ORIGINAL ARTICLE

Reducing Friction in the B2C Digital Purchase Funnel: A Data-Driven Customer Journey Modeling and Intervention Framework for Revenue Optimization

Karim Nabil Haddad[1], Rami Georges Khoury[2], and Tarek Elias Mansour[3]

[1]Lebanese International University, Department of Business Administration, 21 Sami El Solh Avenue, Badaro District, Beirut, Lebanon
[2]Holy Spirit University of Kaslik, Faculty of Economics and Business, Jounieh Highway, Kaslik, Mount Lebanon, Lebanon
[3]Beirut Arab University, Department of Management and Finance, 8 Omar Bayhum Street, Tarik El Jdideh, Beirut, Lebanon

ABSTRACT

Digital consumer purchase journeys increasingly unfold within dense, multi-touch environments where channel proliferation, device fragmentation, and rich merchandising content create opportunities for both engagement and friction. At each stage, small operational or cognitive impediments alter the balance between continuation and abandonment, shaping revenue and customer experience in ways that are measurable but often left unstructured. Friction emerges from latency, information overload, verification demands, input complexity, navigation inconsistency, and perceived risk, interacting with individual sensitivity, intentions, and context. Many organizations deploy localized user interface changes or incentives without a unifying modeling framework that links observed frictions, interventions, and revenue outcomes over heterogeneous paths. This paper presents an integrated, data-driven framework for modeling the B2C digital purchase funnel as a controlled stochastic process combined with causal estimation and constrained optimization. The framework formalizes friction as parameterized costs embedded in state transitions, incorporates behavioral regularization, and supports heterogeneous treatment effect estimation for candidate interventions such as content adjustments, sequencing rules, and incentive strategies. It is designed to operate under typical data limitations, including incomplete identity resolution, sparse conversions, seasonality, and noisy attributions. The resulting structure enables calibrated, uncertainty-aware policies for reducing friction and reallocating intervention budgets across segments, channels, and devices without overstating precision. The discussion emphasizes internal coherence between measurement, modeling, decision rules, and governance, providing a neutral basis on which organizations can align engineering, design, and marketing actions with empirically grounded expectations of incremental revenue impact.

1 | Introduction

Digital B2C purchase funnels operate within environments characterized by dense information, heterogeneous devices, multiple acquisition channels, and rapidly evolving interface patterns [1]. Within these environments, customers move from initial exposure to an offering through sequences of micro-decisions that include continuing to browse, refining preferences, adding items to a cart, providing personal and payment details, and ultimately confirming or abandoning a transaction. At each step, seemingly minor sources of friction, such as latency, visual clutter, non-intuitive form layouts, sudden changes in pricing visibility, or unfamiliar identity verification requirements, can alter the probability that a customer progresses. These effects are cumulative: frictions interact across stages and channels, shaping path distributions rather than acting as isolated obstacles at individual pages. Yet in many operational settings, analysis remains anchored in local aggregates such as stepwise drop-off rates or A/B tests on individual screens, providing only partial insight into how interventions propagate through the full journey and how they translate into revenue and longer-term behavioral changes.

In practice, organizations face an abundance of tactical signals and constraints and a scarcity of synthesized structure [2]. Marketing teams may respond to perceived funnel underperformance by increasing upper-funnel spend or expanding retargeting, while product teams iterate on interface elements, and engineering teams focus on performance optimizations. Each intervention is often evaluated within its own experimental or observational frame, using metrics that are locally meaningful but not consistently linked to a shared model of customer behavior. As a result, friction-reduction efforts can conflict or saturate, interventions can compete for the same limited attentional bandwidth, and budget allocation decisions become dependent on short-run tests that do not generalize well across segments, seasons, or channels. Without a unifying modeling framework that expresses how frictions are measured, how they influence transitions through the funnel, and how interventions modify these relationships under uncertainty, it is difficult to attribute observed revenue changes to specific actions or to coordinate policies across teams.

The notion of friction in digital journeys is conceptually straightforward but operationally subtle [3]. Friction is not limited to long page loads or excess clicks; it includes ambiguous copy, inconsistent navigation hierarchies, unexpected fees revealed late in the process, rigid payment options, intrusive consent banners, repeated login prompts, and verification steps that lack clear justification. These elements impose cognitive or mechanical effort, and their perceived cost depends on context, intent, and individual tolerance. Some customers navigate dense flows with minimal difficulty, while others are sensitive to relatively small obstacles. Furthermore, interventions introduced to mitigate friction, such as reassuring messages, promotional overlays, or dynamic assistance widgets, can contribute their own forms of friction if they interrupt task flow or diminish clarity. Consequently, friction-reduction cannot be defined simply as the minimization of all effort or the proliferation of supportive elements; it must be understood as a calibrated adjustment of costs and signals within a structured behavioral system. [4]

This paper approaches the B2C purchase funnel as a controlled stochastic process in which observable states represent coarse-grained positions in the journey and transitions between states are influenced by frictions, contextual covariates, and explicit interventions. The modeling perspective is deliberately neutral: it does not assume that every intervention is beneficial or that every reduction in measured friction yields net value. Instead, it aims to quantify, in a disciplined manner, how changes in operational or interface parameters affect progression probabilities, abandonment patterns, and revenue outcomes, while propagating uncertainty from measurement through to decision-making. The process view enables the decomposition of performance into components attributable to structure (such as mandatory steps), to environmental conditions (such as device capabilities), and to controllable levers (such as layout simplifications or messaging strategies). By embedding frictions as parameterized quantities within this process, the framework supports comparison of alternative interventions on a common scale.

A central challenge is that funnel outcomes are sparse at the tail and heavily confounded [5]. Only a fraction of sessions convert; many users interact across multiple devices and channels; exogenous influences such as promotions, media coverage, holidays, and macroeconomic shifts modulate intent independently of on-site friction. Standard aggregate metrics are often sensitive to these background variations, making it difficult to disentangle the impact of a specific friction change or intervention. Additionally, experiments tend to be localized in scope and time, leading to context-specific results that are not directly transferable across segments or future periods. The framework proposed here addresses these issues by integrating event-level data architecture, stochastic

Stage	Friction Source	Impact	Metric Proxy
Exposure	Latency, Visual Clutter	Reduced attention	Page load time, bounce rate
Consideration	Cognitive effort, Info overload	Slower decision	Scroll depth, hover duration
Checkout	Form complexity, Price ambiguity	Abandonment	Field count, cart exits
Confirmation	Identity verification, UX shifts	Drop-off risk	Retry count, time-to-complete

Table 1: Representative funnel stages and friction proxies.

Friction Type	Origin	Example	Intervention Leverage
Structural	Regulatory/process	Two-factor auth	Clarify timing/info
Avoidable	Design/implementation	Redundant input	UI simplification
Contextual	Device or network	Slow mobile render	Adaptive layout
Behavioral	User cognition	Confusing copy	Message refinement

Table 2: Taxonomy of digital frictions and potential interventions.

state modeling, and causal inference techniques that explicitly account for selection mechanisms and non-stationarity. It emphasizes parameterizations that are interpretable, estimable from realistic data, and stable enough to support constrained optimization without relying on fragile assumptions. [6]

Within this framework, the construction of friction metrics is a foundational step. Rather than treating friction as an abstract label, the approach defines it through measurable proxies derived from telemetry and interaction patterns. Latency is measured via end-to-end timing statistics; form effort is measured by counts of required fields and observed input corrections; informational frictions are inferred from navigations to ancillary pages, repeated inspection of key text segments, or oscillations between cart and product pages. These proxies are associated with states and transitions, creating a consistent representation that can be incorporated into hazard functions and utility-like formulations. The goal is not to claim that any single proxy perfectly captures experienced friction, but to create a set of observable quantities that can be statistically linked to continuation behavior in ways that remain robust under policy changes. [7]

The paper further recognizes that interventions are themselves decisions within the same system. Common practices such as automatically triggering discount offers, displaying persistent urgency cues, or repeatedly prompting account creation are recast as actions with measurable costs and uncertain effects on both friction and conversion. Treating interventions as explicit actions, rather than as latent conditions, enables their integration into causal estimation and policy learning. It also clarifies trade-offs: an intervention can reduce a particular friction proxy for a subset of users while increasing cognitive load or perceived pressure for others, yielding heterogeneous effects that require careful estimation rather than heuristic interpretation. The framework is structured to incorporate randomized experiments where available and to augment them with doubly robust estimators in observational regions where experimentation is limited or infeasible. [8]

Decision-making under this framework is cast as a constrained optimization problem. Organizations operate under finite incentive budgets, design constraints, regulatory requirements, and reputational considerations that limit permissible action sets. The control layer therefore must maximize expected net revenue subject to constraints on cost, exposure density, and risk measures that bound underperformance relative to established baselines. Importantly, optimization takes place under parameter and model uncertainty: transition intensities, uplift functions, and friction-response elasticities are not treated as fixed values but as estimates with associated variability. The policy design incorporates this uncertainty by discounting high-variance gains, enforcing conservative thresholds, and maintaining guardrails that prevent the system from making large shifts based on weak or unstable evidence. This orientation is intended to reduce oscillatory behavior and to preserve interpretability for stakeholders who must review and approve policy changes. [9]

The introduction of a process-based, friction-aware, and intervention-explicit framework also implies a shift in how organizations interpret their existing metrics and experiments. Many current practices evaluate changes via isolated experiments using single primary metrics such as checkout completion rate or revenue per session. While these experiments remain valuable, their results can be more fully understood when embedded within a structured journey model that

Parameter	Meaning	Observed Proxy	Usage in Model
Latency	Rendering delay	Page timing logs	State transition hazard
Form Effort	Input difficulty	Field edits	Continuation probability
Info Friction	Search complexity	Repeated navigation	Utility modifier
Trust Cue	Security perception	Click-through on policy	Risk adjustment term

Table 3: Key measurable proxies for friction modeling.

Channel	Typical Frictions	Metric Focus	Optimization Goal
Web	Load speed, input lag	Conversion rate	Reduce latency
App	Navigation depth	Session retention	Simplify flows
Email	Link latency, CTA clarity	Click-through rate	Clarify calls-to-action
Social	Ad clutter, mismatch	Engagement time	Tune targeting

Table 4: Cross-channel frictional profiles and optimization targets.

distinguishes where in the funnel the effect arises, how it interacts with other stages, and whether it primarily accelerates existing conversions or attracts additional ones. For example, a change that simplifies checkout and raises same-session conversion may also alter the composition of customers returning through other channels; a promotional overlay that appears to increase conversions for a variant may, when modeled structurally, be found to primarily shift demand timing. The proposed framework does not negate these experiments but instead provides a mechanism to interpret them in a way that aligns short-run findings with long-run, multi-stage behavior. [10]

Instead of advocating maximal optimization or assigning intrinsic desirability to specific techniques, it concentrates on offering a coherent methodology that can be configured to reflect different organizational preferences, regulatory environments, and risk tolerances. The framework can support conservative policies that restrict incentives and intrusive elements, as well as more exploratory configurations, provided constraints and evaluation criteria are clearly specified. Uncertainty quantification, explicit modeling of costs, and transparent mapping from assumptions to recommendations are emphasized as mechanisms to prevent overstatement of conclusions and to facilitate informed oversight.

2 | Conceptualizing Digital Friction in B2C Funnels

Digital friction within B2C purchase funnels is multifaceted and context dependent. From a modeling perspective, friction is more usefully defined as an incremental cost borne by the customer at a particular state that influences the probability of continuation relative to abandonment or deferral [11]. This cost can arise from delays in rendering, cognitive effort in understanding information, effort associated with data entry, perceived uncertainty about total cost or delivery terms, distrust induced by unexpected permissions, or intrusive prompts that interfere with task flow. Treating friction as a state- and context-specific variable allows it to be integrated into stochastic transition models and causal effect estimation.

A practical decomposition distinguishes structural friction, which is inherent to process and compliance constraints, from avoidable friction due to design or implementation. Structural friction includes mandatory disclosures, required regulatory fields, and payment authentication steps. Avoidable friction arises from redundant inputs, fragmented navigation, unstable layouts, and under-optimized infrastructure [12]. For modeling, both forms are measured similarly, but their amenability to intervention differs. Structural friction may be partially mitigated by clarity and timing of information, whereas avoidable friction may be addressed more directly through engineering changes or interface simplification. By mapping friction components to actions, the framework separates descriptive metrics from levers that can be influenced.

The conceptualization also accommodates heterogeneity in friction sensitivity. Different users experience the same operational environment differently depending on device, connection quality, prior familiarity, language proficiency, risk tolerance, and purchase motivation. Thus, friction metrics must not be interpreted as homogenous penalties but as inputs into functions that can generate varying effective costs across segments [13]. A latency of one second can be negligible for some users but contribute

Objective Component	Symbol	Definition	Constraint
Revenue	R	Total realized gain	Maximize
Cost	C	Channel/operation cost	Limit budget
Value	LTV	Discounted future revenue	Weight conservatively
Fairness	δ	Action disparity bound	Must hold $\leq \delta$

Table 5: Optimization components under policy calibration.

Policy Aspect	Description	Risk Control	Interpretation
Uncertainty discounting	Penalize variance	Avoid overreaction	Conservative bias
Regularization	Smooth responses	Prevent oscillation	Stable policy
Segment weighting	Prioritize high-LTV	Limit bias	Fair resource use
Sensitivity analysis	Meta-parameter testing	Detect fragility	Governance insight

Table 6: Neutral calibration mechanisms in decision policy.

meaningfully to abandonment for others in high-urgency or low-trust contexts. The model therefore includes interaction terms between friction proxies and user or context covariates, enabling more nuanced decision rules that avoid uniform interventions.

From the standpoint of journey analytics, friction is tightly linked to temporal structure. Many signals considered as friction, such as repeated validation errors or multiple revisits to the same form field, manifest as patterns in event timing and ordering. To preserve this information, event sequences are transformed into state trajectories that include dwell times and revisitation counts [14]. Rather than summarizing the funnel as a static sequence of aggregated steps, the framework interprets funnel behavior as evolution within a latent cost landscape, where local elevations correspond to friction zones. These zones can be identified via changes in hazard ratios or localized increases in abandonment probability conditioned on entering a certain state under specific friction values.

Conceptual neutrality implies that friction is neither inherently detrimental nor inherently beneficial. Some frictions serve screening functions or increase perceived reliability. For instance, two-factor authentication introduces mechanical effort but may enhance trust for certain segments by signaling security [15]. The framework expresses such trade-offs by allowing friction proxies to have positive or negative coefficients in utility or hazard functions, depending on empirical evidence. This avoids encoding assumptions that all reductions in effort are beneficial, and it supports segment-level interpretations where the same intervention may be helpful in one context and counterproductive in another.

Finally, conceptual clarity about friction underpins the design of interventions. Interventions targeted at friction reduction are categorized not as isolated experiments but as candidate modifications to the cost landscape. Each candidate has hypothesized effects on specific friction metrics and, through them, on continuation probabilities and revenue [16]. By embedding this structure, the framework supports more disciplined comparison of alternative interventions and guards against conflating general promotional noise with genuine decreases in friction-related costs.

3 | Cross-Channel Dynamics, Lifetime Value Integration, and Neutral Policy Calibration

Cross-channel dynamics complicate the design and evaluation of friction-reduction strategies because customers traverse web, app, email, search, social, and offline touchpoints in patterns that are only partially observable. A framework that focuses exclusively on a single digital funnel risks attributing both friction and uplift to proximate interface changes while ignoring substitution and reinforcement effects across channels. To address this, the journey representation can be extended to a multi-channel state space in which each state encodes both stage in the purchase process and dominant channel context, while preserving parsimony through aggregation. This extended representation allows the model to separate cases where friction reduction in one channel truly creates incremental value from cases where it primarily re-routes existing intent from other channels with different cost structures or constraints. In such a configuration, the objective function is defined over joint revenue and cost across channels, and policies are assessed not only for their direct funnel impact but also for their influence on channel mix, acquisition costs, and emerging pathways that may introduce new frictions elsewhere. [17]

In this extended setting, each episode is enriched with channel tags and, when possible, probabilistic linkages to previous and subsequent episodes associated with the same or similar identities. The challenge is that identity resolution remains imperfect and may vary in quality across channels and over time. Rather than forcing deterministic stitching, the framework incorporates linkage probabilities as weights in constructing cross-channel trajectories, thereby propagating uncertainty into transition and value estimates. When the model infers that a substantial portion of app conversions are preceded by web sessions with high friction at checkout, a friction reduction on the web checkout must be interpreted in light of whether it captures demand that would otherwise complete on the app or through assisted channels. This interpretation is operationalized through path-based decomposition of expected revenue under alternative policies, distinguishing incremental paths that would not exist without the intervention from reallocated paths that shift the location of completion while leaving total demand approximately unchanged. [18]

Formally, consider an augmented state variable $S_t = (q_t, z_t)$, where q_t denotes the funnel stage and z_t the channel label. Let R denote aggregate revenue across channels within a defined horizon and C denote channel-specific costs including media spend and support operations. The objective becomes the maximization of net value across channels, subject to constraints on channel capacity and acceptable shifts in mix. Within this setup, friction metrics depend on both stage and channel, capturing, for example, differences in load performance between mobile app and mobile web, or differences in required inputs between kiosk, desktop, and call center flows. Actions include not only intrachannel adjustments but also routing decisions, such as whether to surface app-download prompts, to encourage live support, or to present alternative payment methods that functionally move the customer to different operational stacks.

$$\max_{\pi} \ \mathbb{E}_{\pi}\left[R - C\right]$$ [19]

subject to capacity and exposure constraints that can themselves vary by channel. The expectation is taken over trajectories that reflect both stochastic behavior and identity linkage uncertainty. In this formulation, reallocation of transactions from a high-cost channel to a lower-cost digital funnel can be desirable even without increasing total conversions, whereas shifting from a low-friction, low-cost channel to a more complex environment may be neutral or negative unless offset by other strategic considerations.

Integrating lifetime value into the decision process extends the temporal scope of evaluation beyond single episodes. Friction-reduction interventions that modestly increase immediate conversion among high-churn segments may be less attractive than interventions that foster more stable, lower-friction experiences for segments with higher expected retention, even when short-run conversion effects are similar. To incorporate this, each identity or pseudo-identity is associated with an estimated lifetime value that depends on historical behavior, product mix, and macro context [20]. Because of incomplete tracking, the framework treats lifetime value as a latent variable inferred from partial histories and aggregates, and uses it to weight the marginal impact of interventions.

$$\mathrm{LTV}(x) = \mathbb{E}\left[\sum_{u=0}^{U} \delta^u R_u \mid X = x\right],$$

where R_u is revenue at future period u, δ is a discount factor, and X denotes features summarizing observed characteristics and behaviors. The uplift associated with an intervention at a friction point is then evaluated not only on immediate R but on its induced change in the distribution of future trajectories relevant for LTV. In practice, direct estimation of long-horizon causal effects is limited by data and non-stationarity, so the framework adopts a neutral approach: lifetime value models are used to prioritize segments and to bound plausible longer-term impacts, but short-term, more identifiable effects remain primary drivers of decisions [21]. This guards against overextending interpretations while still allowing policies to avoid concentrating aggressive measures on segments with high strategic relevance.

To preserve tractability, the joint optimization of immediate revenue and lifetime value can be formulated as a weighted objective with conservative weights attached to long-term components. Let $\Delta_0(s,a)$ represent the estimated short-run incremental net revenue of action a in state s, and let $\Delta_\infty(s,a)$ represent a bounded estimate of its contribution to future value, derived from projected retention or cross-sell probabilities. A combined gain metric is defined as

$$G_{\mathrm{LTV}}(s,a) = \Delta_0(s,a) + w \cdot \Delta_\infty(s,a),$$

where w is chosen cautiously to reflect uncertainty. The policy uses $G_{\mathrm{LTV}}(s,a)$ in place of short-run uplift alone when ranking actions, but the conservative weighting prevents minor, noisy signals about long-term value from dominating well-identified near-term effects. This integration enables, for example, differentiated friction-reduction for high-value cohorts in ways that are transparent and bounded, without implying precise knowledge of multi-year outcomes. [22]

Cross-channel optimization further interacts with lifetime value because some channels are more informative about future potential than others. Engagement with high-intent search queries, early adoption of rich features in an app, or responsive behavior to non-incentivized content may signal higher expected value and greater sensitivity to friction. The framework thus allows the mapping from friction changes to value changes to be segment-specific, conditional on cross-channel patterns. However, it maintains neutral interpretation by constructing conservative hierarchical models that avoid abrupt classification. For instance, rather than categorizing users as definitively high or low value, the model treats $\mathrm{LTV}(x)$ as a distribution and derives policies based on probability-weighted effects.

$$\mathbb{E}[G_{\mathrm{LTV}}(s,a)] = \int G_{\mathrm{LTV}}(s,a \mid v)\; p(v \mid X)\, dv, [23]$$

where v indexes possible lifetime value realizations and $p(v \mid X)$ is the inferred distribution given observed features. This expectation is then subjected to the same uncertainty discounts and constraints as other gain estimates. The resulting policy recommendations are thus calibrated both by within-episode uplift evidence and by probabilistic assessments of longer-term impact, without assuming exact foresight.

Policy calibration in this expanded environment involves reconciling multiple, sometimes competing objectives in a stable and interpretable manner. The neutral stance of the framework is reflected in its treatment of optimization outputs as contingent on explicitly modeled assumptions and as subject to governance thresholds rather than as definitive rules. Sensitivity analyses expose how recommendations change when varying the weights on lifetime value, the discounting of uncertainty, or the tolerance for shifts in channel mix [24]. If small changes in these meta-parameters cause large behavioral differences, this indicates fragility, and either simpler policies or stronger regularization are preferred. Calibration therefore includes choosing parameterizations that produce smooth, monotone responses across realistic ranges of inputs, enabling operators to anticipate how interventions will adapt under evolving conditions.

An important consideration in cross-channel and lifetime value calibration is the avoidance of unintended disparities across user groups. Because friction, identity resolution, and channel access differ systematically across populations, naive optimization can concentrate high-friction experiences or aggressive interventions on segments with weaker observability or lower measured value. Although the core framework is not prescriptive about normative choices, it is technically straightforward to incorporate fairness-aware constraints that bound disparities in friction metrics, intervention intensities, or outcome distributions across predefined groups [25]. These constraints are expressed as additional terms in the optimization, enforcing, for example, that the probability of encountering specific high-burden flows does not systematically exceed designated thresholds for any monitored segment.

$$\left| \Pr_{\pi}(A = a \mid G = g_1) - \Pr_{\pi}(A = a \mid G = g_2) \right| \leq \delta,$$

for specified action a, groups g_1, g_2, and tolerance δ. Incorporating such constraints modifies the feasible policy set but does not require altering the underlying structural or causal estimators [26]. Instead, it shifts the selection of actions among those supported by evidence, in a way that remains analyzable and auditable. This further reinforces the emphasis on neutrality: the framework provides mechanisms to align with explicit organizational principles without implicitly encoding unexamined preferences.

Finally, the expanded section highlights that cross-channel dynamics, lifetime value considerations, and neutrality in calibration are not separate layers but interdependent dimensions of the same modeling effort. Decisions about where and how to reduce friction in a digital funnel influence channel substitution patterns, data richness for future modeling, and perceptions that feed back into long-run engagement. By embedding these linkages into a coherent mathematical and operational structure, organizations can interpret observed changes in revenue and experience with greater clarity, adjust interventions with awareness of their multi-period implications, and maintain a disciplined separation between what is estimated from data and what is imposed as policy preference [27]. The framework thus accommodates more complex strategic questions while retaining an emphasis on measured, non-exaggerated inference and transparent control.

4 | Data Architecture and Event-Level Representation

An effective friction-aware framework depends on a data architecture capable of reconstructing individual journeys, characterizing states, and recording intervention exposures. The raw material typically consists of clickstream logs, application telemetry, transactional databases, and campaign delivery logs. These heterogeneous sources are integrated into an event-level representation where each record encodes a timestamped interaction along with user or device identifiers, page or view identifiers, contextual covariates, and any actions applied by the system, such as message displays or layout variants.

Episodes are defined to bundle temporally proximal interactions into coherent journeys. A standard construction defines an episode as a maximal sequence of events separated by no more than a fixed inactivity threshold [28]. Within each episode, events are mapped to a finite set of abstract funnel states, such as landing, exploration, product detail, cart, checkout steps, payment interaction, and order confirmation. This mapping reduces interface-specific variations into analytically tractable categories while preserving the structure necessary for modeling progression and abandonment. The mapping can be implemented through deterministic rules or learned classifiers; in either case, the categorizations should exhibit temporal consistency and robustness to layout experiments.

Friction metrics are then derived as deterministic functions of event sequences and technical telemetry. For latency, the metric may be defined as server-side render times or client-side time-to-interactive [29]. For input complexity, metrics may include the count of mandatory fields, average keystrokes, or rate of validation errors. For informational friction, proxies can be constructed from counts of content blocks viewed, changes in displayed price during the session, or navigations to policy pages. These variables are computed at the state level and associated with transitions so they can enter hazard functions. Where identity resolution is incomplete, journeys are labeled with identifiers accompanied by confidence scores to indicate the degree of cross-device linkage, and analyses can be stratified or weighted accordingly.

A compact notation is used to encode the episode data [30]. Let each episode be indexed by n, with a sequence of states $\{S_{n,t}\}$, covariates $\{X_{n,t}\}$, actions $\{A_{n,t}\}$, and tensioned friction metrics $\{F_{n,t}\}$, all measured at decision times. The realized revenue R_n is attributed to the episode in which an order occurs, including both gross merchandise value and any attributable costs or discounts. To maintain analytic stability, the architecture enforces a consistent definition of decision times, for example at state transitions or at fixed intervals within states, so that the policy mapping and hazard estimation are well defined.

$$\mathcal{D} = \{(S_{n,t}, X_{n,t}, F_{n,t}, A_{n,t}, R_n)\}_{n,t}$$

Within this structure, interventions are treated as explicit actions, not merely variant attributes. For instance, showing a particular banner is encoded as an action with a unique label, rather than as a hidden experimental factor. This explicitness is essential both for causal identification and for subsequent use in optimization, where the policy must reason over an action set. Treatment eligibility rules and exposure throttles are recorded so that the effective policy under which data were generated can be reconstructed [31]. Without this information, uplift estimates risk contamination by undocumented targeting and collision between concurrent campaigns.

Data quality considerations are central. Missing events, clock inconsistencies, duplicated logs, or partial instrumentation can bias estimates of friction effects and treatment impacts. The framework therefore incorporates diagnostic tests for basic conservation properties, such as alignment between order events and payment confirmations, coherence of state transitions, and plausible ranges for friction metrics. Episodes or periods failing these checks can be excluded or down-weighted. This conservatism reflects the preference for neutral, interpretable inferences over aggressive extrapolation from ambiguous records. [32]

5 | Stochastic Journey Modeling and Structural Components

Building on the event-level representation, the journey is modeled as a controlled stochastic process with discrete states, time-varying covariates, and interventions. A semi-Markov decision process formulation is suitable because it accommodates variable dwell times and allows transition intensities to depend on both state and elapsed time. The core structural component is a set of hazard functions that describe the instantaneous propensity to move from one state to another or to abandon, conditional on covariates, friction metrics, and actions.

Let S_t denote the funnel state at decision time t, drawn from a finite state space including at minimum a visit state, product exploration, cart, checkout, payment interaction, completion, and a transient abandonment or inactivity state. Let X_t be a vector of observed covariates and F_t a vector of friction proxies

Data Source	Content Type	Integration Role	Example Metrics
Clickstream Logs	Page events	Journey reconstruction	Dwell time, transitions
App Telemetry	Device interactions	State encoding	Load time, UI latency
Transaction DB	Orders/payments	Revenue attribution	Basket value, discounts
Campaign Logs	Message exposures	Action tracking	Impression count, click ID

Table 7: Primary data sources and their functional roles in event-level integration.

State	Typical Events	Friction Indicators	Outcome Link
Landing	Page load, session start	Latency spikes	Bounce probability
Exploration	Scrolls, product views	Info density	Transition to detail
Cart	Add/remove items	Price changes	Checkout entry
Checkout	Form fills, validation	Input errors	Payment initiation
Payment	Auth, confirmation	Timeout, retries	Purchase completion

Table 8: Abstract funnel states with representative event patterns and frictions.

Friction Metric	Computation Basis	Interpretation	Usage
Latency	Time-to-interactive	Mechanical delay	Hazard scaling
Input Complexity	Required inputs	Cognitive effort	Continuation probability
Info Friction	Page revisits	Decision ambiguity	Abandonment hazard
Trust Proxy	Policy views	Perceived reliability	Risk adjustment

Table 9: Derived friction metrics and modeling interpretations.

Model Component	Symbol	Functional Role	Dependencies
State	S_t	Funnel position	Discrete process
Covariates	X_t	Contextual features	Device, channel
Friction vector	F_t	Effort representation	Derived proxies
Action	A_t	Intervention variable	Policy control

Table 10: Core components of the stochastic journey model.

Causal Element	Notation	Purpose	Estimation Approach
Propensity Score	$e(X)$	Treatment assignment model	Logistic / ML estimator
Outcome Model	$\mu_a(X)$	Expected outcome by action	Regression / neural net
Doubly Robust Score	ψ	Bias-corrected estimator	Combines both models
Uplift Function	$\tau(x)$	Heterogeneous effect surface	Segment-level inference

Table 11: Causal inference components and associated estimators.

Validation Aspect	Diagnostic Test	Purpose	Mitigation
Temporal Coherence	Event ordering check	Ensure time consistency	Drop/reorder events
Transition Logic	State adjacency validation	Preserve structural rules	Flag anomalies
Metric Range	Plausibility bounds	Detect outliers	Winsorize/trim
Policy Traceability	Exposure reconstruction	Avoid hidden bias	Record treatment IDs

Table 12: Data quality and validation safeguards in the analytical pipeline.

at time t, and let A_t represent the chosen intervention. For each pair of states (i, j) that are reachable, a transition hazard is parameterized as

$$\lambda_{i\to j}(t) = \exp\{\beta_{ij}^\top X_t + \gamma_{ij}^\top F_t + \eta_{ij}^\top A_t\}.$$

Abandonment is represented by a dedicated state with its own hazard functions from each active state [33]. Because multiple hazards may be active simultaneously, the next event type is drawn from a competing risks structure determined by the relative magnitudes of hazards. The structure makes it possible to express how variations in friction levels, such as increased latency or additional form fields, systematically influence the probability of proceeding, redirecting, or abandoning.

To prevent overfitting and to incorporate behavioral considerations, the hazard functions are regularized and linked to a latent utility formulation. Define a continuation utility for state i at time t as

$$U_{i,t} = \theta_i^\top X_t - \phi_i^\top F_t + \psi_i^\top A_t + \epsilon_{i,t},$$

where $\epsilon_{i,t}$ is a noise term. Under standard assumptions, continuation probabilities can be linked to hazards via monotone transformations of $U_{i,t}$. This connection encourages consistency between interpreted friction costs and empirically estimated hazard impacts [34]. Regularization can be implemented through penalties on parameter magnitudes or hierarchical priors that shrink extreme values toward shared baselines.

Revenue enters through a terminal value function defined on the completion state and possibly intermediate states if partial monetization, such as deposits or ancillary services, occurs. For an episode, the realized revenue is modeled as

$$R = \sum_t Y_t \mathbb{I}(S_t = y),$$

where Y_t denotes the basket value or margin at time t upon entering the terminal state y. Basket value can itself be modeled as a function of covariates and actions, particularly when incentives or recommendations directly alter composition. In many applications, separating the probability of completion from the conditional distribution of value yields more stable estimates and cleaner decomposition of intervention effects. [35]

The structural model is estimated from historical data using maximum likelihood or Bayesian inference. Likelihood contributions arise from observed transitions and dwell times, with covariate and action dependence explicitly encoded. Where data volume supports it, interaction terms between friction metrics and key covariates are included, enabling the model to capture heterogeneous sensitivity. Inference outputs include point estimates and covariance structures or posterior distributions for parameters, which in turn generate uncertainty intervals for transition probabilities and expected revenues under candidate policies.

Importantly, the structural components are kept interpretable. Parameter groupings map to clear statements such as the incremental effect of an additional second of latency in checkout on abandonment hazard for a specific device class, or the change in completion hazard associated with a reduced field set in payment forms [36]. This interpretability is fundamental for subsequent governance, where stakeholders may request explanations for recommended interventions.

6 | Causal Inference, Heterogeneous Effects, and Uplift

While the structural model describes how covariates, friction, and actions jointly correlate with transitions, it does not by itself guarantee causal identification of intervention effects. To support policy decisions, the framework incorporates explicit causal inference components based on potential outcomes. For each candidate intervention a at a decision point characterized by history H_t, one can define potential revenue outcomes $R^{(a)}$ describing what would occur if action a were chosen. The quantity of interest is the conditional average treatment effect between actions, given observable covariates derived from H_t.

Because interventions in operational environments are rarely assigned uniformly at random, observational data reflect selection biases driven by targeting rules, eligibility filters, and user self-selection. The framework therefore estimates both the propensity of receiving an action and the outcome model, combining them into doubly robust estimators [37]. Let $A \in \{0, 1\}$ indicate whether a given binary intervention is applied at a particular decision point, and let Y represent a suitable outcome, such as episode revenue or a transformed metric aligned with model assumptions. The propensity score $e(X)$ is the conditional probability of treatment given covariates X, approximated through flexible models that incorporate known eligibility constraints.

$$e(X) = \Pr(A = 1 \mid X),$$

and the corresponding doubly robust score for the individual contribution is defined as

$$\psi = \hat{\mu}_1(X) - \hat{\mu}_0(X) + \frac{A(Y - \hat{\mu}_1(X))}{\hat{e}(X)} - \frac{(1 - A)(Y - \hat{\mu}_0(X))}{1 - \hat{e}(X)},$$

where $\hat{\mu}_a(X)$ are estimated outcome regressions for action a. Averaging ψ over groups of observations with similar covariates yields estimates of heterogeneous treatment effects [38]. These estimates can be parameterized as uplift functions $\tau(x)$ that inform which segments are likely to benefit from an intervention.

In multi-action settings, the procedure generalizes. Propensity scores become a vector over actions, and outcome regressions predict potential outcomes under each action. Simplified one-vs-rest approaches can be adopted for scalability, provided policies account for relative rankings. The key requirement is that all actions considered in the policy have been sufficiently explored in historical or experimental data to permit identification within the relevant covariate regions. Where exploration is limited, the framework incorporates uncertainty inflation or explicit constraints that limit policy divergence from historical behavior. [39]

To capture the structure of the funnel, treatment effects may be defined on intermediate quantities such as state-specific hazard modifications rather than only on terminal revenue. For example, the effect of an intervention displayed at the cart stage may be measured primarily on the probability of advancing to checkout within a window, with secondary effects on final revenue. In such cases, the causal component is estimated conditional on entering the relevant state, and potential outcomes are integrated with the structural hazard model to project overall revenue implications.

Given the importance of stability, hierarchical models are used to share statistical strength across segments while permitting heterogeneity. Let τ_g denote the uplift in segment g. One can specify [40]

$$\tau_g \sim \mathcal{N}(\mu_\tau, \sigma_\tau^2),$$

with hyperparameters estimated from data. This structure shrinks extreme estimates toward the global mean and reduces sensitivity to noise in small segments. The resulting posteriors provide both expected uplift and credible intervals, which the policy will later use to make risk-aware decisions.

By combining the structural journey model with causal uplift estimation, the framework distinguishes between correlation induced by friction and other covariates and the incremental impacts of specific interventions that modify friction or its perception. This distinction helps avoid misattributing changes in performance to interventions when they are actually driven by shifts in underlying traffic mix or exogenous factors. [41]

7 | Optimization, Policy Learning, and Control

With structural and causal components in place, attention turns to constructing policies that select interventions along the journey to optimize revenue-related objectives under constraints. The purchase funnel is represented as a Markov decision process with state space including funnel position, relevant covariates, friction metrics, and latent variables summarizing exposure history. Actions consist of allowable interventions, including the option of inaction. Rewards incorporate immediate revenues, expected downstream contributions, and action costs, such as incentive expenses or user experience impacts. The policy optimization problem seeks a mapping $\pi(a \mid s)$ from states to action probabilities that maximizes expected cumulative reward while satisfying resource constraints [42]. Let R denote episode-level net revenue and $C(a)$ denote the cost associated with action a. Let B be a budget on expected total cost over a horizon, and consider a risk constraint on downside deviation from a baseline policy. The optimization can be expressed as

$$\max_{\pi} \ \mathbb{E}_\pi[R]$$

subject to

$$\mathbb{E}_\pi[C(A)] \leq B,$$

and a constraint limiting tail risk in relative performance. To make the problem tractable, a Lagrangian formulation introduces multipliers for the constraints and converts them into penalties within the objective [43]. Approximate dynamic programming or policy gradient methods are then applied, using the estimated transition dynamics and uplift functions as inputs.

The learned policy incorporates uncertainty by adjusting estimated gains downward according to their variability. For a given state s and action a, let $\Delta(s, a)$ be the estimated net uplift in expected reward derived from causal models and structural propagation, and let $\sigma(s, a)$ be a corresponding uncertainty scale. A conservative adjusted gain might be

$$G(s, a) = [44]\Delta(s, a) - z \cdot \sigma(s, a),$$

where z is a nonnegative parameter controlling aversion to uncertainty. The policy can then be represented as a softmax distribution over adjusted gains minus cost and margin terms, constrained by throttling rules that prevent excessive exposures.

$$\pi(a \mid s) = \frac{\exp\{G(s, a)/\tau\}}{\sum_b \exp\{G(s, b)/\tau\}},$$

where τ is a temperature parameter [45]. When resource constraints bind, the associated multipliers effectively reduce gains for high-cost actions and tilt the policy toward lower-cost alternatives. This construction is compatible with both offline training and online adjustment.

In practice, the optimization need not rely solely on high-dimensional reinforcement learning. Given the interpretability requirements and the need for governance, simpler myopic or two-step policies can be sufficient. For example, at each decision point the policy might select the feasible action with the highest conservative uplift among those passing eligibility checks and exposure caps [46]. Structural modeling provides a way to map local effects to expected episode-level outcomes, so even a greedy policy can be evaluated in terms of its global consequences.

A critical property of the control layer is monotonicity with respect to constraints. As budgets tighten or risk aversion increases, the policy should adjust smoothly rather than exhibiting abrupt changes that are difficult to interpret or implement. Sensitivity to hyperparameters is therefore examined through simulation and diagnostic tools described in subsequent sections. The intention is to maintain a neutral and stable response surface that facilitates organizational acceptance and avoids inadvertent shocks to user experience caused by model or parameter fluctuations.

8 | Simulation, Diagnostics, and Governance

Simulation plays a central role in understanding how the modeling and policy stack behaves under varying conditions [47]. Because real-world environments contain confounding, nonstationarities, and constraints on experimentation, a calibrated simulator can provide a complementary view of plausible outcomes. The simulator is constructed by instantiating the structural hazard and value models with parameter values estimated from data and then generating synthetic journeys under alternative policies and friction configurations.

Let x_t denote covariate vectors drawn from an empirical or parametric distribution, and let hazards $\lambda_{i\to j}(t)$ be computed using previously estimated parameters. Actions at each decision point are drawn according to a candidate policy π. The simulator then advances the state until completion or absorption, accruing costs and rewards as specified. By repeating this process over many synthetic episodes, one obtains distributions of revenue, cost, and auxiliary metrics such as number of prompts, incentive utilization, and abandonment rates [48]. This environment allows evaluation of how different policies and friction reductions influence both central tendencies and tails.

Simulation is also used to study misspecification sensitivity. For instance, one can perturb friction effect parameters within plausible intervals and observe how optimal or learned policies respond. If small perturbations cause large changes in recommended actions, the system may be judged unstable and in need of further regularization. Similarly, one can introduce exogenous shocks to covariate distributions, such as shifts in device mix or traffic sources, to examine robustness [49]. Diagnostics track metrics such as regret relative to an oracle policy that has access to true parameters, frequency of constraint violations, and variability of outcomes over repeated runs.

$$\mathrm{Regret} = \mathbb{E}[R^\star] - \mathbb{E}_\pi[R],$$

where $R^\star$ is the reward under the oracle policy. While oracle performance is not available in production, simulator-based approximations support internal consistency checks and comparisons among candidate policies. If a simpler interpretable policy exhibits low regret in plausible environments, it may be preferred over a more complex alternative.

Governance overlays these technical elements with processes that ensure transparency, accountability, and adherence to organizational and regulatory constraints. Policy recommendations, uplift estimates, and friction metrics should be inspectable [50]. For any action deployed at scale, stakeholders should be able to see the associated estimated effects, uncertainty intervals, and how the action interacts with budgets and risk limits. Logs must record which policy version was active at each time, which actions were available, and which constraints were binding.

To avoid inadvertent escalation toward aggressive tactics, governance may impose meta-constraints such as maximum impression frequencies for disruptive elements, upper bounds on incentives per customer over defined periods, and prohibitions on conditioning friction on sensitive attributes. The modeling framework accommodates these by restricting the action set or by incorporating hard constraints into optimization. Observed divergences between predicted and realized metrics trigger diagnostic investigations, including re-examination of instrumentation, drift in user behavior, or unexpected interactions among concurrent features.

The overall role of simulation and diagnostics is not to guarantee precise predictions but to provide a structured environment in which assumptions are made explicit, sensitivities are measured, and governance bodies can evaluate trade-offs with a clear line of sight. [51]

9 | Implementation Considerations and Organizational Integration

Implementing the proposed framework requires both technical infrastructure and organizational alignment. On the technical side, data pipelines must support reliable and timely ingestion of event logs, consistent mapping to funnel states, computation of friction metrics, and assembly of modeling datasets. Feature

Optimization Element	Symbol	Role	Interpretation
Policy	$\pi(a \mid s)$	Action distribution	Mapping from states to interventions
Reward	R	Objective signal	Net revenue outcome
Cost	$C(a)$	Resource expenditure	Incentive or UX impact
Budget	B	Constraint	Expected total cost cap

Table 13: Key variables and constraints in policy optimization.

Policy Mechanism	Equation	Purpose	Control Parameter
Adjusted Gain	$G(s,a) = \Delta(s,a) - z\sigma(s,a)$	Penalize uncertain rewards	z (risk aversion)
Softmax Policy	$\pi(a\|s) \propto \exp\{G(s,a)/\tau\}$	Smooth action selection	τ (temperature)
Constraint Penalty	$L = \mathbb{E}[R] - \lambda(\mathbb{E}[C] - B)$	Enforce cost limits	λ (multiplier)
Exposure Throttle	$\pi(a\|s) \leq \kappa$	Cap action frequency	κ (max exposure)

Table 14: Policy control constructs and associated tuning parameters.

Simulation Component	Input Source	Generated Output	Use Case
Structural Model	Estimated hazards	Synthetic transitions	Journey reconstruction
Policy Function	Learned $\pi(a\|s)$	Action sequences	Comparative evaluation
Covariate Stream	Empirical data	User/device mix	Scenario realism
Revenue Function	Value model	Simulated earnings	Regret and variance checks

Table 15: Simulation architecture and its analytical objectives.

Diagnostic Metric	Definition	Insight Provided	Governance Use
Regret	$\mathbb{E}[R^\star] - \mathbb{E}_\pi[R]$	Efficiency loss	Policy comparison
Constraint Violations	$\Pr(C > B)$	Budget/risk breaches	Enforcement checks
Stability Sensitivity	$\partial\pi/\partial\theta$	Hyperparameter robustness	Regularization tuning
Outcome Variance	$\mathrm{Var}(R)$	Predictive reliability	Policy smoothness

Table 16: Diagnostics for simulation-based validation and oversight.

Implementation Layer	Function	Core Requirement	Example Component
Data Pipeline	Event ingestion	Consistent mapping	Stream processors
Model Training	Hazard/uplift estimation	Versioned features	ML workflows
Policy Serving	Real-time decisioning	Fallback safety	Microservice APIs
Monitoring	Performance tracking	Drift detection	Dashboards, alerts

Table 17: Technical implementation layers supporting the framework.

Governance Aspect	Objective	Mechanism	Example Policy
Transparency	Explain decisions	Log + report active policies	Version tagging
Accountability	Track outcomes	Audit intervention effects	Uncertainty bounds
Fairness	Bound sensitive exposure	Hard constraints in $\pi(a\|s)$	Frequency limits
Adaptability	Handle drift	Periodic retraining	Rolling window updates

Table 18: Governance and oversight mechanisms for sustainable deployment.

representations for covariates, friction, and interventions must be versioned to ensure that model training and policy execution use compatible definitions. The serving architecture should support deployment of policies in real time or near real time, with fallbacks to default behaviors in case of degradation.

Model training workflows are organized around regular cycles that reflect business rhythms and data freshness [52]. Structural hazard models and value estimators can be retrained on rolling windows to capture evolving relationships while smoothing transients.

Causal uplift models can incorporate both observational and experimental data, systematically including new experiments as they conclude. Policy optimization can be run on updated models to generate candidate policies whose performance is compared against baselines through off-policy evaluation and, where feasible, controlled rollouts.
Integration with existing experimentation platforms is important. Interventions recommended by the policy framework may coincide with treatments in concurrent experiments [53]. To maintain identifiability, assignment logs and eligibility criteria must be unified. The framework can consume experiment metadata to treat randomized variants as exogenous sources of variation when estimating causal effects. Conversely, when deploying new policies, organizations can choose to maintain randomized exploration at controlled rates to sustain learning about underrepresented segments and actions.
From an organizational perspective, adoption requires clarity regarding decision rights and responsibilities. Data science, product, engineering, and marketing teams must share an understanding of friction metrics, intervention definitions, and objective functions [54]. Governance committees or designated owners review proposed changes in action sets, constraint configurations, and risk thresholds. Model documentation should describe key assumptions and limitations in non-promotional language, emphasizing where empirical support is strong and where it remains limited.
Monitoring is implemented through dashboards aligned with the framework. Core panels track conversion, revenue, average incentive cost, friction metrics at key states, and distributions of exposure to interventions, all segmented by relevant attributes such as device type or traffic source. Additional panels track divergence between model predictions and observed outcomes, providing early indications of drift. When discrepancies persist beyond thresholds, retraining or recalibration is initiated, and new simulations are run to reassess stability. [55]
The framework is designed to coexist with human judgment rather than displace it. Policy outputs can be viewed as recommendations augmented with uncertainty and constraint information. Stakeholders may override specific recommendations or adjust constraints in response to qualitative insights, risk appetite, or strategic considerations. Because the system is explicit about how such changes affect objective functions and constraints, these adjustments remain auditable and can be folded back into future modeling cycles.
Finally, implementation should acknowledge non-stationarity in consumer expectations and technology [56]. As new payment methods, devices, or regulatory obligations emerge, friction profiles change. The modular nature of the framework simplifies adaptation: new states can be introduced, new friction metrics defined, and new actions considered, while preserving continuity in foundational components such as hazard estimation and policy optimization.

10 | Conclusion

This paper has presented a neutral and integrated framework for reducing friction in the B2C digital purchase funnel through data-driven customer journey modeling and intervention design. By conceptualizing friction as state- and context-specific costs embedded in a controlled stochastic process, and by aligning event-level data architecture with this representation, the framework supports explicit modeling of progression, abandonment, and revenue formation. Structural hazard models and behavioral regularization provide interpretable links between friction metrics and funnel transitions [57]. Causal inference components estimate heterogeneous treatment effects for candidate interventions under realistic assignment mechanisms, while constrained optimization translates these estimates into policies that respect budgets, risk tolerances, and exposure constraints.
Simulation and diagnostic tools offer a controlled environment to evaluate policy behavior under uncertainty and potential misspecification, enabling informed trade-offs between complexity, robustness, and interpretability. Implementation considerations emphasize reliable instrumentation, consistent feature definitions, integration with experimentation systems, and governance processes that ensure transparency and accountability. The framework does not assume that all frictions must be eliminated or that interventions yield uniform benefits. Instead, it provides a structured method to quantify plausible incremental effects, propagate them through journey dynamics, and support measured adjustments as conditions evolve. In doing so, it aims to facilitate coherent alignment between design, engineering, and marketing actions and empirically grounded expectations of revenue and experience outcomes in digital B2C environments. [58]

References

[1] T. V. Rocha, C. L. S. Jansen, E. Lofti, and R. R. Fraga, "An exploratory study on the use of social

networks in building customer relationships," *Review of Business Management*, vol. 15, pp. 262–282, 6 2013.

[2] T. Miesler, C. Wimschneider, A. Brem, and L. Meinel, "Frugal innovation for point-of-care diagnostics controlling outbreaks and epidemics.," *ACS biomaterials science & engineering*, vol. 6, pp. 2709–2725, 4 2020.

[3] S. K. Mufti and A. Aprianingsih, "Customer analysis using business buying behavior framework to develop a b2b marketing strategy for user experience (ux) studio, uxgala," *European Journal of Business and Management Research*, vol. 7, pp. 73–77, 1 2022.

[4] A. Elsafty and M. Elshahed, "The changes in online buying intention as a determinant of behavior during covid-19 pandemic in the ready-made garments industry in egypt," *International Journal of Business and Management*, vol. 16, pp. 1–, 4 2021.

[5] J. Song and J. Baker, "An integrated model exploring sellers' strategies in ebay auctions," *Electronic Commerce Research*, vol. 7, pp. 165–187, 5 2007.

[6] B. Jin and D. C. Shin, "The power of 4th industrial revolution in the fashion industry: what, why, and how has the industry changed?," *Fashion and Textiles*, vol. 8, pp. 1–25, 10 2021.

[7] X. Cai, M. Milojevic, D. Syromyatnikov, A. Kurilova, and B. Ślusarczyk, "Mathematical interpretation of global competition between payment systems," *Mathematics*, vol. 9, pp. 2070–, 8 2021.

[8] N. Ind and M. C. Riondino, "Branding on the web: A real revolution?," *Journal of Brand Management*, vol. 9, pp. 8–19, 9 2001.

[9] M. Mahdzar, A. Khairuddin, W. A. W. A. Zahri, N. M. Fakri, R. A. Rahman, and M. S. H. Saifudin, "The relationship between traditional travel agencies website design on customers' e-loyalty," *Journal of International Business, Economics and Entrepreneurship*, 12 2022.

[10] E. G. Sözer, M. E. Civelek, and M. Çemberci, "Strategic excellence in post-digital ecosystems: A b2c perspective," *Zea Books*, 1 2018.

[11] A. R. R. S. AL-Hawari, A. P. Balasa, and Z. Slimi, "Covid-19 impact on online purchasing behaviour in oman and the future of online groceries," *European Journal of Business and Management Research*, vol. 6, pp. 74–83, 7 2021.

[12] H. Hou, M. Hu, L. Chen, and J. Y. Choi, "An enhanced model framework of personalized material flow services," *Information Technology and Management*, vol. 12, pp. 149–159, 5 2011.

[13] A. Parshukov, I. Aleksandrov, and M. Fedorova, "Universal classification of goods and services for marketing and logistics," *E3S Web of Conferences*, vol. 217, pp. 07001–, 12 2020.

[14] T. Bosona, "Urban freight last mile logistics-challenges and opportunities to improve sustainability: A literature review," *Sustainability*, vol. 12, pp. 8769–, 10 2020.

[15] A. Hazarika, "E-commerce: Expectations and shortfalls," *Christ University Law Journal*, vol. 2, pp. 59–78, 12 2013.

[16] null Nuryakin, I. Fatmawati, and K. Siriyota, "Why does environmental uncertainty have a moderating effect on marketing performance? a comparative analysis of indonesian and thai service industries," *Engineering Management in Production and Services*, vol. 14, pp. 57–67, 9 2022.

[17] Y. Popova and O. Cernisevs, "Smart city: Sharing of financial services," *Social Sciences*, vol. 12, pp. 8–8, 12 2022.

[18] Z. Huo and J. Du, "Analysis of e-commerce business strategy - take jd.com, inc. as the case," *Asian Business Research*, vol. 7, pp. 41–41, 1 2022.

[19] 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.

[20] X.-X. Yan, Z.-Q. Hu, J. Xu, and J.-Y. Liu, "Research on the social e-commerce marketing model based on sicas model in china," *International Journal of Marketing Studies*, vol. 9, pp. 113–123, 5 2017.

[21] H. Yadav, U. Soni, S. Gupta, and G. Kumar, "Evaluation of barriers in the adoption of e-commerce technology in smes," *Journal of Electronic Commerce in Organizations*, vol. 20, pp. 1–18, 12 2021.

[22] Y.-Y. Chen, Y.-S. Duan, J.-J. Ni, and J.-L. Liang, "Perceived risk in tuanzhu group buying and traditional online buying in taiwan," *International Business Research*, vol. 9, pp. 92–102, 11 2016.

[23] S. Saniuk and S. Grabowska, "The concept of cyber-physical networks of small and medium enterprises under personalized manufacturing," *Energies*, vol. 14, pp. 5273–, 8 2021.

[24] J. Wang, "Research on the impact of customer participation in virtual community on service innovation performance- the role of knowledge transfer.," *Frontiers in psychology*, vol. 13, pp. 847713–, 3 2022.

[25] I. Onyusheva and J. Thongaim, "Customer satisfaction in online shopping: A case study of bangkok urban area," *The EUrASEANs: journal on global socio-economic dynamics*, pp. 13–23, 3 2018.

[26] D. Zavyalov, N. Zavyalova, O. Saginova, and N. Kireeva, "Digitalisation processes and their impact on the development of commodity trading companies," *SHS Web of Conferences*, vol. 93, pp. 01016–, 1 2021.

[27] F. A. de Araújo, J. G. M. dos Reis, M. T. da Silva, and E. Aktas, "A fuzzy analytic hierarchy process model to evaluate logistics service expectations and delivery methods in last-mile delivery in brazil," *Sustainability*, vol. 14, pp. 5753–5753, 5 2022.

[28] E. IJBM, "International journal of business and management, vol. 4, no. 10, october 2009, all in one pdf file," *International Journal of Business and Management*, vol. 4, 9 2009.

[29] R. Mengen, M. S. Ramli, N. H. Ayob, and Z. N. Roslan, "The influence of e-service quality toward online shopping behaviour amongst uitm students in malaysia," *International Journal of Academic Research in Business and Social Sciences*, vol. 11, pp. 175–189, 10 2021.

[30] A. Mor and M. G. Speranza, "Vehicle routing problems over time: a survey," *4OR*, vol. 18, pp. 129–149, 3 2020.

[31] A. Rady, M. E. Rady, and null Ezzat, "Exploring the usage of digital content marketing in hotels and travel agencies in egypt," , vol. 15, pp. 56–72, 10 2018.

[32] P. Handa, M. Mahapatra, and A. Yadav, "Brand communication through facebook altering consumer psychology," *Journal of Content, Community and Communication*, vol. 4, pp. 34–42, 6 2018.

[33] J. N. Nakhumwa, "Adoption of e-commerce payment systems by commercial banks in kenya," *INTERNATIONAL JOURNAL OF MANAGEMENT & INFORMATION TECHNOLOGY*, vol. 9, pp. 1600–1622, 4 2014.

[34] D. Dimyati and J. Juvenia, "Facebook usage towards reading interest at news portal merdeka.com in north cikarang," *Jurnal ASPIKOM*, vol. 3, pp. 550–569, 9 2017.

[35] D. Juniansyah, A. H. P. K. Putra, H. Syahnur, S. Hasan, and S. Nujum, "Symmetrical and asymmetrical of tam: Consumer emotional value and service innovation on consumer purchase decisions," *Golden Ratio of Mapping Idea and Literature Format*, vol. 2, pp. 8–35, 1 2021.

[36] O. Vynogradova and N. Drokina, "The structure of an integrated internet marketing complex, based on the marketing-mix concept," *Acta Scientiarum Polonorum. Oeconomia*, vol. 19, pp. 117–126, 9 2020.

[37] 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.

[38] W. Hong, C. Zheng, L. Wu, and X. Pu, "Analyzing the relationship between consumer satisfaction and fresh e-commerce logistics service using text mining techniques," *Sustainability*, vol. 11, pp. 3570–, 6 2019.

[39] V. Fedorenko, N. Filipenko, I. Shumilo, V. Nesterovych, and S. Nischymna, "Entrepreneurial activity of the it sector in the conditions of the covid-19 pandemic and in the post-quarantine period," *Entrepreneurship and Sustainability Issues*, vol. 8, pp. 697–712, 6 2021.

[40] L. Wang, H. Tang, D. Liu, and C. Xing, "Study on customer experience dimension based on online shopping reviews of mi's products," *MATEC Web of Conferences*, vol. 100, pp. 05024–, 3 2017.

[41] T. N. Varma and D. A. Khan, "Information technology and e-risk of supply chain

management," *African Journal of Business Management*, vol. 9, pp. 243–258, 3 2015.

[42] D. Dumitriu, G. Militaru, D. C. Deselnicu, A. Niculescu, and M. Popescu, "A perspective over modern smes: Managing brand equity, growth and sustainability through digital marketing tools and techniques," *Sustainability*, vol. 11, pp. 2111–, 4 2019.

[43] D. Raheem, M. G. Shishaev, and V. Dikovitsky, "Food system digitalization as a means to promote food and nutrition security in the barents region," *Agriculture*, vol. 9, pp. 168–, 8 2019.

[44] null Hartoyo, F. Karambut, R. Nurmalina, and M. Najib, "The intention in online submission of micro credit," *EUROPEAN RESEARCH STUDIES JOURNAL*, vol. XXII, pp. 186–200, 8 2019.

[45] S. G. Gavrila and A. de Lucas Ancillo, "Spanish smes' digitalization enablers: E-receipt applications to the offline retail market.," *Technological forecasting and social change*, vol. 162, pp. 120381–120381, 10 2020.

[46] I. Hidayat and S. Agustina, "The effect of decreased transaction volume on the low increase in e-commurce revenue during the pandemic on the shope application in 2020," *Jurnal Penelitian Ekonomi Dan Akuntansi*, vol. 7, pp. 103–111, 2 2022.

[47] S. Vinerean and A. Opreana, "Measuring customer engagement in social media marketing: A higher-order model," *Journal of Theoretical and Applied Electronic Commerce Research*, vol. 16, pp. 2633–2654, 10 2021.

[48] Y. Hu, S. Chen, and J. Wang, "Managerial humanistic attention and csr: Do firm characteristics matter?," *Sustainability*, vol. 10, pp. 4029–, 11 2018.

[49] R. A. Zaker, "The effect of social networking tools on the smes domain: an empirical study in the nineveh governorate-iraq," *TANMIYAT AL-RAFIDAIN*, vol. 39, pp. 9–20, 6 2020.

[50] 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.

[51] M. Matarazzo, L. Penco, and G. Profumo, "How is digital transformation changing business models and internationalisation in made in italy smes," *Sinergie Italian Journal of Management*, vol. 38, pp. 89–107, 1 2021.

[52] Z. Kalinic, V. Rankovic, and L. Kalinić, "Challenges in cross-border e-commerce in the european union," *Zeszyty Naukowe Uniwersytetu Ekonomicznego w Krakowie*, pp. 159–170, 2 2019.

[53] D. L. Tungkup, "The importance of online transportation effectiveness for business resistance strategies during the covid-19 pandemic," *KnE Social Sciences*, pp. 396–414–396–414, 1 2021.

[54] M. B. Khan and M. P. Martin, "Barriers to b2c segment of e-business," *Journal of Business & Economics Research (JBER)*, vol. 2, 2 2011.

[55] S. Periyasami and A. P. Periyasamy, "Metaverse as future promising platform business model: Case study on fashion value chain," *Businesses*, vol. 2, pp. 527–545, 11 2022.

[56] K. Kraus, N. Kraus, M. Hryhorkiv, I. Kuzmuk, and O. Shtepa, "Artificial intelligence in established of industry 4.0," *WSEAS TRANSACTIONS ON BUSINESS AND ECONOMICS*, vol. 19, pp. 1884–1900, 11 2022.

[57] J. Duan, R. Xie, S. Yuan, and J. Zhao, "Review of the influence of marketing innovation on consumer behavior in digital economy," *BCP Business & Management*, vol. 23, pp. 723–731, 8 2022.

[58] F. F. Mr, "Intelligent process automation of industries using artificial intelligence and machine learning," *Journal of Computing and Natural Science*, pp. 45–56, 4 2021.