HeilArchive : Pages:1–16

Original Research

Detection and Resolution of Personalization Conflicts and Fatigue in B2C Digital Journeys Using Unified Customer 360 Behavioral and Preference Data

Nguyen Hoang Minh[1], Tran Quoc Dat[2] and Le Van Khoi[3]

[1]Thu Duc Institute of Economics, Department of Business Administration, 145 Le Van Viet Street, Thu Duc City, Ho Chi Minh City 700000, Vietnam.
[2]Hai Duong College of Finance, Department of Accounting and Finance, 92 Nguyen Luong Bang Street, Hai Duong 170000, Vietnam.
[3]Da Nang School of Commerce and Management, Department of International Business, 38 Vo Nguyen Giap Road, Da Nang 550000, Vietnam.

Abstract
Personalization in business to consumer digital channels has become widespread as organizations integrate behavioral, transactional, and preference data into unified customer profiles. Many journeys now involve multiple orchestrated touchpoints across web, mobile, email, messaging, and in product interfaces. As the density and adaptivity of these journeys increase, unintended personalization conflicts and user fatigue become more frequent. Conflicts can appear as inconsistent messages, competing calls to action, or contradictory offers, while fatigue can arise when the volume, frequency, or repetitiveness of personalized content exceeds individual tolerance thresholds. Unified customer 360 data provides a basis to reason about such phenomena at the level of individual users and sessions, but practical detection and resolution mechanisms remain relatively under formalized. This work studies how to use unified behavioral streams, preference declarations, and inferred latent signals to detect and resolve conflicts and fatigue in B2C digital journeys. The paper conceptualizes personalization conflicts and fatigue as properties of sequences of actions and responses over time rather than isolated events. It proposes a set of quantitative representations and model families that operate on customer 360 data and produce interpretable scores and policy recommendations. The emphasis is on connecting modeling constructs to operational decision points such as eligibility, ranking, throttling, and escalation rules. The discussion remains neutral regarding deployment strategies and focuses on outlining design options, possible detection pipelines, and trade offs between sensitivity, precision, and user experience stability.

1. Introduction

Personalization in digital commerce has evolved from simple segment based targeting to highly granular and dynamic decisioning driven by real time signals [1]. Typical B2C organizations now operate a variety of channels including websites, mobile applications, email programs, paid and owned media, in product messaging, and contact center interactions. Each of these channels often has its own optimization objectives and personalization logic. At the same time, customer data platforms aggregate user level behavioral traces, transactional histories, and preference settings into unified profiles. This environment supports more tailored communication but also increases the risk of delivering conflicting or excessive experiences to individual users.

Personalization conflicts arise when decisioning systems produce messages or offers that are inconsistent with each other at the level of one customer journey. A simple form is when two channels simultaneously show incompatible prices or discounts for the same item. More subtle forms include conflicting tones, mismatched lifecycle stages, or contradictory eligibility statuses [2]. These conflicts may be driven by data latency, misaligned business rules, or independently tuned optimization algorithms. In mature ecosystems, conflicts often appear as complex interactions among multiple real time systems rather than as single rule violations.

Personalization fatigue refers to a state where the intensity or style of personalized stimuli leads to diminished engagement or even active avoidance. Users exposed to a high frequency of targeted messages may start to ignore them, unsubscribe from channels, or adopt behaviors that minimize exposure. Fatigue is not solely a function of contact volume. It is shaped by the perceived relevance, diversity, intrusiveness, and predictability of content, as well as by individual sensitivity levels and contextual factors such as time of day or concurrent life events. In practice, fatigue manifests in gradual changes in click, open, dwell, and conversion patterns, as well as in explicit signals such as preference updates or opt outs.

Customer 360 platforms consolidate behavioral log streams, transactional data, preference centers, and inferred attributes into unified representations [3]. These systems are often positioned as the foundation for intelligent personalization and journey orchestration. However, they are equally applicable as a substrate for monitoring and controlling the side effects of intensive personalization. By treating the customer 360 as a longitudinal panel of interactions and responses, one can define quantitative indicators of conflicts and fatigue, and can attach them to decision policies that manage exposure and resolve inconsistencies.

The objective of this paper is to describe a modeling and decision framework for detecting and resolving personalization conflicts and fatigue in B2C journeys using unified customer 360 behavioral and preference data. The focus is on a neutral exposition of definitions, representations, and algorithmic options rather than on advocating a specific architecture. The paper formulates conflicts and fatigue using sequence level constructs, introduces candidate detection models and scores, and outlines how these can feed into resolution policies implemented in real time decision engines. It highlights trade offs among complexity, interpretability, and operational viability and discusses how such models can be evaluated in controlled experiments and observational studies.

2. Customer 360 Behavioral and Preference Data in B2C Digital Journeys

Customer 360 data models are designed to provide a longitudinal, multi channel view of user behavior and attributes. At a high level, such models integrate identities, events, and features into a unified schema. Identity resolution links multiple identifiers, such as cookies, device identifiers, email addresses, and account credentials, into a single customer level entity. Event ingestion brings in page views, app events, message deliveries, opens, clicks, purchases, returns, and service interactions. Feature computation derives aggregates and embeddings at various temporal resolutions. Preference data includes explicit opt ins, opt outs, topic interests, channel preferences, and consent metadata.

In the context of B2C digital journeys, behavioral events exhibit substantial heterogeneity across channels. Web and app events capture browsing and in product usage, often at high frequency. Messaging events capture interactions with outbound campaigns, including exposure to personalized subject lines, recommendations, or offers [4]. Transactional events encode purchases, reservations, or other conversions and typically carry rich metadata about products, prices, and context. Support or service events provide signals about friction and dissatisfaction. A customer 360 platform must organize these disparate streams into well defined tables or feature spaces that can support both batch and real time computation.

Preference data is particularly important for personalization conflict and fatigue analysis. Explicit preferences, such as channel specific frequency caps provided by the user, define hard boundaries that decision systems should respect. Topic or interest preferences constrain content selection and can show contradictions when combined with behavioral inferences. Consent and privacy preferences impose legal and ethical constraints that may override personalization strategies [5]. When preferences exist in multiple isolated systems, conflicts emerge because different channels may hold inconsistent records of what a user has expressed. A unified preference view reduces such inconsistencies and supports consistent enforcement across channels.

From a modeling standpoint, the customer 360 entity can be simplified as a sequence of time stamped tuples. Each tuple contains the customer identifier, the channel or surface, the action taken by the system, the observed user response, and a set of contextual features. This sequential representation enables the use of time series and sequence modeling methods for detecting fatigue and conflicts. It also supports the computation of exposure histories for specific classes of messages, such as promotional offers, reminders, or recommendations, which can be summarized into compact features like rolling counts or recency measures.

The customer 360 must also capture the configuration of personalization logic applied at each interaction. This includes eligibility rules, ranking scores, exploration parameters, and experiment assignments [6]. Without such metadata, it is difficult to attribute conflicts to specific causes. For example, two personalized prices may conflict because different promotional programs were applied, or because one channel uses a cached price list while another uses real time pricing. Including decision metadata in the 360 dataset allows for analysis of which rules and models tend to generate conflicts.

A practical challenge lies in the temporal alignment of data from heterogeneous systems. Events may arrive with different delays and clocks may not be perfectly synchronized. For conflict detection, the temporal ordering of messages is critical, as conflicts are defined relative to windows of user experience. If two conflicting messages are separated by several days, they may not lead to substantial confusion. Therefore, the 360 representation must encode event times with adequate precision and support queries that retrieve sequences within specified windows [7] [8].. For fatigue modeling, cumulative exposure metrics depend on accurate counts and recency within moving horizons.

Finally, the customer 360 dataset should expose both raw and derived attributes related to user state, such as estimated lifecycle stage, predicted propensity scores, churn risk indicators, and latent embedding vectors. These attributes are often produced by separate models and themselves drive personalization decisions. When conflicts and fatigue develop, they may be connected to miscalibrated propensities or embeddings. By incorporating these signals into the detection framework, one can examine whether certain model outputs correlate with elevated conflict or fatigue risk and adjust policies accordingly.

3. Formulation of Personalization Conflicts and Fatigue

To reason about conflicts and fatigue within B2C digital journeys, it is useful to adopt a formal representation of user interaction sequences. Consider a set of customers indexed by i, and a discrete time index t representing ordered interaction steps rather than absolute clock time. Each interaction can be represented as a tuple $(a_{i,t}, c_{i,t}, r_{i,t})$ where $a_{i,t}$ denotes the personalization action chosen by the system, $c_{i,t}$ denotes contextual features for that step, and $r_{i,t}$ denotes the observed user response. The sequence for a customer is $(a_{i,1}, c_{i,1}, r_{i,1}), (a_{i,2}, c_{i,2}, r_{i,2}), \ldots$ and remains open ended.

Personalization actions can be categorized into families such as offers, recommendations, prompts, educational content, and transactional notices [9]. Each action has attributes such as channel, surface position, value proposition, and optionality. To define conflicts, one may introduce a conflict relation over pairs of actions. For two actions a and a' the relation $K(a, a') = 1$ may indicate that these actions are mutually conflicting. The relation may be symmetric or asymmetric. Conflicts can be defined at the level of attributes. For example, two offers with different price points on the same item may be deemed conflicting, as may two messages that simultaneously assert incompatible account states.

Given a conflict relation, a conflict event at step t can be defined when there exists a previous step s within a window W such that $K(a_{i,s}, a_{i,t}) = 1$ and the pair falls within a defined temporal or journey context. The window may be defined in terms of elapsed time, number of interactions, or business concepts like phases within an onboarding journey [10]. This leads to conflict indicators such as

$$k_{i,t} = \max_{s \in W} K(a_{i,s}, a_{i,t}).$$

Here $k_{i,t} = 1$ represents that the action at step t participates in at least one conflict within the window.

Personalization fatigue is less naturally definable as a binary relation and is better conceptualized as a latent state that evolves over time. Let $z_{i,t}$ denote a latent fatigue state for customer i at step t, represented on a continuous scale where higher values indicate greater fatigue. The state evolution may depend on both exposure to personalization and on contextual factors. A simple autoregressive update can be written as

$$z_{i,t} = \gamma z_{i,t-1} + \phi e_{i,t},$$

where $e_{i,t}$ is an exposure measure summarizing the intensity of personalized content delivered between steps $t-1$ and t. Parameters γ and ϕ control decay and accumulation. This formulation treats fatigue as gradually accumulating with exposure and decaying over subsequent periods.

The link between the latent fatigue state and observed user behavior can be expressed through response models. For a binary engagement outcome $r_{i,t} \in \{0, 1\}$ one can posit a logistic form

$$\Pr(r_{i,t} = 1) = \sigma([11] u_{i,t} - \beta z_{i,t}),$$

where $u_{i,t}$ captures baseline utility from the action given features and σ is the logistic function. The coefficient β indicates the degree to which higher fatigue reduces the probability of engagement. While this model is simplified, it illustrates how fatigue can be treated as a latent variable that modulates responsiveness.

Conflicts and fatigue can interact in multiple ways. Conflicts can contribute to fatigue if users perceive them as confusing or manipulative. This can be represented by incorporating conflict indicators into the exposure measure $e_{i,t}$, for example

$$e_{i,t} = w_1 v_{i,t} + w_2 k_{i,t},$$

where $v_{i,t}$ is a volume based exposure metric and $k_{i,t}$ is the conflict indicator. In this conception, a conflict at step t adds extra incremental fatigue compared with a non conflicting action of similar volume. Conversely, elevated fatigue may increase the sensitivity of users to conflicts, leading to a higher likelihood of negative reactions when inconsistencies occur.

From a practical standpoint, the latent state $z_{i,t}$ is not directly observed. Instead, it can be inferred from observed trajectories of $r_{i,t}$ using state space models, recurrent neural networks, or simpler surrogate features such as moving averages of non engagement. The precise form of the model can vary, but the core requirement is to produce a fatigue related score at the individual level that can be updated as new events arrive. This score can then be compared to thresholds or used as an input into decision policies that modulate exposure [12].

The formalization of conflicts and fatigue as functions over sequences makes it possible to define quantitative objectives for decision making. For example, an organization might aim to maximize expected revenue or long term engagement subject to constraints on the frequency of conflicts and the distribution of fatigue scores across the population. By treating conflicts and fatigue as measurable quantities, one can design optimization and control strategies that explicitly account for them rather than treating them as informal concerns.

Concept	Description	Key Risk
Personalization Evolution	Move from segments to granular, real time actions	Increased conflict likelihood
Channel Diversity	Web, app, email, media, in product, contact center	Fragmented decision logic
Unified Profiles	Aggregated behavioral and preference data	Higher exposure intensity

Conflict Type	Origin	Manifestation
Price or Offer Conflict	Data latency, inconsistent pricing systems	Different discounts for same item
Lifecycle Mismatch	Independent journey rules	Contradictory stage messages
Tone or Eligibility Clash	Asynchronous models	Mixed signals on account status

Fatigue Driver	Description	Behavioral Signal
High Exposure Volume	Frequent personalized messages	Lower opens, clicks
Low Content Diversity	Repetitive or predictable content	Diminished responsiveness
Contextual Sensitivity	Time, channel, or user context mismatch	Preference changes, opt outs

Customer 360 Component	Role	Relevance
Behavioral Logs	Sequence of interactions	Exposure history computation
Preference Data	Explicit user choices	Conflict boundary enforcement
Feature and Model Outputs	Embeddings, propensities	Fatigue and inconsistency signals

4. Detection Models Using Unified Behavioral Signals

Detection of personalization conflicts and fatigue requires mapping the formal definitions onto concrete signals derived from the customer 360 dataset. This mapping can be achieved through engineered features, probabilistic models, or representation learning techniques applied to the sequence of interactions and responses. In practice, a layered approach is often required, combining deterministic rule based detection with statistical modeling of latent patterns.

For conflict detection, a starting point is the explicit encoding of conflict relations in a structured catalog. This catalog can be represented as a set of pairs of action templates that are deemed inconsistent [13]. Each template describes a class of actions such as a discount level on a product or a subscription status message. When an interaction occurs, the catalog can be used to generate conflict indicators by inspecting recent and concurrent actions. Rule based detection using such catalogs is straightforward and interpretable but may miss subtle conflicts that arise from combinations of features not covered by predefined templates.

To extend beyond explicit templates, one can learn conflict patterns from data. For example, one can represent actions as vectors in an embedding space derived from content attributes, and define conflicts as pairs of actions whose combined effect on user behavior deviates from expectations. If the joint engagement to two actions delivered close together is systematically lower than what would be predicted from their individual effects, this may indicate a conflict like confusion or annoyance. Let $g(a_{i,s}, a_{i,t})$ denote the observed engagement outcome when two actions co occur. A learned model can approximate the expected engagement $\hat{g}$ under independence and define an interaction residual

$$\delta_{i,s,t} = g(a_{i,s}, a_{i,t}) - \hat{g}(a_{i,s}, a_{i,t}).$$

Negative residuals sustained over many customers and episodes can suggest conflictive combinations [14].

Fatigue detection relies more heavily on temporal patterns of response than on the content of individual messages. A simple feature based method computes rolling statistics such as the moving average of non engagement over recent exposures, the trend in open or click rates, and the time since last positive action. For a given horizon H, one can define a recent engagement rate

$$p_{i,t} = \frac{1}{H} \sum_{h=1}^{H} r_{i,t-h}.$$

Declines in $p_{i,t}$ relative to a baseline can indicate emerging fatigue. However, such raw features can be confounded by seasonality, changes in content quality, or shifts in targeting criteria.

To separate individual fatigue effects from background variation, hierarchical models can be used. For example, a logistic regression with random effects can relate engagement outcomes to exposure intensity and time while controlling for content attributes. At a given step, one may model the log odds of engagement as a linear function of features that include cumulative exposure counts, recency, and context indicators. The random effects capture heterogeneity across customers [15]. From this model, one can derive individual level fatigue scores based on the estimated marginal impact of additional exposure on engagement probability.

More flexible detection models use sequence modeling techniques such as gated recurrent units or temporal convolutional networks to map past exposures and responses to a predicted next step engagement probability and optionally a fatigue score. The input sequence encodes for each step the action type, channel, content embedding, and observed response. The network learns a latent representation of the user state that subsumes fatigue alongside other aspects of readiness or interest. While such models can achieve high predictive accuracy, their outputs may be less interpretable than those of simpler models, which can be a concern when they are used to trigger throttling or escalation decisions that affect large numbers of users.

An intermediate approach is to use hidden Markov models or related state space structures with a small number of discrete latent states representing levels of fatigue or receptivity. For each customer, the model infers probabilities over these states at each step. Transition probabilities depend on exposure intensity and conflict indicators, while emission probabilities describe the likelihood of engagement given the latent state [16]. The resulting state probabilities can be used as concise fatigue indicators. For example, the probability that a user is in a high fatigue state can be compared to a threshold to trigger suppression of additional personalized contacts.

Regardless of the chosen modeling technique, detection outputs must be produced at latencies compatible with operational decisioning. For some channels, such as batch email campaigns, it may be acceptable to compute fatigue and conflict scores in daily or hourly jobs. For others, such as in session web personalization, scores must be updated in near real time as each new event arrives. This places constraints on the complexity of models and the size of features that can be computed on the fly. It also motivates hybrid schemes where a slowly updated base score is combined with fast incremental adjustments.

Calibration and validation of detection models require careful experimental design [17]. One must distinguish genuine fatigue and conflict effects from artifacts of content quality or targeting. Controlled experiments that deliberately vary exposure intensity and combinations of actions, while holding other factors constant, can reveal causal impacts on engagement and satisfaction. Observational analyses that adjust for confounders using matching or weighting techniques can complement experiments when tests are not feasible or ethical. Detection models should be evaluated not only on predictive metrics but also on their stability, interpretability, and robustness to changes in personalization strategies.

Optimization Element	Role in Resolution	Key Consideration
Utility Function $U_{i,t}(a)$	Balances reward, conflict cost, fatigue cost	Requires tunable λ_1, λ_2
Feasible Action Set $\mathcal{A}_{i,t}$	Defines allowable interventions	Constrained by rules and context
Conflict Cost $C_{i,t}(a)$	Penalizes inconsistent actions	Depends on conflict likelihood $k_{i,t}(a)$

Fatigue-Focused Control	Effect	Usage Condition
Throttling Policies	Reduces personalized contact frequency	High fatigue scores or rapid exposure build up
Content Diversification	Introduces variety to mitigate monotony	Repeated recommendation loops
Rest Periods	Pauses specific message types for decay	Users exhibiting sharp engagement decline

Evaluation Layer	Purpose	Primary Output
Simulation	Test policies under hypothetical dynamics	Robustness across parameter ranges
Retrospective Analysis	Identify historical conflict and fatigue patterns	Threshold estimation and risk mapping
Prospective Experiments	Directly measure impact of policy variants	Causal estimates of engagement

5. Resolution Policies and Multi Objective Optimization

Detection of personalization conflicts and fatigue is only useful if it informs decisions that adjust future actions. Resolution in this context refers to the selection of actions or policies that mitigate identified conflicts and fatigue while still pursuing business objectives such as revenue or retention. This can be framed as a multi objective optimization problem where one seeks to balance engagement, economic outcomes, and experience quality metrics [18].

Consider a simplified decision problem at step t for a customer i. The system can choose an action a from a feasible set $\mathcal{A}_{i,t}$, which depends on eligibility rules and contextual constraints. Let $R_{i,t}(a)$ denote the random reward associated with taking action a, capturing outcomes such as conversion or revenue. Let $C_{i,t}(a)$ denote a conflict cost, which may depend on the history of actions and conflict indicators, and let $F_{i,t}(a)$ denote a fatigue cost, which reflects the incremental impact of the action on the fatigue state. A simple scalar utility function can be written as

$$U_{i,t}(a) = \mathbb{E}[R_{i,t}(a)] - \lambda_1 C_{i,t}(a) - \lambda_2 F_{i,t}(a),$$

with weights λ_1 and λ_2 encoding trade offs.

The conflict cost can be operationalized using the conflict indicators described earlier. If $k_{i,t}(a)$ denotes the probability that choosing action a at step t will create a conflict with recent or concurrent actions, one can define

$$C_{i,t}(a) = \eta k_{i,t}(a),$$

where η is an estimated impact of conflict on long term outcomes. Fatigue cost can be approximated using the predicted change in the fatigue state $z_{i,t}$ or directly using an expected reduction in future engagement.

In principle, one could select at each step the action that maximizes $U_{i,t}(a)$ within the feasible set. However, exact computation of expected rewards and costs is generally not tractable in complex journeys. Instead, one can approximate $U_{i,t}(a)$ using predictive models and heuristics. For example, a ranking system can define a score

$$s_{i,t}(a) = \hat{r}_{i,t}(a) - \lambda_1 \hat{k}_{i,t}(a) - \lambda_2 \hat{f}_{i,t}(a),$$

where $\hat{r}$, $\hat{k}$, and $\hat{f}$ are model based estimates of immediate reward, conflict probability, and fatigue impact. The action with the highest score is selected, potentially subject to additional constraints like channel frequency caps.

Resolution policies also operate at aggregate levels [19]. At the channel level, one may allocate overall contact budgets with constraints on conflict and fatigue metrics. For instance, an email channel may be allowed to send messages up to a population level expected conflict rate below a given threshold and an expected fraction of users above a fatigue score threshold below another bound. These constraints can be incorporated into batch optimization problems that determine sending volumes for segments or campaigns.

Dynamic resolution strategies leverage reinforcement learning and control theory. One can frame the multi channel journey as a Markov decision process where the state includes the user profile, recent exposures, conflict indicators, and fatigue score. Actions correspond to choices of messages across channels. Rewards incorporate business outcomes and penalties for conflicts and fatigue. Policy learning methods such as policy gradient or actor critic algorithms can then be used to approximate policies that implicitly balance the multiple objectives [20]. In practice, the state space is large and partially observed, so approximate representations and offline reinforcement learning techniques are often needed.

Hard constraints are frequently required to ensure compliance with legal or ethical requirements and to maintain user trust. For example, explicit channel level preference settings define absolute caps that may not be exceeded even when models suggest potential revenue gains. Conflicts related to regulatory language or critical account status messages must be resolved in favor of safety. Such constraints can be implemented as rule layers that override algorithmic suggestions. In the optimization formulation, hard constraints correspond to restricting the feasible set $\mathcal{A}_{i,t}$ rather than adding penalties.

Resolution of conflicts can take several forms. One is suppression of one of the conflicting actions, typically by prioritizing actions based on a hierarchy of importance. Another is harmonization, where content is adjusted so that messages with different objectives share consistent framing or parameters [21]. A third is deferral, where actions are scheduled at different times to avoid overlapping windows of user experience. The choice among these options can depend on the relative value of actions, their urgency, and the estimated sensitivity of the user.

Resolution of fatigue focuses primarily on modulation of exposure intensity and diversity. Throttling policies reduce the frequency of personalized contacts for users with high fatigue scores, either globally or by channel. Diversification policies alter the mix of content types to introduce more variety or informational value, which can alleviate boredom even when contact volume remains constant. Rest periods, during which certain types of personalized messages are paused, can allow fatigue scores to decay. The design of these policies must consider the potential for missed opportunities and for shifts in user expectations [22].

Multi objective optimization frameworks can help articulate trade offs but often require simplifications and approximations to be implementable. Organizations may choose to pre define acceptable ranges for conflict and fatigue metrics and tune decision systems within these ranges. Alternatively, they may periodically reassess trade offs based on experimental results and customer feedback. In either case, the combination of detection models and resolution policies forms a feedback loop in which observed consequences of decisions inform updates to both models and policies over time.

6. Simulation and Evaluation Framework

Evaluating detection and resolution mechanisms for personalization conflicts and fatigue requires systematic frameworks that can handle both controlled experiments and observational data. Because direct experimentation on negative experiences has ethical and business constraints, simulation plays an important complementary role. Simulated environments allow exploration of policy behavior under a range of assumptions about user dynamics without exposing real customers to potentially problematic treatments.

A basic component of such a framework is a generative model of user behavior in response to personalized actions [23]. This model need not be fully realistic but should capture key aspects of how conflicts and fatigue influence engagement and conversion. For instance, one can define a latent state for each simulated user comprised of a stable preference vector and a dynamic fatigue variable, evolving according to the kind of autoregressive update described earlier. The probability of engagement with an action can be defined as a function of the match between content and preferences, adjusted for fatigue and conflict indicators.

In a simulation environment, one can implement alternative decision policies that use or ignore detection signals. One policy might maximize short term revenue based solely on propensity scores, without regard to conflicts or fatigue. Another might incorporate constraints on exposure intensity and use conflict indicators to suppress inconsistent actions. By simulating large numbers of journeys under these policies, one can estimate distributions of long term outcomes such as cumulative revenue, engagement, opt out rates, and average fatigue levels. Differences among policies reveal trade offs and the potential value of detection and resolution mechanisms [24].

Simulation also supports sensitivity analysis. Parameters controlling the effect of fatigue on engagement, the rate of fatigue decay, and the impact of conflicts on user trust can be varied to observe how policy performance changes. This is important because in real systems these parameters are uncertain and likely heterogeneous across users and contexts. A policy that is robust across a range of parameter values is more desirable than one that is optimal under a narrow set of assumptions but fragile under perturbations.

Beyond simulation, empirical evaluation with real data is essential. One approach is retrospective analysis, where historical logs are used to compute conflict and fatigue indicators for past journeys and to examine their associations with outcomes. For example, one can estimate how the probability of unsubscribe or complaint varies as a function of cumulative exposure and the presence of conflicts in recent windows. Such analyses can suggest thresholds for detection scores that are associated with elevated risk and can inform the design of throttling or suppression policies [25].

Prospective experiments provide more direct evidence of causal effects. For conflict resolution, experiments might compare the existing personalization logic with a version that applies additional conflict detection and suppression. For fatigue management, experiments can compare standard frequency caps with adaptive caps based on fatigue scores. Randomized assignment of users to policy variants permits estimation of causal impacts on engagement, revenue, and satisfaction metrics. Experiments should be designed with safeguards to avoid excessive risk, such as caps on maximum exposure levels and monitoring for adverse signals.

Evaluation metrics should reflect multiple dimensions of performance. Business outcomes such as revenue per user, conversion rates, and retention are important but do not fully capture experience quality [26]. Indicators of negative reactions, including unsubscribe and complaint rates, must be included. Intermediate metrics related to conflicts and fatigue themselves, such as average conflict count per journey or the distribution of fatigue scores, provide more direct views into the behavior of detection and resolution mechanisms. Stability metrics, such as the fraction of users experiencing abrupt changes in exposure patterns, can reveal potential unintended consequences.

Offline evaluation methods can accelerate development cycles by using logged data to approximate the performance of alternative policies without full deployment. Inverse propensity weighting and related counterfactual estimation techniques can adjust for the fact that logged data reflect the behavior of an existing policy. These methods require accurate estimation of the propensities with which actions were chosen in the historical system. When this condition is satisfied, one can evaluate hypothetical policies that would have taken different actions in the same contexts. However, offline evaluation is limited by support issues; policies that diverge too far from the historical decision rules may be evaluated with high variance and bias [27].

Combining simulation, retrospective analysis, offline counterfactual estimation, and prospective experiments can provide a more complete understanding of detection and resolution mechanisms. Each method has strengths and limitations, and their results should be interpreted collectively rather than in isolation. For example, simulation results can guide the design of experiments by identifying promising policy variants and parameter ranges, while retrospective analyses can deliver prior information about plausible effect sizes. Prospective experiments then provide the most direct evidence, which can be used to update models and refine simulation assumptions.

7. Implementation Considerations, Limitations, and Extensions

Translating conceptual detection and resolution frameworks into operational systems involves numerous implementation choices. One central consideration is the integration of detection models with existing decisioning infrastructure. Many organizations already operate recommendation engines, offer management systems, and campaign orchestration platforms. Detection components for conflicts and fatigue must interact with these systems without introducing unacceptable latency or complexity [28]. This often leads to architectures in which detection outputs are materialized as additional features or scores that downstream decision engines can consume via standard interfaces.

Data quality is a critical factor for reliable detection. Incomplete or inconsistent event logs can lead to inaccurate measurement of exposure histories and conflict occurrences. Identity resolution errors can cause the same user to be treated as multiple profiles, fragmenting their true history and undermining fatigue modeling. Preference data that is not synchronized across systems can create apparent conflicts that reflect data misalignment rather than user experience. Implementation therefore requires substantial investment in data validation, reconciliation, and monitoring. Mechanisms to detect anomalies in event rates, identity linkages, and preference distributions are important complements to the specific models for conflicts and fatigue.

Model governance and interpretability present additional challenges [29]. Because detection scores influence exposure to marketing and other messages, they can have material impacts on user experience and business outcomes. Stakeholders may require explanations for why certain users receive fewer messages or why particular actions are suppressed as conflicting. Simple rule based conflict detectors and transparent fatigue scores, such as those based on rolling engagement rates, offer greater interpretability but may be less sensitive than complex models. Organizations may therefore adopt layered approaches where simple indicators provide default safeguards and more sophisticated models offer incremental refinement.

There are inherent limitations in how precisely fatigue and conflicts can be inferred from behavioral data alone. Many external factors influence engagement, including macroeconomic conditions, seasonality, and individual life events, which may be only weakly reflected in observable signals. A decrease in engagement may be misattributed to fatigue when it actually reflects diminished relevance of the product itself or increased competition [30]. Likewise, a conflict that appears obvious in catalog terms may have minimal real impact on user perceptions. These uncertainties suggest caution in interpreting detection outputs and in applying aggressive throttling or suppression based on them.

Ethical and regulatory considerations also shape implementation choices. Personalized experiences interact with privacy regulations that govern consent, data retention, and profiling. Some jurisdictions may treat certain forms of automated decision making that substantially affect users as subject to additional transparency or oversight requirements. Fatigue oriented throttling may reduce the volume of messages sent, which can be seen as beneficial from a privacy perspective, but the underlying inferences about user tolerance may themselves be considered sensitive. Implementations should therefore align with legal guidance and organizational principles regarding fairness, transparency, and user control.

Extensions of the basic framework can incorporate richer notions of user well being and long term relationship health [31]. Fatigue modeling can be broadened to include signals of overload or stress beyond engagement metrics, such as patterns of customer support contacts or sentiment expressed in feedback forms. Conflict detection can be expanded to include semantic inconsistencies in messaging tone or promises rather than only structural contradictions in offers or account states. Natural language processing techniques can be applied to message content to extract representations that can be used in learned conflict models.

Another extension is the explicit modeling of heterogeneity in fatigue dynamics and conflict sensitivity. Different customers may have different baseline tolerance for personalization intensity and different reactions to conflicts. Hierarchical models and clustering approaches can allow the system to learn segments with distinct fatigue profiles. For example, one segment may show steep declines in engagement after a small number of exposures, while another may maintain stable engagement even under frequent contact. Personalized resolution policies can then be tailored to segment specific thresholds and strategies rather than relying on uniform rules [32].

Cross organizational collaboration presents both opportunities and complications. In some industries, users receive personalized messages from multiple brands that may share infrastructure or participate in joint programs. Conflicts and fatigue in such environments may arise from the combined effect of messages from different senders. A unified customer 360 within one organization may not capture these external influences, limiting the accuracy of detection models. Data sharing arrangements or federated learning approaches could, in principle, mitigate this limitation while respecting privacy constraints, though they introduce additional governance and technical challenges.

The framework described in this paper is also subject to temporal drift. As personalization practices, channel mix, and user expectations evolve, the patterns of conflicts and fatigue may change. Models and policies that are effective under one regime may become misaligned under another [33]. Implementation should therefore include mechanisms for ongoing monitoring and recalibration. Drift detection methods can signal when the distribution of inputs or outputs of detection models shifts significantly, prompting review and potential retraining. Periodic re evaluation of thresholds and constraints in resolution policies can ensure that they remain aligned with current conditions.

Finally, it is important to acknowledge that resolving personalization conflicts and fatigue is not solely a technical modeling problem. Organizational processes, governance structures, and incentive systems play a substantial role. For instance, if different channels or product teams are rewarded primarily on short term performance metrics without regard to cross channel conflicts or fatigue, technical solutions may have limited impact. Aligning incentives and decision rights so that conflict and fatigue metrics are treated as shared responsibilities can enhance the effectiveness of detection and resolution mechanisms [34].

8. Organizational Integration and Governance

Operationalizing detection and resolution of personalization conflicts and fatigue in B2C digital journeys requires more than the deployment of models and decision rules. It depends on how these technical components are embedded in organizational structures, governance processes, and day to day operational practices. In many organizations, personalization logic emerges from a combination of marketing teams, product groups, data science units, engineering, and legal or compliance functions. Each group brings its own objectives, vocabularies, and constraints. Without deliberate coordination, the resulting ecosystem of rules and models can be fragmented even when a unified customer 360 data platform exists. This section examines how organizational integration and governance can support coherent management of conflicts and fatigue, focusing on roles, processes, and tooling rather than on algorithmic details.

A central consideration is the definition of ownership for conflict and fatigue metrics. In some organizations, channel teams maintain independent frequency caps and quality controls, while centralized analytics teams produce cross channel performance dashboards [35]. When detection of conflicts and fatigue is added, ambiguity can arise over who is responsible for acting on elevated risk signals and for adjusting policies that may affect multiple channels simultaneously. One approach is to designate a cross functional journey governance group or council that holds accountability for aggregate conflict and fatigue indicators. This group may include representatives from marketing, product, analytics, engineering, customer support, and privacy or compliance. It can review metrics, approve policy changes, and arbitrate trade offs when channel objectives and experience quality considerations diverge. The existence of such a group does not remove local responsibility from individual teams but provides a forum in which cross channel issues are surfaced and resolved.

The design of metrics for operational use is another key element of governance. Detection models can generate many scores and indicators at different levels of granularity, but operational stakeholders often require a concise set of measures that can be monitored and discussed. For conflicts, organizations may define measures such as average number of conflicting messages per user per period, or proportion of journeys with at least one conflict in a given phase [36]. For fatigue, they may track distributions of fatigue scores, proportions of users above certain thresholds, or changes in engagement as a function of exposure intensity. Translating complex model outputs into stable metrics that can be incorporated into key performance indicator frameworks is non trivial. It requires collaboration between data science teams, who understand model behavior, and business stakeholders, who understand how metrics drive incentives and decision making [37].

Governance processes should include explicit stages for the introduction and modification of detection models and resolution policies. When a new fatigue score is proposed, for example, there should be a documented procedure for validation, including back testing on historical data, controlled experiments where feasible, and reviews by relevant stakeholders. Similarly, when changes to throttling rules or conflict suppression policies are planned, the potential impacts on revenue, engagement, and customer experience should be analyzed in advance. Change proposals can be captured in lightweight design documents that articulate objectives, assumptions, expected effects, and monitoring plans. These documents can be reviewed by the cross functional governance group and archived for future reference [38]. Over time, the organization can build a knowledge base of policy changes and their observed consequences.

Tooling plays a significant role in making detection and governance actionable. Dashboards that visualize conflict and fatigue metrics at multiple levels of aggregation can help channel teams and executives understand where issues are concentrated. For example, a dashboard might show conflict rates by channel pair, revealing that inconsistencies between email and mobile app notifications are more frequent than those between web and call center interactions. Another dashboard might show the distribution of fatigue scores across user segments, highlighting whether certain cohorts, such as new subscribers or high value customers, experience systematically higher fatigue. Effective dashboards allow users to drill down from aggregate metrics to specific journeys or message combinations, enabling root cause analysis and targeted remediation.

Integrating conflict and fatigue signals into existing campaign and personalization tools is equally important [39]. If detection outputs are only visible in analytical dashboards but do not influence the tools used to design and launch campaigns, the burden falls on human operators to interpret and act on the signals. To reduce reliance on manual interpretation, conflict and fatigue indicators can be surfaced contextually within tooling. For instance, when a marketer configures a new campaign targeting a certain cohort at a high frequency, the system can display warnings or guidance based on the current fatigue profile of that cohort or on historical conflicts involving similar content. In more advanced implementations, the system can automatically adjust send volumes or eligibility criteria based on predefined rules that reference detection scores, while still allowing humans to review and approve changes.

The relationship between detection driven controls and legal or regulatory obligations requires careful management. Privacy and consumer protection regulations may influence how personalization is conducted, which data can be used, and how automated decisions must be explained. Governance structures need to ensure that conflict and fatigue management remains aligned with these requirements. For example, if users have explicit rights to limit profiling or automated decision making, the mechanisms that infer fatigue states or suppress messages based on such inferences should be consistent with the consent provided [40]. Legal and compliance teams can participate in the design of detection models and policies, reviewing how user data is used, how inferences are generated, and how decisions are communicated. This collaboration can reduce the risk of implementing technically sophisticated but legally questionable controls.

Organizational incentives and performance measurement systems can either support or undermine conflict and fatigue controls. If channel teams are rewarded predominantly on short term metrics such as immediate revenue or click rates, they may resist policies that limit contact volumes or suppress potentially profitable offers in order to reduce conflicts. Conversely, if cross channel metrics that incorporate conflict and fatigue indicators are included in performance reviews and planning, teams may be more inclined to cooperate. This does not imply that all existing incentives must be replaced, but rather that they should be adjusted to acknowledge the importance of experience quality and long term relationship health. Leadership communication can reinforce that conflict and fatigue management is a shared responsibility that aligns with broader organizational goals such as customer satisfaction and brand trust.

Training and knowledge sharing support the practical adoption of detection and resolution mechanisms [41]. Non technical stakeholders may be unfamiliar with concepts such as latent fatigue states, conflict relations among actions, or state based throttling. Without sufficient understanding, they may misinterpret detection outputs or perceive them as arbitrary constraints imposed by models. Training sessions that explain the intuition behind fatigue scores, the kinds of conflicts that have been observed, and the trade offs inherent in resolution decisions can build confidence. Case studies derived from internal data, even when anonymized and aggregated, can illustrate how certain combinations of actions produced conflicts or how high exposure for some segments preceded rises in unsubscribe rates. These narratives can make abstract concepts more tangible and support more informed decisions.

Incident management procedures provide a complementary governance layer for situations where conflicts or fatigue related issues become acute. For example, if a configuration error leads to a surge in conflicting messages or an unexpected spike in contact volume to a particular segment, organizations need clear processes for detection, triage, and remediation. This may include automatic alerts when conflict or fatigue metrics exceed defined thresholds, escalation paths to relevant teams, and playbooks that describe immediate mitigation steps, such as suspending certain campaigns or overriding default throttling rules [42]. Post incident reviews can analyze root causes, including both technical and organizational factors, and propose improvements in detection models, policies, or operational processes.

Over time, organizations may seek to standardize their approach to conflict and fatigue governance across regions, brands, or business units. Standardization can simplify training, tooling, and compliance, but it must be balanced against local differences in user expectations, regulatory regimes, and business models. A centralized governance framework can define core principles, such as maximum acceptable conflict rates or general approaches to fatigue aware throttling, while allowing local teams to adapt specific thresholds and resolution strategies. Shared repositories of detection models and policy templates can facilitate reuse and reduce duplication of effort, while regular cross unit reviews can identify patterns and learnings that span contexts.

There are also strategic considerations about how transparent to be with users regarding personalization intensity and conflict resolution. Some organizations choose to expose preference centers where users can adjust frequency settings, choose channels, or specify topics of interest. Detection models can inform the design of these interfaces, for example by suggesting default settings that reflect observed tolerance patterns, or by prompting users in high fatigue states to adjust their preferences proactively. Communication about how personalization works, including acknowledgement that systems may adjust contact intensity to avoid overload, can contribute to trust if designed carefully. At the same time, excessive detail or technical language may confuse users, so collaboration between product, design, and legal teams is needed to craft appropriate messages.

Finally, organizational integration and governance should be viewed as evolving rather than static. As new channels emerge, as customer expectations shift, and as regulatory frameworks develop, the definitions of conflicts, the manifestations of fatigue, and the levers available for resolution will change. Governance bodies may need to periodically revisit metrics, thresholds, and roles. Detection models may require retraining or redesign. Resolution policies may need recalibration as business strategies or macroeconomic conditions shift [43]. Embedding a culture of continuous learning, where detection outputs, experimental results, and incident analyses feed back into governance deliberations, can help organizations adapt. In this way, conflict and fatigue management becomes an ongoing discipline integrated into broader personalization and customer experience strategies rather than a one time technical project.

9. Conclusion

Personalization in B2C digital journeys has progressed to a level of complexity where unintended conflicts and user fatigue are natural byproducts of dense and adaptive interaction patterns. Unified customer 360 behavioral and preference data offers a foundation for systematically detecting and mitigating these phenomena. By representing journeys as sequences of actions, contexts, and responses, and by defining conflicts and fatigue as properties of these sequences, organizations can move beyond ad hoc rules toward more principled monitoring and control.

The paper has outlined how customer 360 data can be organized and used to define conflict relations among actions and to construct latent models of fatigue that evolve with exposure. It has discussed detection techniques that range from rule based catalogs and engineered features to probabilistic state space models and sequence based representation learning. These techniques produce scores and indicators that can be integrated into decision systems to adjust eligibility, ranking, throttling, and escalation policies in ways that account for conflicts and fatigue alongside traditional business objectives [44].

Resolution was framed as a multi objective optimization problem in which expected rewards are balanced against conflict and fatigue costs, under both hard and soft constraints. Approaches such as heuristic scoring, constrained optimization, and reinforcement learning can be used to approximate policies that respect defined thresholds and trade offs. Simulation and empirical evaluation frameworks were described as tools to understand the behavior of detection and resolution mechanisms under varying assumptions and to measure their impact on engagement, revenue, and experience quality.

Implementation considerations highlight that data quality, model governance, and alignment with legal and ethical requirements are integral to practical deployment. Limitations in observability, uncertainties in causal pathways, and temporal drift imply that detection outputs and resolution policies should be applied with appropriate caution and subject to ongoing monitoring. Extensions such as richer well being models, heterogeneity aware fatigue dynamics, and cross organizational perspectives suggest directions for further development.

Overall, the use of unified customer 360 data for the detection and resolution of personalization conflicts and fatigue can provide a structured way to manage some of the side effects of intensive personalization. While modeling and optimization techniques can support more informed decision making, their effectiveness depends on integration with broader organizational practices and on continued reassessment as digital ecosystems evolve [45].

References

[1] U. Gretzel, M. Sigala, Z. Xiang, and C. Koo, "Smart tourism: foundations and developments," *Electronic Markets*, vol. 25, pp. 179–188, 8 2015.

[2] C. Åberg and M. Torchia, "Do boards of directors foster strategic change? a dynamic managerial capabilities perspective," *Journal of Management and Governance*, vol. 24, pp. 655–684, 4 2019.

[3] D. W. Stewart and I. M. Martin, "Advertising disclosures: Clear and conspicuous or understood and used?:," *Journal of Public Policy & Marketing*, vol. 23, pp. 183–192, 9 2004.

[4] N. Singh and S. K. Kundu, "Explaining the growth of e-commerce corporations (eccs): An extension and application of the eclectic paradigm," *Journal of International Business Studies*, vol. 33, pp. 679–697, 12 2002.

[5] S. Quach, P. Thaichon, K. D. Martin, S. Weaven, and R. W. Palmatier, "Digital technologies: tensions in privacy and data.," *Journal of the Academy of Marketing Science*, vol. 50, pp. 1299–1323, 3 2022.

[6] J. W. Moffett, J. A. G. Folse, and R. W. Palmatier, "A theory of multiformat communication: mechanisms, dynamics, and strategies," *Journal of the Academy of Marketing Science*, vol. 49, pp. 1–21, 11 2020.

[7] M. K. Erramilli, "The experience factor in foreign market entry behavior of service firms," *Journal of International Business Studies*, vol. 22, pp. 479–501, 9 1991.

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

[9] M. D. Giebelhausen, B. Lawrence, and H. H. Chun, "Doing good while behaving badly: Checkout charity process mechanisms," *Journal of Business Ethics*, vol. 172, pp. 133–149, 1 2020.

[10] S. Sridhar and E. Fang, "New vistas for marketing strategy: digital, data-rich, and developing market (d3) environments," *Journal of the Academy of Marketing Science*, vol. 47, pp. 977–985, 10 2019.

[11] C. R. Plouffe, J. Hulland, and T. Wachner, "Customer-directed selling behaviors and performance: A comparison of existing perspectives," *Journal of the Academy of Marketing Science*, vol. 37, pp. 422–439, 4 2009.

[12] D. Easley, M. O'Hara, and L. Yang, "Differential access to price information in financial markets," *Journal of Financial and Quantitative Analysis*, vol. 51, pp. 1071–1110, 11 2016.

[13] P. V. Zhang, S. Kim, and A. Chakravarty, "Influence of pull marketing actions on marketing action effectiveness of multichannel firms: A meta-analysis," *Journal of the Academy of Marketing Science*, vol. 51, pp. 310–333, 6 2022.

[14] R. R. Burke, "Technology and the customer interface: What consumers want in the physical and virtual store:," *Journal of the Academy of Marketing Science*, vol. 30, pp. 411–432, 10 2002.

[15] E. Markos, G. R. Milne, and J. W. Peltier, "Information sensitivity and willingness to provide continua: A comparative privacy study of the united states and brazil:," *Journal of Public Policy & Marketing*, vol. 36, pp. 79–96, 4 2017.

[16] S. M. Lee and D. Lee, "Opportunities and challenges for contactless healthcare services in the post-covid-19 era," *Technological forecasting and social change*, vol. 167, pp. 120712–120712, 2 2021.

[17] A. S. Gabriel, J. D. Acosta, and A. A. Grandey, "The value of a smile: Does emotional performance matter more in familiar or unfamiliar exchanges?," *Journal of Business and Psychology*, vol. 30, pp. 37–50, 10 2013.

[18] M. Harvey and D. A. Griffith, "Developing effective intercultural relationships: The importance of communication strategies," *Thunderbird International Business Review*, vol. 44, pp. 455–476, 6 2002.

[19] A. P. D'Costa, "Institutions and industrial governance in india: Learning to cooperate the japanese way," *Asian Business & Management*, vol. 2, pp. 63–89, 4 2003.

[20] A. L. Roggeveen, M. Tsiros, and D. Grewal, "Understanding the co-creation effect: When does collaborating with customers provide a lift to service recovery?," *Journal of the Academy of Marketing Science*, vol. 40, pp. 771–790, 8 2011.

[21] K. Martin and A. Waldman, "Are algorithmic decisions legitimate? the effect of process and outcomes on perceptions of legitimacy of ai decisions," *Journal of Business Ethics*, vol. 183, pp. 653–670, 2 2022.

[22] M. Zhalechian, E. Keyvanshokooh, C. Shi, and M. P. V. Oyen, "Online resource allocation with personalized learning," *SSRN Electronic Journal*, 1 2020.

[23] 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.

[24] J. Abaluck and J. Gruber, "Evolving choice inconsistencies in choice of prescription drug insurance.," *The American economic review*, vol. 106, pp. 2145–2184, 8 2016.

[25] L. Anderson, J. Spanjol, J. G. Jefferies, A. L. Ostrom, C. N. Baker, S. A. Bone, H. Downey, M. Mende, and J. M. Rapp, "Responsibility and well-being: Resource integration under responsibilization in expert services," *Journal of Public Policy & Marketing*, vol. 35, pp. 262–279, 9 2016.

[26] T. Loper and V. L. Crittenden, "Energy security: Shaping the consumer decision making process in emerging economies," *Organizations and Markets in Emerging Economies*, vol. 8, pp. 8–32, 5 2017.

[27] B. M. Noone and K. A. McGuire, "Impact of attitudinal loyalty on the frequent unmanaged business traveler's use of price and consumer reviews in hotel choice," *Journal of Revenue and Pricing Management*, vol. 15, pp. 20–36, 2 2016.

[28] M. Granovetter, "The impact of social structure on economic outcomes," *Journal of Economic Perspectives*, vol. 19, pp. 33–50, 2 2005.

[29] C.-H. Tan, H.-H. Teo, and H. Xu, "Online auction: the effects of transaction probability and listing price on a seller's decision-making behavior," *Electronic Markets*, vol. 20, pp. 67–79, 2 2010.

[30] S. Spiekermann, R. Böhme, A. Acquisti, and K. L. Hui, "Personal data markets," *Electronic Markets*, vol. 25, pp. 91–93, 4 2015.

[31] G. Ellison and S. F. Ellison, "Lessons about markets from the internet," *Journal of Economic Perspectives*, vol. 19, pp. 139–158, 4 2005.

[32] C. F. Miao and K. R. Evans, "The interactive effects of sales control systems on salesperson performance: a job demands-resources perspective," *Journal of the Academy of Marketing Science*, vol. 41, pp. 73–90, 10 2012.

[33] Z. S. Byrne, K. Cave, and S. D. Raymer, "Using a generalizable photo-coding methodology for assessing organizational culture artifacts.," *Journal of business and psychology*, vol. 37, pp. 1–15, 10 2021.

[34] C. R. Sunstein, "Ruining popcorn? the welfare effects of information," *Journal of Risk and Uncertainty*, vol. 58, pp. 121–142, 5 2019.

[35] D. L. Haytko, "Firm-to-firm and interpersonal relationships: Perspectives from advertising agency account managers," *Journal of the Academy of Marketing Science*, vol. 32, pp. 312–328, 7 2004.

[36] G. M. Eckhardt and S. Dobscha, "The consumer experience of responsibilization: The case of panera cares," *Journal of Business Ethics*, vol. 159, pp. 651–663, 1 2018.

[37] 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.

[38] J. Liu, H. Zhao, J. Li, and X. Yue, “Operational strategy of customized bus considering customers’ variety seeking behavior and service level,” *International journal of production economics*, vol. 231, pp. 107856–107856, 7 2020.

[39] M. Benmamoun, H. Alhor, C. Ascencio, and W. Sim, “Social enterprises in electronic markets: web localization or standardization,” *Electronic Markets*, vol. 31, pp. 215–231, 7 2020.

[40] L. Jiang, M. Jun, and Z. Yang, “Customer-perceived value and loyalty: how do key service quality dimensions matter in the context of b2c e-commerce?,” *Service Business*, vol. 10, pp. 301–317, 2 2015.

[41] R. T. Wigand, “Whatever happened to disintermediation,” *Electronic Markets*, vol. 30, pp. 39–47, 2 2020.

[42] M. E. Schweitzer and D. E. Gibson, “Fairness, feelings, and ethical decision- making: Consequences of violating community standards of fairness,” *Journal of Business Ethics*, vol. 77, pp. 287–301, 3 2007.

[43] K. K. Myers and K. Sadaghiani, “Millennials in the workplace: A communication perspective on millennials’ organizational relationships and performance,” *Journal of business and psychology*, vol. 25, pp. 225–238, 3 2010.

[44] A. Saini and J. L. Johnson, “Organizational capabilities in e-commerce: An empirical investigation of e-brokerage service providers,” *Journal of the Academy of Marketing Science*, vol. 33, pp. 360–375, 7 2005.

[45] M. G. Dekimpe, I. Geyskens, and K. Gielens, “Using technology to bring online convenience to offline shopping,” *Marketing Letters*, vol. 31, pp. 25–29, 12 2019.