplier $\lambda \geq 0$ and optimizing

$$L(\pi, \lambda) = J(\pi) - \lambda\big(C(\pi) - B(22)\big).$$

Gradient based methods can be applied to update both policy parameters and the multiplier.

The friction enriched state representation plays a critical role in the effectiveness of the policy. Interventions that are costly should be concentrated in states where friction analytics indicate high hazard of abandonment and where Customer 360 embeddings suggest responsiveness. For example, if a user shows strong historical preference for a particular payment method that is unavailable in the current context, friction features may capture repeated back and forth navigation around the payment step. A policy conditioned on these features may prioritize interventions such as displaying alternative payment information or contact options rather than generic discounts.

Finally, one must consider robustness and safety. Policy models trained on offline data may propose actions that have not been observed in certain state regions, leading to uncertain outcomes. Conservative policy improvement methods attempt to restrict policy changes to remain close to the logging policy in regions of high uncertainty (*23*) (*24*). A simple measure of divergence between the candidate policy π and the logging policy μ is the Kullback Leiber divergence,

$$D(\pi \parallel \mu) = \mathbb{E}\left[\sum_{a} \pi(a \mid s) \log \frac{\pi(a \mid s)}{\mu(a \mid s)}\right].$$

By constraining this divergence below a threshold, one can limit the degree of extrapolation. Alternatively, uncertainty estimates in value predictions can be used to penalize actions with high epistemic uncertainty (*25*).

Experimental Design and Evaluation Considerations

Evaluating a personalization framework that aims to reduce cart abandonment involves both offline and online components. Offline evaluation uses historical logs to estimate the performance of candidate policies without deploying them in production, while online evaluation relies on controlled experiments such as randomized trials. The design of offline evaluation must address the counterfactual nature of the problem: the logs record outcomes under the logging policy μ, but one wishes to estimate outcomes under a different policy π.

In the contextual bandit setting, off policy evaluation methods such as inverse propensity scoring and doubly robust estimators can be applied directly. The basic requirement is knowledge or estimation of the logging policy probabilities $\mu(a \mid x)$. In complex systems, the logging policy may be a mixture of deterministic rules and stochastic components, and reconstructing its action probabilities can be nontrivial. When action propensities are unknown, one may approximate them using a parametric model fitted to the observed action frequencies. The accuracy of such approximations affects the bias and variance of off policy estimators.

In sequential settings, where multiple interventions may occur within a session, trajectory based estimators are needed (*26*). Off policy evaluation can be posed in terms of estimating the expected discounted return under the new policy using importance sampling over entire trajectories. However, as noted previously, naive importance sampling suffers from high variance, especially when the new policy deviates substantially from the logging policy. Per decision estimators, where importance ratios are applied stepwise, can reduce variance but still face challenges in long horizons. Recent advances in model based and model assisted off policy evaluation can be leveraged by learning transition models and using them to simulate the effect of new policies, combined with adjustments to control model bias.

From a metrics perspective, the most direct performance indicator is the cart abandonment rate, defined as the proportion of sessions that do not culminate in purchase. If N sessions are observed, with conversion indicators $y_1, \ldots, y_N$, the empirical abandonment rate is

$$\hat{A} = 1 - \frac{1}{N} \sum_{n=1}^{N} y_n.$$

A personalization policy is expected to reduce $\hat{A}$, but evaluation must consider additional metrics. These include revenue per session, average order value, distribution of discount usage, and indicators of customer satisfaction such as post purchase behaviors. Excessive use of incentives to reduce abandonment may erode margin and create expectations of constant discounts.

Offline evaluation of the interplay between cart abandonment and revenue can be facilitated by modeling not only the binary conversion variable but also the order value conditional on conversion. Let $v_{i,j}$ denote the order value for session j of user i, with $v_{i,j} = 0$ for abandoned sessions. The expected revenue per session under policy π is (*27*)

$$R(\pi) = \mathbb{E}[v \mid \pi] \,.$$

Off policy estimators for $R(\pi)$ can be constructed analogously to those for conversion probability. When order

Evaluation Mode	Key Requirement	Limitation
Offline OPE (IPS/DR)	Needs logging propensities $\mu(a \mid x)$	Sensitive to misspecified propensities, high variance
Trajectory Based OPE	Stepwise or full ratio estimation	Variance grows with horizon length
Model Assisted OPE	Requires learned transition model	Model bias affects counterfactual estimates

Metric	Definition	Consideration
Abandonment Rate	$1 - \frac{1}{N}\sum_n y_n$	Must balance with revenue effects
Revenue per Session	$\mathbb{E}[v]$ under π	Requires variance control for skewed order values
Friction Distribution Shift	Change in friction indices across steps	Reveals mechanisms of intervention impact

Risk Area	Description	Evaluation Need
Fairness Drift	Unequal incentives or nudges across segments	Segment stratification of abandonment and revenue
Data Drift	Shifts in device mix or friction patterns	Continuous monitoring and retraining triggers
Operational Load	Latency added by personalization logic	Stress tests and system level profiling

Online Experiment Parameter	Role	Issue
Traffic Allocation	Assigns users to policies	Requires stability across segments
Endpoint Measurement	Captures conversion, revenue, friction shifts	Must track both primary and secondary outcomes
Longitudinal Consistency	User level randomization persistence	Important for multi visit behavioral effects

values are highly skewed, variance reduction techniques or modeling of the conditional distribution of order values may be needed.

Online evaluation typically involves an A B test or multi arm experiment where a fraction of traffic is assigned to the new policy and the remainder to control conditions. The assignment can be random at the user or session level, with stratification to ensure balance across segments. The primary endpoint is an estimate of difference in abandonment rate between treatment and control, along with confidence intervals. The friction analytics can also be used as secondary endpoints to understand how interventions change the distribution of friction indices and which parts of the journey are most affected.

The presence of Customer 360 and friction based personalization raises potential concerns about fairness and unintended biases (28). For instance, certain customer segments may be systematically offered more generous incentives due to historical behavior patterns, while others may be exposed to more intrusive prompts. Evaluation should therefore include analyses of policy impact across segments defined by stable attributes such as geography or tenure, even when the optimization objective does not explicitly encode fairness constraints. This can be approached by stratified estimation of metrics, where one computes separate estimates of abandonment rate and revenue for each segment and examines differences under new policies.

Robustness to data drift is another evaluation concern. Customer preferences, device distributions, and environmental conditions evolve over time, affecting both Customer 360 data and friction analytics. A policy trained on historical data under certain conditions may perform differently when deployed under changed conditions. Monitoring schemes can track deviations in input feature distributions and outcome metrics over time, supporting retraining or adaptation when necessary. Scenario testing using counterfactual simulations that perturb friction patterns or demand distributions can provide additional insight into policy stability.

Finally, practical evaluation must account for operational constraints such as latency and system load (29). Personalization decisions based on complex models may introduce additional processing time, potentially creating new sources of friction. Measurements of end to end latency under load, combined with stress testing of model serving components, can ensure that the personalization framework does not degrade the baseline experience. When necessary, simplified or cached decision logic can be used in latency sensitive parts of the journey, while more complex models operate in less critical paths.

System Integration and Deployment Considerations

Implementing the proposed framework in a real commerce platform requires integrating several components: data ingestion pipelines, Customer 360 representation learning infrastructure, friction analytics computation, model training services, and online inference endpoints. Each component has design choices that influence both performance and maintainability. Although specifics vary across organizations, certain abstract considerations can be discussed without reference to any particular system.

Data ingestion must unify logs from web and mobile applications, backend services, payment gateways, and marketing systems. Events must be linked to user identifiers, which may involve probabilistic resolution when multiple identifiers are present across channels. The construction of Customer 360 vectors c_i relies on periodic aggregation jobs that summarize interactions over sliding time windows and produce features such as recency and frequency. These features must be consistent between offline training and online inference to avoid training serving skew. At the same time, certain state components such as friction descriptors are inherently real time and must be computed on the fly during sessions (*30*).

The representation learning models for Customer 360 may be trained in batch mode on large historical datasets. Training can be scheduled at regular intervals, with embeddings updated using the latest data. For users with evolving behavior, one may employ incremental updates that adjust embeddings based on new events without retraining the entire model. This can be modeled as a state space system where embeddings follow dynamics

$$z_{i,t+1} = Az_{i,t} + Bu_{i,t} + \epsilon_{i,t},$$

where $u_{i,t}$ summarizes new events in a short interval and $\epsilon_{i,t}$ is a noise term. Parameters A and B can be learned from data, enabling efficient recursive updates of $z_{i,t}$ as fresh events arrive.

Journey friction analytics require instrumentation at the application level to capture signals such as dwell times, focus changes, and error occurrences. These signals must be transmitted to backend systems with minimal overhead and latency. Feature computation functions h_η and smoothing mechanisms must operate under real time constraints, often on streaming infrastructures that aggregate and transform event streams. The resulting friction representations $\phi_{i,j,t}$ and indices $\bar{F}_{i,j,t}$ must be available to the personalization engine at each decision point.

Online inference of the policy π_θ involves retrieving the current Customer 360 embedding, computing friction based features for the current state, and evaluating the policy network to obtain action probabilities. Latency budgets dictate constraints on the complexity of the policy network and the depth of feature transformations. Caching strategies can mitigate some costs by precomputing parts of the state representation. For example, the Customer 360 embedding z_i can be cached per user, leaving only friction computations and light transformations at session time.

Deployment also requires mechanisms for safe experimentation and gradual rollout (*31*). Traffic allocation components must support routing requests to different policy variants based on experimental assignments. Logging must capture sufficient information to reconstruct states, actions, and rewards for offline analyses, including the action probabilities under the deployed policies. Longitudinal experiments, where effects of personalization policies accumulate over repeated visits, may require assignment at the user level rather than the session level to preserve consistency.

Monitoring is an essential component of deployment. Key performance indicators such as abandonment rate, conversion, revenue per session, and incentive spend must be tracked with appropriate time resolution. Additional diagnostics specific to the personalization framework include distributions of action frequencies across states, estimates of off policy metrics for candidate policies, and alerts for deviations in input feature distributions. For friction analytics, dashboards that display friction indices across journey steps and segments can reveal shifts in system performance or design.

Finally, governance considerations are relevant when Customer 360 data include sensitive attributes. Even in the absence of explicit protected attributes, correlation structures may lead to differential treatment of segments (*32*). Governance processes may therefore require that certain attributes be excluded from the state representation or that constraints be imposed on policy behavior. From a technical standpoint, one can enforce constraints on action probabilities for specific segments or include fairness penalties in the optimization objective. However, such interventions may trade off some efficiency in abandonment reduction against more balanced treatment, and careful analysis is required to quantify these trade offs.

Statistical and Theoretical Considerations

A more formal understanding of the learning problem described earlier requires explicit characterization of the data generating process, the function classes used for representation and policy learning, and the sources of estimation error that arise from finite logged data. The personalization engine observes a collection of trajectories generated by the logging policy, and these trajectories contain both Customer 360 information and journey friction analytics. Let the random variables S, A, R, and S' denote respectively the state, action, immediate reward, and next state at a generic decision step, where the state includes both the Customer 360 embedding and the friction augmented session descriptors. Under a fixed logging policy μ, the joint distribution of (S, A, R, S') can be treated as stationary for analytical purposes, even though in practice it is induced by a complex environment. A dataset of n decision points can be written abstractly as independent draws (S_k, A_k, R_k, S'_k) for $k = 1, \ldots, n$, acknowledging that strict independence is an approximation when data are sampled from trajectories but can be justified via standard mixing arguments for long sessions.

The performance of a parameterized policy π_θ with parameters θ is measured by its expected discounted return. Using the shorthand $J(\theta)$ for this performance, one can write a compact definition as

$$J(\theta)(33) = \mathbb{E}[R_{\pi_\theta}],$$

where the expectation is taken over trajectories generated when actions follow π_θ. Directly estimating $J(\theta)$ from data generated by the logging policy is not possible without additional assumptions, because the data do not contain outcomes for actions that were not chosen. Off policy evaluation methods introduce estimators $\hat{J}_n(\theta)$ that approximate $J(\theta)$ by reweighting observed rewards or by combining reward models with propensity information. The discrepancy between $J(\theta)$ and $\hat{J}_n(\theta)$ can be decomposed into approximation error, due to model misspecification, and estimation error, due to finite sample fluctuations.

In the contextual bandit specialization, where a single decision per session is modeled, the analysis can be made more explicit. Let X denote the context, which in this setting contains the Customer 360 embedding z and a summary of friction analytics over the session. The logging policy μ draws an action A based on X, and a reward R is observed. For a candidate policy π_θ, the inverse propensity score estimator is

$$\hat{J}_n(\theta) = \frac{1}{n}\sum_{k=1}^{n} R_k W_k(\theta),$$

where the importance weight at sample k is

$$W_k(\theta) = \frac{\pi_\theta(A_k \mid X_k)}{\mu(A_k \mid X_k)}.$$

Assuming the logging propensities $\mu(A_k \mid X_k)$ are known and strictly positive, the estimator is unbiased in the idealized setting. However, the variance of $\hat{J}_n(\theta)$ can be large when $W_k(\theta)$ varies strongly, especially when π_θ places high probability on actions that were rarely selected under μ.

To quantify this behavior, consider the random variable $W(\theta)$ representing the importance weight at a generic decision point, and denote its variance by $\sigma^2(\theta)$. A simple application of concentration inequalities yields that, with high probability over the sample of size n, the deviation of the estimator from its expectation satisfies

$$\big|\hat{J}_n(\theta) - J_{\text{IPS}}(\theta)\big|(34) \le \epsilon_n(\theta),$$

where the random deviation term $\epsilon_n(\theta)$ scales on the order of $\sigma(\theta)/\sqrt{n}$. Here $J_{\text{IPS}}(\theta)$ denotes the ideal value obtained if expectations under the logging distribution were computed exactly. This relation highlights that variance control is central in offline policy optimization, and that design choices for the policy class, such as restricting deviations from the logging policy, have direct implications for the stability of estimators.

One way to formalize this restriction is to constrain the divergence between π_θ and μ. A common divergence measure is the Kullback Leiber divergence, which in this context can be summarized as

$$D(\theta) = \mathbb{E}\big[d(\theta, S)\big],$$

where, for a single state s, the local divergence is defined as

$$d(\theta, s) = \sum_a \pi_\theta(a \mid s) \log \frac{(35)\pi_\theta(a \mid s)}{\mu(a \mid s)}.$$

The expectation in $D(\theta)$ is with respect to the distribution of states under the logging process. Restricting θ to a feasible set where $D(\theta)$ remains below a chosen threshold enforces a trust region around the logging policy and indirectly controls the spread of importance weights, thereby reducing variance. The trust region interpretation provides a bridge between heuristic regularization schemes and quantities that appear in generalization bounds.

Beyond variance, generalization error is influenced by the complexity of the policy class. In the contextual bandit case, this can be analyzed through Rademacher complexity of the function class that maps contexts and actions to predicted rewards or propensity adjusted quantities. For a class of real valued functions $\mathcal{F}$ defined on the space of contexts and actions, the empirical Rademacher complexity on a sample $\{(X_k, A_k)\}_{k=1}^n$ is defined as

$$\hat{\mathcal{R}}_n(\mathcal{F}) = \mathbb{E}\left[\sup_{f \in \mathcal{F}} \frac{1}{n}\sum_{k=1}^{n} \sigma_k f(X_k, A_k)\right],$$

where $\sigma_1, \ldots, \sigma_n$ are independent Rademacher variables taking values ± 1. When function values represent clipped importance weighted rewards, standard arguments lead to bounds of the form (*36*)

$$\big|J(\hat{\theta}) - J(\theta^*)\big| \le C_1 \hat{\mathcal{R}}_n(\mathcal{F}) + C_2\sqrt{\frac{\log(1/\delta)}{n}},$$

holding with probability at least $1 - \delta$, where C_1 and C_2 are constants that depend on reward and weight ranges, and θ^* is an optimal parameter in the class. While such inequalities abstract away many details, they indicate that the representation of states through Customer 360 embeddings and friction features enters the analysis via the induced complexity of the function class. Richer state representations enable more expressive policies but also enlarge $\mathcal{F}$, potentially increasing the complexity term.

In the sequential decision setting, the analysis often proceeds through value functions. For a fixed policy π, the state value function $V^\pi(s)$ is defined through the relation

$$V^\pi(s) = \mathbb{E}\big[R \mid S_0 = s\big],$$

where the expectation is taken over trajectories starting at state s and following policy π. Likewise, the action value function $Q^\pi(s,a)$ satisfies

$$Q^\pi(s,a) = \mathbb{E}\big[(37)R \mid S_0 = s, A_0 = a\big].$$

Approximate dynamic programming methods aim to learn parameterized approximations Q_θ that are close to Q^π in some norm. Denote the approximation error in the sup norm by

$$\Delta_\infty = \sup_{s,a}\big|Q_\theta(s,a) - Q^\pi(s,a)\big|.$$

Under standard assumptions on the discount factor and transition dynamics, one can derive bounds that relate the performance of the greedy policy with respect to Q_θ to the performance of the optimal policy. The classical form of such results asserts that the performance loss is at most proportional to Δ_∞ divided by a factor depending on the discount. Although the constants are sensitive to the exact assumptions, the general message is that accurate approximation of value functions over the space of Customer 360 and friction augmented states is essential for reliable policy improvement.

The presence of friction analytics in the state representation introduces additional structure into these approximations. Friction features are often highly informative about imminent abandonment, meaning that local variations in friction values can produce large changes in the conditional distribution of future rewards. In value function terms, this implies high curvature of Q^π along directions associated with friction variables. Function classes used to approximate Q^π must therefore be sufficiently flexible along these directions to avoid large approximation error. At the same time, friction features can be noisy and transient, especially when derived from detailed telemetry such as cursor movements. From a statistical standpoint, this leads to a trade off between expressiveness and variance: aggressive fitting of friction features may capture noise rather than signal, while excessive smoothing may obscure true risk patterns (*38*).

A related issue concerns identifiability and confounding. Because logged data are generated under a particular logging policy, the distribution of states and actions may reflect systematic choices in the deployment environment. For example, the logging policy may already incorporate simple heuristics that trigger interventions in sessions with high friction indices. When attempting to learn a new policy, it is necessary to account for the fact that observed outcomes in high friction states may be disproportionately associated with certain actions. The formal condition often invoked is a form of sequential ignorability, which in a single step abstraction can be written as

$$R(a) \perp A \mid S,$$

where $R(a)$ denotes the potential reward that would have been observed had action a been taken, and the symbol $\perp$ indicates conditional independence. This condition asserts that, given the state including Customer 360 and friction descriptors, the choice of action is as good as random. In practice, the condition may not hold exactly, but friction features improve its plausibility by absorbing some of the unobserved determinants of abandonment, such as impatience induced by latency (*39*).

From a learning perspective, the assumption suggests that the state representation should be rich enough that residual dependence between actions and unmeasured drivers of outcomes is limited (*40*). Customer 360 embeddings contribute historical context such as sensitivity to shipping cost or preferred payment methods, while friction analytics reflect real time factors such as temporary connectivity problems. Together, these features expand the conditioning set in the potential outcome relation and can reduce confounding, although they do not guarantee its absence. When the assumption is only approximately valid, biases in off policy evaluation and policy learning can be bounded under additional structural constraints, such as limits on the strength of residual dependence or monotonicity properties of intervention effects.

Another dimension where theory informs design is sample allocation across users and sessions. In many commerce environments, a small subset of users generates a large fraction of sessions and conversions. If policy learning gives disproportionate weight to such users, it may over optimize for frequent visitors while underrepresenting rare segments. Formally, the data distribution may place high mass on a subset of the state space corresponding to frequent users, and function approximation will be more accurate in these regions. Generalization guarantees that rely on uniform convergence can be refined by considering margin conditions or local Rademacher complexities that adapt to the effective support of the distribution. In qualitative terms, this suggests that the expressive capacity of models can be focused where data density is high, while regularization and conservative policy updates are preferable in sparsely observed regions, particularly for combinations of Customer 360 profiles and friction patterns that appear rarely (*41*).

Representation learning for Customer 360 embeddings also has theoretical implications. When embeddings are trained to optimize a supervised objective related to conversion or revenue, they compress the original high dimensional features into a lower dimensional manifold that preserves information relevant to the chosen labels. Under standard assumptions on the encoder and decoder classes, one can relate the mutual information between original features and embeddings to reconstruction error and prediction performance. Writing the embedding as $Z = f_\theta(C)$, where C denotes the raw Customer 360 vector, one is often interested in ensuring that Z retains sufficient information about outcomes Y related to abandonment. While exact information theoretic quantities are difficult to compute in practice, regularizers that penalize excessive variance in embeddings or encourage smoothness with respect to input features can be interpreted as controlling effective information capacity. When friction features are concatenated to embeddings at session time, the combined state might be seen as augmenting this manifold with directions corresponding to short term variability.

The choice of how aggressively to compress Customer 360 data has a direct effect on the complexity of downstream policy classes. A very low dimensional embedding simplifies policy learning, potentially tightening generalization bounds, but may omit subtle patterns that distinguish between users who respond differently to interventions under similar friction conditions. A higher dimensional embedding captures more structure but increases the size of the hypothesis space. In the bandit analysis, this manifests as a larger function class $\mathcal{F}$ and larger Rademacher complexity. In the value function perspective, high dimensional embeddings require more parameters in approximation architectures and more data to accurately estimate value surfaces across the state space. Theoretical trade offs of this kind are typically expressed through error decompositions where total loss is the sum of approximation error, estimation error, and optimization error, each depending on representation choices.

Optimization dynamics form another aspect of theoretical interest. Policy and value function parameters are often trained with variants of stochastic gradient descent based on mini batches of logged data. Even when loss functions are convex in certain parameters, off policy objectives with importance weighting introduce non convexities and heavy tailed gradients. Let θ_t denote the parameter vector at iteration t, and write a generic update as

$$\theta_{t+1} = \theta_t - \eta_t G_t,$$

where η_t is the learning rate and G_t is a stochastic gradient estimate. When gradients incorporate products of rewards and importance weights, the norm of G_t may fluctuate considerably, especially if weights are only loosely controlled. Stability analyses that bound the expected squared norm of gradients can be used to derive conditions on learning rates that ensure convergence to stationary points of the empirical objective. In practice, gradient clipping and normalization of importance weights are employed to approximate such conditions, and friction aware state representations can reduce the need for extreme weighting by bringing the logging and target policies closer in the augmented space.

Finally, there is the question of how theoretical analyses grounded in idealized assumptions translate into design guidelines for real systems. The preceding discussion suggests several broad principles. First, state representations that pragmatically combine Customer 360 embeddings with friction analytics can make assumptions such as sequential ignorability more plausible and reduce confounding, thereby improving the validity of offline evaluation (*42*). Second, the complexity of policy and value function classes should be controlled in light of available data volume and the dimensionality of state representations, with explicit attention to variance of importance weights and concentration bounds. Third, constraints on divergence from the logging policy and conservative update schemes can temper extrapolation in sparsely observed regions of the state space, particularly where combinations of user attributes and friction patterns are rare. Fourth, optimization procedures should account for the heavy tailed nature of importance weighted gradients to maintain stability.

While these principles do not yield closed form sample complexity formulas tailored to every deployment, they connect the abstract machinery of statistical learning theory and dynamic programming with the concrete challenges of using Customer 360 data and journey friction analytics for personalization aimed at cart abandonment reduction. This connection provides a structured perspective on why certain design patterns, such as trust region constraints, representation learning with regularization, and friction sensitive policies, tend to perform more robustly when trained from large but imperfect historical logs.

Conclusion

This paper presented a technical framework for learning to reduce cart abandonment through personalization based on Customer 360 representations and journey friction analytics. The formulation treats cart abandonment as a sequential decision problem in which states encode both long term customer attributes and short term friction conditions. Customer 360 data are transformed into compact embeddings through supervised, self supervised, or variational representation learning, while friction analytics convert detailed interaction telemetry into structured descriptors and indices. These components together define a state representation that supports the construction of policies mapping states to interventions intended to influence progression toward checkout (*43*).

The framework accommodates both contextual bandit and reinforcement learning perspectives, reflecting varying degrees of sequential intervention structure. Policy optimization from offline data is enabled by off policy evaluation and learning methods that account for logged action propensities and selection effects. Resource constraints, such as limits on incentive usage, can be incorporated through constrained optimization formulations. Safety and robustness considerations motivate conservative policy improvement strategies and uncertainty aware value estimation.

Evaluation and deployment aspects are integral to the overall framework. Offline evaluation relies on counterfactual estimators of abandonment and revenue metrics under candidate policies, while online evaluation uses controlled experiments to estimate treatment effects. The inclusion of fairness analyses and monitoring of data drift supports more stable and transparent deployments. System integration considerations emphasize the importance of consistent feature computation between offline and online environments, efficient streaming computation of friction features, and low latency policy inference.

Overall, the proposed framework illustrates how Customer 360 assets and journey friction analytics can be combined within a principled learning based approach to personalization targeted at cart abandonment reduction. While concrete deployments require adaptation to specific platforms, data characteristics, and governance requirements, the mathematical and system level constructs described here provide a basis for designing and analyzing such systems in a structured way (*44*).

References

1. M. C. Johlke, D. F. Duhan, R. D. Howell, and R. W. Wilkes, "An integrated model of sales managers' communication practices," *Journal of the Academy of Marketing Science*, vol. 28, pp. 263–277, 4 2000.
2. M. J. Tippins and R. S. Sohi, "It competency and firm performance: is organizational learning a missing link?," *Strategic Management Journal*, vol. 24, pp. 745–761, 6 2003.
3. J. Ganesh, K. E. Reynolds, and M. Luckett, "Retail patronage behavior and shopper typologies: a replication and extension using a multi-format, multi-method approach," *Journal of the Academy of Marketing Science*, vol. 35, pp. 369–381, 5 2007.
4. D. T. Guion, D. L. Scammon, and A. L. Borders, "Weathering the storm: A social marketing perspective on disaster preparedness and response with lessons from hurricane katrina," *Journal of Public Policy & Marketing*, vol. 26, pp. 20–32, 4 2007.
5. E. Carrier and L. Weatherford, "Implementation of mnl choice models in the passenger origin-destination simulator (pods)," *Journal of Revenue and Pricing Management*, vol. 14, pp. 400–407, 10 2015.
6. S. W. Jeong, B. Jin, J.-E. Chung, and H. Yang, "Network evolution and cultivation patterns during the internationalization process: case analyses from korean smes," *Asian Business & Management*, vol. 16, pp. 323–351, 10 2017.
7. C. L. Newman, M. D. Cinelli, D. W. Vorhies, and J. A. G. Folse, "Benefitting a few at the expense of many? exclusive promotions and their impact on untargeted customers," *Journal of the Academy of Marketing Science*, vol. 47, pp. 76–96, 8 2018.
8. 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.
9. J. M. Balmer, K. Fukukawa, and E. R. Gray, "The nature and management of ethical corporate identity: A commentary on corporate identity, corporate social responsibility and ethics," *Journal of Business Ethics*, vol. 76, pp. 7–15, 2 2007.
10. H. Antlov, D. W. Brinkerhoff, and E. Rapp, "Civil society capacity building for democratic reform: Experience and lessons from indonesia," *VOLUNTAS: International Journal of Voluntary and Nonprofit Organizations*, vol. 21, pp. 417–439, 5 2010.
11. A. O. Laplume, B. Petersen, and J. M. Pearce, "Global value chains from a 3d printing perspective," *Journal of International Business Studies*, vol. 47, pp. 595–609, 1 2016.
12. T. A. Burnham, J. Frels, and V. Mahajan, "Consumer switching costs: A typology, antecedents, and consequences," *Journal of the Academy of Marketing Science*, vol. 31, pp. 109–126, 4 2003.
13. M. Malika, D. Maheswaran, and S. P. Jain, "Perceived financial constraints and normative influence: discretionary purchase decisions across cultures," *Journal of the Academy of Marketing Science*, vol. 50, pp. 1–20, 10 2021.
14. M. Schreiner, P. Kale, and D. Corsten, "What really is alliance management capability and how does it impact alliance outcomes and success," *Strategic Management Journal*, vol. 30, pp. 1395–1419, 7 2009.
15. A. Adhikary, K. S. Diatha, S. B. Borah, and A. Sharma, "How does the adoption of digital payment technologies influence unorganized retailers' performance? an investigation in an emerging market," *Journal of the Academy of Marketing Science*, vol. 49, pp. 882–902, 3 2021.
16. M. Wedel and P. Kannan, "Marketing analytics for data-rich environments," *Journal of Marketing*, vol. 80, pp. 97–121, 11 2016.
17. N. Arora, X. Drèze, A. Ghose, J. D. Hess, R. Iyengar, B. Jing, Y. V. Joshi, V. Kumar, N. H. Lurie, S. A. Neslin, S. Sajeesh, M. Su, N. B. Syam, J. S. Thomas, and Z. J. Zhang, "Putting one-to-one marketing to work: Personalization, customization, and choice," *Marketing Letters*, vol. 19, pp. 305–321, 9 2008.
18. T. McCausland, E. B. King, L. N. Bartholomew, R. P. Feyre, A. S. Ahmad, and L. M. Finkelstein, "The technological age: The effects of perceived age in technology training," *Journal of Business and Psychology*, vol. 30, pp. 693–708, 2 2015.
19. A. Gnanlet and H. M. Yayla-Kullu, "Impact of national culture on the quality of information delivery in services," *Service Business*, vol. 8, pp. 135–169, 6 2013.
20. B.-C. Kim and J. P. Choi, "Customer information sharing: Strategic incentives and new implications," *Journal of Economics & Management Strategy*, vol. 19, pp. 403–433, 5 2010.

21. D. E. Schultz and P. J. Kitchen, “Managing the changes in corporate branding and communication: Closing and re-opening the corporate umbrella,” *Corporate Reputation Review*, vol. 6, pp. 347–366, 1 2004.
22. C. Bradler and S. Neckermann, “The magic of the personal touch: Field experimental evidence on money appreciation as gifts,” *The Scandinavian Journal of Economics*, vol. 121, pp. 1189–1221, 3 2019.
23. D. J. Teece, “Managing the university: Why “organized anarchy” is unacceptable in the age of massive open online courses:,” *Strategic Organization*, vol. 16, pp. 92–102, 9 2017.
24. 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.
25. B. M. Noone and T. A. Maier, “A decision framework for restaurant revenue management,” *Journal of Revenue and Pricing Management*, vol. 14, pp. 231–244, 4 2015.
26. S. A. Neslin, K. Jerath, A. V. Bodapati, E. T. Bradlow, J. Deighton, S. Gensler, L. Lee, E. Montaguti, R. Telang, R. Venkatesan, P. C. Verhoef, and Z. J. Zhang, “The interrelationships between brand and channel choice,” *Marketing Letters*, vol. 25, pp. 319–330, 6 2014.
27. S. Motyka, D. Grewal, E. Aguirre, D. Mahr, K. de Ruyter, and M. Wetzels, “The emotional review–reward effect: how do reviews increase impulsivity?,” *Journal of the Academy of Marketing Science*, vol. 46, pp. 1032–1051, 5 2018.
28. A. Muk, “Perceptions of mobile coupons: A cross-national study,” *Journal of Direct, Data and Digital Marketing Practice*, vol. 13, pp. 311–324, 4 2012.
29. S. Spiekermann, A. Acquisti, R. Böhme, and K. L. Hui, “The challenges of personal data markets and privacy,” *Electronic Markets*, vol. 25, pp. 161–167, 4 2015.
30. C. C. Chen, B. D. Medlin, and R.-S. Shaw, “A cross-cultural investigation of situational information security awareness programs,” *Information Management & Computer Security*, vol. 16, pp. 360–376, 10 2008.
31. S. Park, H.-W. Lee, and C. Nite, “Underscoring trainer’s lack of effort or talent in service failure of personal fitness training: customers’ negative word-of-mouth response,” *Service Business*, vol. 17, pp. 1025–1050, 9 2023.
32. J. Mahlich, “Patents and performance in the japanese pharmaceutical industry: An institution-based view,” *Asia Pacific Journal of Management*, vol. 27, pp. 99–113, 1 2009.
33. S. Helm, V. Ligon, T. Stovall, and S. V. Riper, “Consumer interpretations of digital ownership in the book market,” *Electronic Markets*, vol. 28, pp. 177–189, 4 2018.
34. V. Kumar, B. Rajan, S. Gupta, and I. D. Pozza, “Customer engagement in service,” *Journal of the Academy of Marketing Science*, vol. 47, pp. 138–160, 10 2017.
35. A. Acquisti, C. R. Taylor, and L. Wagman, “The economics of privacy,” *Journal of Economic Literature*, vol. 54, pp. 442–492, 6 2016.
36. D. Bergemann and A. Bonatti, “Markets for information: An introduction,” *Annual Review of Economics*, vol. 11, pp. 85–107, 8 2019.
37. M. Duggan, P. Healy, and F. S. Morton, “Providing prescription drug coverage to the elderly: America’s experiment with medicare part d,” *The journal of economic perspectives : a journal of the American Economic Association*, vol. 22, pp. 69–92, 10 2008.
38. A. Godfrey, K. Seiders, and G. B. Voss, “Enough is enough! the fine line in executing multichannel relational communication,” *Journal of Marketing*, vol. 75, pp. 94–109, 7 2011.
39. 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.
40. B. Barry, M. Olekalns, and L. Rees, “An ethical analysis of emotional labor,” *Journal of Business Ethics*, vol. 160, pp. 17–34, 5 2018.
41. T. Demirciftci and A. Belarmino, “A cross-cultural study of competitive intelligence in revenue management,” *Journal of Revenue and Pricing Management*, vol. 21, pp. 691–699, 5 2022.
42. H. J. Schau, A. M. Muñiz, and E. J. Arnould, “How brand community practices create value,” *Journal of Marketing*, vol. 73, pp. 30–51, 9 2009.
43. E. Hermann, G. Y. Williams, and S. Puntoni, “Deploying artificial intelligence in services to aid vulnerable consumers,” *Journal of the Academy of Marketing Science*, vol. 52, pp. 1431–1451, 11 2023.
44. Y. Wu, S. Balasubramanian, and V. Mahajan, “When is a preannounced new product likely to be delayed,” *Journal of Marketing*, vol. 68, pp. 101–113, 4 2004.