Research

EPOCH Journals

# A Data-Driven Approach Using Customer 360, Next-Best-Action Models, and Real-Time Analytics

Youssef El Amrani[1] and Ali Benjelloun[2]

[1]Cadi Ayyad University, Faculty of Economics and Management, Avenue Abdelkrim Khattabi, Gueliz, Marrakech 40000, Morocco
[2]Hassan II University of Casablanca, Department of Business Administration, Km 7 Route dEl Jadida, Casablanca 20000, Morocco



## Abstract

Organizations seeking consistent and contextually aligned interactions across channels have increasingly adopted architectures that unify operational, behavioral, and contextual data into comprehensive views of individual customers. Advances in scalable storage, event streaming, and machine learning have made it feasible to maintain such consolidated representations while responding to interactions as they occur. At the same time, customer engagement strategies have shifted from static, rule-based campaigns toward decisions that are continuously optimized given constraints such as regulatory requirements, capacity limits, and fairness considerations. In this setting, next-best-action models and real-time analytics are often presented as complementary capabilities, yet their operational integration with Customer 360 platforms remains nontrivial and subject to technical trade-offs. This paper examines a data-driven approach that integrates Customer 360 representations, decisioning models for next best action, and low-latency analytical pipelines into a coherent architecture suitable for high-volume environments. The proposed formulation emphasizes stable entity resolution, feature governance, and action arbitration across multiple objectives and horizons without relying on a single optimization paradigm. The discussion covers online learning, feedback attribution, and robustness under partial observability and delayed responses, together with the implications of hard latency and throughput constraints. The goal is to characterize how such systems can be designed, analyzed, and tuned so that decisions remain interpretable, computationally feasible, and aligned with explicit performance criteria, without assuming a unique or universal optimal configuration.

## Introduction

Modern customer engagement infrastructures frequently emerge from the convergence of three technological strands: integrated customer data platforms capable of synthesizing heterogeneous sources into persistent profiles, algorithmic decision engines that compute next best actions conditioned on those profiles, and real-time analytics layers that track, evaluate, and adapt behaviors under streaming interaction data [1]. The joint evolution of these strands reflects a shift from offline, aggregated reporting toward continuous, data-driven control of individualized interactions. However, this shift introduces architectural, statistical, and operational challenges that are not fully addressed by isolated treatments of recommendation systems, campaign management, or streaming analytics. The interactions among data latency, model uncertainty, constraint handling, and feedback mechanisms tend to be complex, context-dependent, and sensitive to small deviations in pipeline assumptions, making ad hoc integration strategies difficult to validate and maintain.

A central construct often adopted in practice is the Customer 360 representation, defined informally as the unified, longitudinal, and multi-channel view of an identifiable entity, along with the derived signals that summarize its historical and current states [2]. Although widely referenced, the precise formulation of Customer 360 is frequently implicit, leading to ambiguity around what constitutes a sufficient state for downstream decisioning and evaluation. In operational environments, this view is further complicated by probabilistic identity resolution, dynamic consent states, evolving schemas, missingness patterns, and asynchronous event arrival across systems of record. Any attempt to layer next-best-action models and real-time analytics on top of this representation relies on an explicit understanding of these limitations, since misspecified or unstable state estimates propagate directly into policy behaviors and derived performance metrics. A data-driven approach in this setting benefits from treating the construction and update rules for Customer 360 as elements of a formally specified and testable state evolution process rather than as undocumented configuration artifacts.

Next-best-action models are commonly introduced as mechanisms that recommend actions, offers, or messages at the individual level, given objectives such as incremental revenue, retention, satisfaction, risk mitigation, or operational efficiency [3]. A neutral perspective treats these models as parameterized policy functions over a structured state space rather than as prescriptive guarantees of optimal behavior. Under this view, policies map current representations of customer state and context into probability distributions over candidate actions. They must reconcile multiple objectives, adhere to regulatory and contractual requirements, respect channel capacities and contact governance, and adapt to feedback signals that may be indirect, delayed, censored, or confounded by concurrent interventions and exogenous changes. Simple ranking or uplift models can yield adequate behavior in narrowly scoped applications, yet their properties in large-scale, multi-objective, multi-channel environments are difficult to characterize without a more explicit mathematical framework that represents uncertainty, temporal dependencies, and the influence of earlier actions on subsequent opportunities.

Real-time analytics close the loop by providing short-latency measurement, monitoring, and adaptation signals that inform both representation maintenance and policy adjustment [4]. While frequently described in terms of dashboards and key performance indicators, their structural role in an integrated system is to define how information about environment responses, operational states, and constraint utilization flows back into the decision pipeline. These feedback channels affect stability, identifiability of treatment effects, and the ability to detect, attribute, and respond to anomalies in a timely manner. When real-time analytics are weakly coupled to decisioning, degradation modes arise in which actions depend on stale or inconsistent information, empirical evaluations become biased, and constraint breaches are discovered only after prolonged violation. Conversely, when analytics drive automated adjustments without regard to uncertainty, measurement error, or latency, the risk of oscillatory or unstable behavior increases, potentially obscuring whether observed performance changes reflect genuine improvements or artifacts of reactive tuning.

The context for this work is the growing prevalence of high-volume interaction environments in which organizations apply algorithmic decisioning across web properties, mobile applications, notification channels, contact centers, and assisted sales or service processes [5]. In such settings, decisions are executed at fine temporal resolutions yet are governed by strategic objectives and policy guidelines defined at much coarser horizons. A data-driven architecture that integrates Customer 360, next-best-action models, and real-time analytics must therefore support heterogeneous temporal scales, from millisecond-level scoring and routing to quarterly or annual assessments of portfolio health, regulatory compliance, and brand impact. This heterogeneity complicates alignment between local decision rules and aggregate outcomes. Policies tuned solely on short-term responses risk underrepresenting actions whose primary impact lies in medium- or long-term outcomes, such as reduced churn, improved satisfaction, or mitigated risk exposures, unless these horizons are explicitly encoded in state, objectives, and monitoring. [6]

The formulation adopted here approaches these challenges by viewing the combined system as a composition of state construction, decision policy, and feedback estimation operators acting on streaming data. Rather than focusing on specific technologies or domain-specific implementations, the discussion abstracts to representations, mappings, and constraints. Customer 360 is treated as a dynamic state whose evolution depends on ingested events, identity graphs, eligibility and consent logic, and feature governance; next-best-action policies are functions that map this state into probabilistic recommendations under explicit or configurable objectives; and real-time analytics are operators that estimate performance functionals, detect drift and anomalies, and track constraint indicators along the realized trajectories. This abstraction enables analysis of consistency conditions between training and deployment environments, examination of how architectural choices influence statistical properties of estimators and learned policies, and structured reasoning about trade-offs between responsiveness and robustness.

Several themes motivate a more technical treatment of this integration [7]. First, the reliance on streaming interaction data introduces nontrivial temporal structures, including out-of-order event arrival, delayed logging, right-censoring of outcomes, and overlapping attribution windows across multiple actions. Ignoring these structures when constructing features or estimating effects can yield biased or operationally infeasible models, particularly when learned functions implicitly rely on information that is not available at decision time. Second, the same events that update Customer 360 states often serve as the basis for computing evaluation metrics and for driving model updates, creating closed feedback loops in which policies shape data and data shape policies. Without explicit modeling of these loops, it becomes difficult to separate genuine performance changes from shifts induced by targeting strategies, drift in identity resolution, or evolving traffic compositions. Third, practical deployments must respect constraints that do not naturally emerge from unconstrained optimization, including contact frequency and suppression rules, channel and workforce capacities, and fairness or inclusion guidelines defined across segments and protected attributes. [8]

A central objective of this work is to articulate an integrated framework that captures these interactions while remaining neutral with respect to particular algorithmic or vendor choices. The emphasis is on characterizing design spaces and dependencies rather than prescribing a single solution. Customer 360 constructions may rely on deterministic aggregations of transactional histories, learned sequence embeddings, or hybrid feature sets; next-best-action policies may implement parametric scoring rules, contextual bandit algorithms, or approximate reinforcement learning; real-time analytics may employ simple sliding window estimates or more elaborate streaming estimators with stratification and variance control. What is common across such variants is that they operate on shared data flows, place implicit assumptions on one another, and jointly determine observable behavior. By representing them within a unified mathematical structure, one can identify conditions under which joint behavior is stable, constraints are enforced coherently, performance is estimable with acceptable uncertainty, and updates are interpretable. [9]

Within this framework, several guiding questions emerge. How should the Customer 360 state be defined so that it is sufficiently informative for decisioning while remaining stable under incremental updates, changes in identity linkage, and schema evolution. How should next-best-action policies be specified so that they exploit available information without relying on unverifiable causal assumptions, and how should their adaptation mechanisms be regularized to avoid overreaction to short-term noise or transient shifts. In what ways can real-time analytics be configured to provide timely yet statistically grounded signals for adaptation, anomaly detection, and risk monitoring, and how should those signals interface with both representation and policy update procedures. What experimentation strategies are suitable when policies, state representations, and feedback estimators are all evolving, and how can those strategies be reconciled with operational and ethical constraints. [10]

A further consideration is the distinction between descriptive unification of customer data and prescriptive use of that data for automated decisions. Many implementations achieve comprehensive consolidation of identifiers and events without specifying how that information will drive action selection, calibration, and monitoring. Conversely, decision engines may operate effectively over stable engineered features while remaining agnostic to how those features are produced or governed. In the absence of an explicit contract between these layers, modifications in upstream logic can silently alter the semantics of downstream inputs, complicating evaluation and incident analysis [11]. By treating Customer 360 as an integral component of the decision system, the framework encourages explicit definitions of which features are stable, which are experimental, and how lineage, masks, and governance metadata propagate into training, validation, and runtime contexts.

Another issue is the interaction between algorithmic decision policies and human agents. Many organizations deploy next-best-action models not as fully autonomous controllers but as decision support tools for service or sales representatives. In such hybrid workflows, observed actions emerge from a composition of model-generated recommendations, user interfaces, and human judgment. The effective policy is therefore a joint function of scores and agent behavior [12]. Evaluating and improving decision strategies in this setting requires acknowledging that overrides, selective attention, and learned responses by human users can systematically shape both action and outcome distributions. The conceptualization adopted here accommodates such hybrid configurations by treating observed actions as samples from an implicit joint policy and by allowing real-time analytics to characterize and monitor this composite behavior.

The introduction also recognizes that integrated decision systems operate within nonstationary external environments. Customer preferences, competitive landscapes, macroeconomic conditions, and regulatory frameworks evolve, inducing shifts in both covariate and outcome distributions. These shifts can undermine assumptions encoded in Customer 360 features, invalidate calibrations in next-best-action models, and alter the baselines against which real-time metrics are interpreted [13]. Periodic offline retraining alone may be insufficient when change is rapid or structural. Explicit drift detection, temporal validation schemes, cautious adaptation rules, and scenario-based stress testing become relevant components of a data-driven approach. The aim is not to eliminate drift, but to incorporate it as an expected property of the environment that informs model design, monitoring thresholds, and governance practices.

Finally, there is a practical motivation for a unified formulation in the modular composition of contemporary technology stacks. Capabilities for identity management, event streaming, feature computation, decision orchestration, and analytics are frequently distributed across distinct platforms or services maintained by different teams [14]. Without shared abstractions, it is difficult to reason consistently about their interactions or to diagnose emergent behaviors across boundaries. By grounding the discussion in a common vocabulary of states, policies, constraints, and feedback operators, the framework supports clearer interface specifications, more predictable impact assessment of localized changes, and more systematic collaboration among stakeholders responsible for data, models, operations, and compliance. The introduction therefore positions the subsequent technical developments not as endorsements of a particular product or organizational structure, but as tools for clarifying and structuring integration patterns that already arise in diverse settings.

In summary, the introduction frames the problem of integrating Customer 360, next-best-action models, and real-time analytics as one of designing and analyzing coupled data and decision systems under realistic technical and organizational constraints [15]. It highlights that the central questions extend beyond predictive accuracy or short-term response optimization to encompass representational adequacy, robustness to drift and delays, coherence of automated and human decision-making, and reliability of monitoring. The following sections build on this foundation by specifying the data assumptions, mathematical structures, and implementation considerations that can support such systems, with attention to both their capabilities and their limitations in varied operational contexts.

**Integration of Data, Decision, and Analytics**

Modern customer engagement systems integrate three key strands: **Customer 360** (data synthesis), **Next-Best-Action models** (decision engines), and **Real-Time Analytics** (feedback loops). Their convergence enables continuous individualized control but introduces architectural and statistical challenges.

**Customer 360 as Dynamic State Representation**

Customer 360 is a unified yet evolving customer profile that aggregates multi-channel, [16] longitudinal data. It must handle identity uncertainty, asynchronous event arrival, and regulatory masking. Treating it as a formally defined, testable state evolution process improves stability and interpretability.

**Next-Best-Action Policies**

Next-best-action models act as parameterized policy functions mapping states to action probabilities. They must reconcile multiple objectives, regulatory limits, and feedback delays while maintaining robustness [17] under uncertainty and temporal dependencies.

**Role of Real-Time Analytics**

Real-time analytics provide low-latency performance estimation and feedback for decision adaptation. When tightly coupled, they ensure timely detection of drift or constraint breaches; when weakly coupled, they risk instability, delayed correction, and biased evaluations.

**Multi-Temporal and Organizational Challenges**

Integrated systems must operate across multiple time scalesfrom milliseconds (individual scoring) [18] to months (strategic metrics). Misalignment between local optimizations and long-term goals can bias behavior unless explicitly modeled.

**Unified Framework and Guiding Questions**

A unified framework treats Customer 360, policy, and analytics as interacting operators on streaming data. Key questions include: How to define stable states? How to adapt policies under uncertainty? How to ensure feedback signals are valid and ethical in evolving environments? [19]

## Data Foundations for Customer 360

A Customer 360 representation can be formalized as a state vector associated with an entity index, constructed from heterogeneous data sources that include transactions, interactions, product holdings, channel behaviors, and contextual attributes. Let the set of customers be denoted by an index set, and for each entity define an evolving state that aggregates historical and current information. The construction of this state is subject to several technical constraints. Identity resolution is inherently probabilistic whenever multiple identifiers, imperfect keys, or device-level signals are present [20]. Event streams from operational systems arrive with variable and sometimes unknown delays, and attributes may be censored or missing in nonrandom patterns. These constraints imply that the Customer 360 state is not simply a deterministic concatenation of observations, but rather a filtered estimate of underlying latent properties.

Consider an entity-level state representation at time $t$, denoted by

$$x_t \in \mathbb{R}^d,$$

which is derived from all events and attributes observed up to that time. The mapping from the event history to $x_t$ can be characterized by a function that may include transformations, embeddings, and temporal aggregations. In many implementations, this mapping is parameterized and learned, effectively serving as a representation learning layer that compresses high-dimensional interaction histories into tractable feature spaces [21]. When such mappings are learned jointly with downstream models, care is needed to avoid leakage between training and evaluation windows, particularly in the presence of temporal drift and intervention effects.

In a more explicit formulation, identity resolution can be treated as computation over a graph of identifiers. Let nodes represent identifiers and edges represent probabilistic matches. The resolution process induces a partition of nodes into clusters, each corresponding to an inferred entity. The Customer 360 state is then defined at the cluster level [22]. If $G$ denotes the identifier graph, and $C$ a cluster, one may define

$$x_{C,t} = f\big(\{e \in E : e \text{ linked to } C, e.\tau \le t\}\big),$$

where $E$ is the set of events and $\tau$ their timestamps, and $f$ is an aggregation function that remains bounded and stable under incremental updates. The stability requirement is important: small variations in linkage probabilities should not cause disproportionate shifts in state, since next-best-action policies will operate on these representations continuously.

Data governance and schema evolution further shape the state. New attributes, channels, and products are introduced over time; regulations impose constraints on data retention and usage; consent choices vary across individuals and epochs [23]. A neutral but explicit treatment regards the Customer 360 store as a versioned state space, where for each time point, only attributes consistent with applicable constraints are injected into $x_t$. Let $m_t$ denote a binary mask encoding allowed attributes at time $t$, then an operational representation can be written as

$$\tilde{x}_t = m_t \odot x_t,$$

where $\odot$ is the elementwise product. All subsequent modeling and decisioning must be defined in terms of $\tilde{x}_t$ rather than $x_t$, as the latter may include conceptually unavailable features. This distinction is relevant when models trained on historical data are deployed into regimes with different masks, potentially altering their effective input distributions.

Another consideration arises from the multi-resolution nature of Customer 360. Interactions occur at event level, while many actions are defined at session, account, household, or segment level [24]. Consistency conditions link these resolutions. A simple constraint is that certain aggregated metrics computed from event-level records should match the values stored in profile-level features within tolerance, given streaming latencies and update schedules. Inconsistencies between aggregates and event streams can signal ingestion issues, backfills, or transformation errors. Real-time analytics, discussed later, may monitor such invariants, but from the perspective of data foundations they motivate explicit functional relations between granular data and state representation.

Overall, the Customer 360 can be viewed as a dynamic, masked, and partially latent state space, updated via a streaming pipeline subject to operational and regulatory constraints [25]. This perspective avoids assuming a perfectly complete or static view and instead accommodates uncertainty, missingness, and evolution directly in the representation upon which next-best-action and real-time analytics are constructed.

## Stability, Governance, and Risk Management in Integrated Decisioning

The integration of Customer 360 architectures, next-best-action decision policies, and real-time analytics induces a coupled dynamical system whose behavior depends on data flows, model updates, operational constraints, and monitoring mechanisms. This coupling can be examined in terms of stability, governance, and risk management without presupposing that the resulting configuration is inherently optimal. Instead, the objective is to describe conditions under which such a system exhibits well-behaved trajectories for key metrics, retains interpretability, and remains responsive to deviations or failures. In this perspective, stability refers to boundedness and predictability of decision and outcome processes under perturbations; governance captures the structures by which policies are defined, audited, and evolved; and risk management concerns the identification and mitigation of failure modes originating from data quality, model drift, feedback loops, and constraint violations. [26]

A high-level abstraction views the system as iterating over three transformations at decision time scales: state construction from the Customer 360 pipeline, action selection from the next-best-action policy, and metric estimation from real-time analytics. Let $s_t$ denote the operational state at time $t$, $a_t$ the selected action, and $\hat{m}_t$ the vector of monitored metrics. The joint evolution can be summarized as a mapping from $(s_t, \theta_t)$ to $(s_{t+1}, \theta_{t+1})$, where $\theta_t$ are the parameters governing the policy and possibly some components of the representation. Under online or frequent retraining, $\theta_t$ evolves according to observed discrepancies between target values and realized outcomes. When these updates interact with delayed, noisy, or biased feedback, the risk arises that the closed-loop system oscillates, drifts, or converges to undesirable equilibria. Stability analysis thus becomes an integral component of governance, rather than a purely theoretical exercise.

A simplified representation of parameter updates for a next-best-action model under streaming feedback is given by an iteration such as

$$\theta_{t+1} = \theta_t - \alpha_t \nabla_\theta \ell_t,$$

where $\alpha_t$ is a learning rate and $\ell_t$ is a loss derived from real-time or recent outcomes. While this form is generic, the effective loss surface depends on how the Customer 360 state is constructed and how metrics are estimated, both of which may be changing [27]. If the learning rate schedule, metric definitions, or feature transformations are modified without explicit compatibility checks, the iterative process may cease to approximate descent on a stable objective. In practice, this can manifest as oscillating contact strategies, unstable score distributions, or abrupt shifts in which segments are prioritized. Governance processes that require controlled changes and explicit documentation of metric and feature evolution help ensure that updates to $\theta_t$ remain interpretable and diagnostically tractable.

Customer 360 pipelines themselves introduce dynamic elements into stability considerations. Since $s_t$ is derived from streaming events, identity resolution, and feature transformations, perturbations in upstream systems propagate into state space. Small changes in match thresholds, schema mappings, or late arrival handling rules can lead to discrete shifts in representations for certain entities. Let $\Delta s_t$ denote the variation in state due to such changes at time $t$. Even when $\Delta s_t$ is localized, the policy $\pi_{\theta_t}$ may amplify these variations if decision boundaries are sharp or heavily reliant on specific features. A basic robustness requirement is that the induced change in action distributions remains bounded, which can be expressed as a condition such as [28]

$$\sup_t \mathbb{E}\left[\|\pi_{\theta_t}(\cdot \mid s_t + \Delta s_t) - \pi_{\theta_t}(\cdot \mid s_t)\|_1\right] < \kappa,$$

for some tolerable bound $\kappa$. While this condition is not enforced directly in most production setups, it suggests regularization and calibration practices that reduce sensitivity of decisions to small representation perturbations. Governance measures that review proposed feature changes against such sensitivity analyses can reduce the risk of unintended action volatility.

Real-time analytics contribute both to stability and to potential instability [29]. On one hand, they enable rapid detection of anomalies, drift, and constraint violations; on the other, aggressive or uncalibrated adaptation rules driven by short-term fluctuations can introduce feedback loops. Suppose a performance metric estimate $\hat{m}_t$ controls a policy parameter $\beta_t$, which in turn influences future actions and thereby future values of $\hat{m}_t$. An adaptation rule of the form

$$\beta_{t+1} = \beta_t + \eta(\hat{m}_t - \bar{m}),$$

where $\bar{m}$ is a target value and $\eta$ a step size, defines a feedback mechanism. If $\eta$ is large relative to the noise and delay structure of $\hat{m}_t$, the closed-loop may oscillate. To mitigate this, real-time analytics can expose confidence intervals, effective sample sizes, and delay distributions, and adaptation rules can be constrained by bounds on $\eta$ or by incorporating smoothing and hysteresis. Governance frameworks can formalize acceptable ranges for these parameters and require simulations or shadow tests before deploying new adaptation logic.

Risk management in integrated decisioning systems must consider several classes of failure modes arising from the interaction of Customer 360, next-best-action policies, and analytics [30]. Data risks include missing or corrupted event streams, identity graph misconfigurations, and inconsistent masks for consent or eligibility. Model risks include overfitting to transient patterns, neglected confounders in uplift or bandit models, and insufficient coverage of state-action space for off-policy evaluation. Feedback risks involve attribution errors when multiple actions, channels, or external factors jointly influence outcomes. Operational risks relate to capacity misestimation, misalignment between configured constraints and implemented checks, and failures in streaming or model-serving infrastructure. A neutral treatment does not assume that such risks are eliminated; instead, it emphasizes explicit identification, monitoring, and containment strategies. [31]

One way to characterize systemic risk is through stress scenarios applied at the level of the integrated system rather than isolated components. These scenarios involve perturbing inputs, delaying feedback, or modifying action sets and observing the resulting trajectories of key metrics and constraint indicators. Mathematically, let $\omega$ index a family of perturbation scenarios, each defining altered dynamics for $F$, $\pi_\theta$, or measurement processes. For each $\omega$, the system induces a trajectory of metrics $\{\hat{m}_t^{(\omega)}\}$. Governance processes can define acceptability regions such that, for a set of scenarios deemed plausible, the trajectories remain within specified bounds. For example, for a moderate delay exaggeration scenario one may require that the deviation

$$\sup_t \|\hat{m}_t^{(\omega)} - \hat{m}_t^{(0)}\|$$

stays below a limit for primary metrics, where $(0)$ denotes baseline [32]. This formulation does not guarantee robustness in all conditions, but it provides a structured method to assess sensitivity before changes are promoted to production.

Another dimension is compliance with explicit constraints such as contact policies, fairness guidelines, or segment-level quotas. These constraints can be represented as inequalities on aggregates over actions and states. If $q_{g,t}$ denotes the fraction of actions allocated to group $g$ at time $t$, then a fairness or coverage rule might specify

$$|q_{g,t} - q_g^*| \le \delta,$$

for a target $q_g^*$ and tolerance $\delta$. Real-time analytics estimate $q_{g,t}$, while the policy and arbitration logic determine its evolution. Violations can arise not only from modeling choices but also from shifts in group composition, data quality issues, or technical failures [33]. Governance mechanisms may therefore require that any policy updates be accompanied by projected impacts on such constraints, using replay or simulation. This integrated view treats constraints as first-class objects in both design and monitoring, rather than as afterthoughts to be checked intermittently.

Interpretability plays a role in both stability and governance. When policies are high-dimensional or involve complex representation learning, explaining individual recommendations or aggregate patterns becomes challenging. However, some level of interpretability is needed for model risk assessment, incident investigations, and decision audits [34]. This can be supported by maintaining auxiliary models or diagnostics that approximate the main policy with simpler functions on a stable subset of features. Although such approximations may not be used for decisioning, they can reveal whether changes in behavior arise primarily from shifts in a few interpretable directions or from opaque interactions. Mathematically, if $\pi_\theta$ is complex, an auxiliary mapping $\tilde{\pi}_\varphi$ can be fitted to minimize discrepancies over sampled states, subject to structural constraints that enhance interpretability. Monitoring the divergence between $\pi_\theta$ and $\tilde{\pi}_\varphi$ over time can signal when the effective behavior of the system moves beyond regions that have been qualitatively understood.

The interaction between governance processes and real-time analytics also affects how quickly issues are detected and addressed. Alerting thresholds chosen solely on point estimates of metrics can be either hypersensitive or insensitive, depending on noise and traffic levels. Incorporating variance estimates, distributional shifts, and contextual indicators into alerting logic introduces additional complexity, but it can reduce unnecessary interventions while preserving sensitivity to structurally significant deviations [35]. For instance, if a response rate metric is estimated with an associated uncertainty measure, governance rules can require that alerts be triggered only when deviations exceed combined statistical and practical significance thresholds. This practice aligns monitoring with the probabilistic nature of streaming data and avoids overreaction to random fluctuations.

Capacity and resource constraints create another feedback path. Next-best-action models that aggressively select actions requiring human or system resources can saturate those capacities, leading to degraded service levels and altered realized outcomes. Let $C_t$ denote available capacity for a class of actions at time $t$, and $D_t$ the demand induced by the policy. A simple feasibility condition is $D_t \leq C_t$, but in practice both are stochastic and time-varying. Real-time analytics estimate utilization patterns, while policies may include terms that penalize actions likely to exceed capacities [36]. Adaptation logic must ensure that updates intended to improve short-term metrics do not systematically generate overloads that compromise overall performance. Modeling this interaction explicitly, even in simplified form, encourages conservative and transparent handling of capacity-sensitive actions.

In aggregate, a stability, governance, and risk management perspective on integrated decisioning systems underscores the importance of treating Customer 360, next-best-action models, and real-time analytics as coupled components governed by coherent rules. Rather than assuming uniform reliability or optimality, the system is viewed as an evolving construct whose behavior can be characterized by its reactions to perturbations, its adherence to declared constraints, and its capacity for controlled adaptation. Mathematical formulations of feedback, robustness, and constraints offer a vocabulary for these assessments, while governance structures define how such formulations are applied in practice [37]. This approach supports incremental refinement and informed adjustments over time, with the understanding that integrated data-driven architectures are subject to both statistical and operational uncertainties that warrant systematic, rather than ad hoc, management.

## Next-Best-Action Models as Decision Policies

Next-best-action models can be framed as parameterized policies that map the current Customer 360 state and contextual variables to a probability distribution over candidate actions. Let $\mathcal{A}$ denote the action set, where an action may represent an offer, a message, a treatment, a service intervention, or a decision not to intervene. At decision time $t$, given state $\tilde{x}_t$ and possibly exogenous context $c_t$, the policy outputs a distribution

$$\pi_\theta(a \mid \tilde{x}_t, c_t)$$

for $a \in \mathcal{A}$, where $\theta$ are parameters. This policy may be stochastic to facilitate exploration, handle uncertainty, and represent indifference across near-equivalent choices. From an operational standpoint, $\pi_\theta$ is often factorized or constrained to satisfy action eligibility rules, contact policies, and resource limits.

In the simplest case, next-best-action selection is posed as a one-step optimization of an expected reward function [38]. Suppose that for each action $a$, a model estimates the conditional expectation

$$\mu_\phi(a, \tilde{x}_t, c_t)$$

under parameters $\phi$. Then an immediate reward policy selects actions that maximize this estimate, subject to constraints. However, this approach can be sensitive to confounding in historical logs, changes in environment, and unobserved heterogeneity. Furthermore, it does not account for delayed or long-term consequences of actions [39]. To mitigate some of these limitations, counterfactual modeling, uplift estimation, and contextual bandit methods can be employed, but each approach rests on assumptions about assignment mechanisms and outcome models that may or may not hold.

A more general formulation treats the interaction process as a Markov decision process where the state includes the Customer 360 representation, the action is chosen by the policy, and the environment produces a next state and reward. Denote state by $s_t$ and action by $a_t$. A discounted objective for a policy $\pi$ can be expressed as

$$J(\pi) = \mathbb{E}\Big[\sum_{t=0}^{T} \gamma^t r_t\Big],$$

with discount factor $\gamma \in (0, 1)$ and horizon $T$. The role of $\pi$ is to map states to action distributions [40]. In practice, state observability is partial, reward definitions are composite, and the environment is nonstationary, which limits the direct applicability of canonical reinforcement learning algorithms without adaptation to these constraints. However, viewing next-best-action through this lens helps structure the presence of sequential dependencies and long-term objectives.

In many environments, multiple objectives must be balanced, such as revenue, satisfaction, risk, and fairness metrics. One way to encode this is through a vector-valued reward function, combined into a scalar via a weight vector that can be adjusted. For example, let [41]

$$r_t = w^\top z_t,$$

where $z_t$ contains component metrics and $w$ is specified or tuned. The neutrality of this construction underscores that the resulting policy is contingent on chosen weights rather than inherently optimal. Constraints can be handled using Lagrangian or projection-based methods, where policies are adapted so that empirical metrics remain within acceptable intervals.

Another operational dimension is eligibility and arbitration among multiple models. Different domains or channels may propose candidate actions simultaneously [42]. An arbitration layer can be represented as a function

$$g(a_t^{(1)}, \ldots, a_t^{(K)}, \tilde{x}_t, c_t),$$

where each $a_t^{(k)}$ is the suggestion from model $k$. The function $g$ selects or combines these suggestions while enforcing global rules, capacities, and pacing constraints. This level of indirection is useful in large organizations, where separate teams maintain specialized models; representing arbitration as a formal mapping avoids implicit precedence rules that are difficult to audit or adjust.

Learning next-best-action policies from logged data requires off-policy evaluation techniques. Propensity scores, importance weighting, and doubly robust estimators are often proposed, yet they can exhibit high variance or bias under small propensities or model misspecification [43]. Real-time analytics integrated with experimentation can help maintain coverage over relevant portions of the state-action space, thereby stabilizing estimates, but only if exploration is managed within acceptable operational constraints. Thus, the definition of next-best-action models as policies is inseparable from the mechanisms used to evaluate and update them under streaming data conditions.

## Real-Time Analytics and Streaming Feedback

Real-time analytics in this context are not limited to descriptive reporting; they represent the mechanism linking observed outcomes back into the decision-making loop. Given a stream of events, each associated with customer identifiers, timestamps, actions, and outcomes, the analytics layer maintains short-latency aggregates, performance indicators, anomaly signals, and inputs for online or nearline learning processes. The technical characteristics of this layer, such as end-to-end latency, windowing semantics, and fault tolerance, directly influence the effective observability and timeliness of feedback available to next-best-action models. [44]

Let the incoming event stream be denoted by $\{e_j\}$, each event characterized minimally by an identifier, a timestamp, and content fields. For a given metric function $h$ and window $W_t$, the streaming estimate at time $t$ can be written as

$$\hat{m}_t = \frac{1}{|W_t|} \sum_{e_j \in W_t} h(e_j),$$

with the understanding that $W_t$ is defined by event time or processing time, subject to watermarking and late arrival policies. The precise semantics of $W_t$ influence how rapidly shifts in behavior are detected and how robust the estimates are to skewed or delayed data. If next-best-action policies adapt parameters using such metrics, then misalignment between the conceptual observation window and the implemented window can lead to biased updates.

Real-time analytics also play a role in monitoring operational constraints, such as maximum contact frequencies, channel capacities, and regulatory or contractual limits. Let $N_{a,t}$ denote the number of times action $a$ has been applied within a specified period up to time $t$. Streaming counters can be maintained such that

$$N_{a,t+1} = N_{a,t} + \delta_a(e_{t+1}),$$

where $\delta_a$ increments when an event corresponds to action $a$. When $N_{a,t}$ approaches a configured capacity, either globally or for specific segments, the arbitration logic adjusts action probabilities to remain within boundaries. This connection between accounting metrics and decision policies is essential for ensuring that local, state-based optimization does not inadvertently violate system-level constraints.

From a modeling perspective, real-time analytics supply covariates and calibration targets for online learning [45]. For instance, short-term response rates to actions, conditional on coarse state partitions, can be estimated continuously, providing empirical calibration for outcome models. Let $R_t(a)$ represent the empirical response rate estimate for action $a$ in a recent time window. A simple update can be represented as

$$R_{t+1}(a) = (1 - \eta) R_t(a) + \eta y_{t+1}(a),$$

where $y_{t+1}(a)$ is the observed outcome indicator and $\eta$ is a step size. Although elementary, such recursions embody the basic trade-off between stability and responsiveness that more complex algorithms also face.

An additional function of the analytics layer is anomaly detection. Unexpected shifts in distributions of actions, outcomes, or state features may indicate upstream issues or environment changes [46]. Rather than relying solely on batch diagnostics, streaming statistics can feed simple decision rules or more advanced detectors. For example, deviations between predicted and realized aggregate outcomes can be tracked via residual processes. If $\hat{y}_t$ is the predicted outcome aggregate and $\tilde{y}_t$ the realized value, one may define the residual

$$\epsilon_t = \tilde{y}_t - \hat{y}_t,$$

and monitor its distribution over rolling windows. Persistent deviations beyond expected ranges can trigger investigations or automated mitigations, such as reverting to more conservative policies. [47]

Real-time analytics are implemented in distributed streaming frameworks where exactly-once guarantees, state snapshots, and backpressure handling define practical boundaries. The Customer 360 and next-best-action layers must be aligned with these properties. If decisioning relies on state that is updated on different cadence or semantics than the analytics used for evaluation, discrepancies may arise. Therefore, both conceptual and technical coherence between these layers is required to maintain interpretable and stable behavior.

## Integrated Mathematical Formulation

To reason about the joint behavior of Customer 360, next-best-action policies, and real-time analytics, an integrated mathematical formulation can be constructed [48]. Consider a discrete-time process where at each step $t$ the system observes or infers a state $s_t$, selects an action $a_t$ according to policy $\pi_\theta$, and receives outcome information that is incorporated into both the Customer 360 state and the analytics layer. The state $s_t$ may be a function of underlying latent variables and historical events; for tractability, the model treats $s_t$ as the operational representation $\tilde{x}_t$ augmented with additional contextual variables.

| Aspect | Definition | Representation | Constraint |
|---|---|---|---|
| Customer 360 State | Aggregated customer representation | $x_t \in \mathbb{R}^d$ | Latent, probabilistic |
| Identity Resolution | Mapping identifiers to entities | Graph-based clusters $C$ | Probabilistic linkage |
| Temporal Masking | Attribute filtering over time | $\tilde{x}_t = m_t \odot x_t$ | Consent & regulation |
| Multi-resolution | Event-to-entity consistency | Aggregation invariants | Streaming latency |

**Table 1** Key Components of Data Foundations for Customer 360

| Concept | Notation | Description | Purpose |
|---|---|---|---|
| Entity State | $x_{C,t}$ | Aggregated cluster features | Entity-level modeling |
| Event Stream | $E$ | Transactional inputs with timestamps | Temporal data foundation |
| Aggregation Function | $f(\cdot)$ | Mapping of events to state | Stability over updates |
| Mask Vector | $m_t$ | Regulatory attribute filter | Versioned state management |

**Table 2** Formal Elements of the Customer 360 Representation

Formally, define the evolution of the state as

$$s_{t+1} = F(s_t, a_t, o_{t+1}),$$

where $o_{t+1}$ represents observed feedback and exogenous signals between $t$ and $t+1$, and $F$ is a possibly stochastic update operator representing the Customer 360 ingestion and transformation pipeline. The policy is given by

$$a_t \sim \pi_\theta(\cdot \mid s_t),$$

and the analytics layer maintains estimators $\hat{m}_t$ of performance functionals $m$. The combined system can be viewed as a controlled stochastic process parameterized by $\theta$ and by the configuration of $F$ and the estimators.

The design problem is then to select $\theta$ and the structural elements of $F$ and the analytics layer such that certain criteria are met. These criteria may include convergence of performance metrics to stable ranges, satisfaction of constraints, robustness to moderate distribution shifts, and interpretability conditions [49]. A generic objective can be expressed as

$$\max_\theta \Phi(\{\hat{m}_t\}),$$

where $\Phi$ is a functional of the time series of estimated metrics. For example, $\Phi$ might capture long-run averages of outcomes, penalized by constraint violations or volatility. Unlike classical formulations that optimize a known analytical form, here $\Phi$ is typically approximated from data and depends on feedback loops among components.

A particular concern in integrated systems is off-policy evaluation under logging policies that evolve over time [50]. Let $b_t(a \mid s)$ denote the behavior policy active at time $t$. When estimating the value of a target policy $\pi$ from logs generated by $b_t$, importance sampling ratios of the form

$$w_t = \frac{\pi(a_t \mid s_t)}{b_t(a_t \mid s_t)}$$

arise. However, if $b_t$ is itself adaptive and only partially recorded, these ratios may be inaccurately specified, leading to unstable estimates. One possible mitigation is to constrain both $b_t$ and $\pi$ to remain within a bounded divergence. For instance, enforcing that

$$D_{\mathrm{KL}}(\pi(\cdot \mid s) \,\|\, b_t(\cdot \mid s)) \le \delta$$

for some small $\delta$ can limit variance in off-policy estimators. This constraint, or analogous ones based on other divergences, can be integrated into the policy optimization process and monitored by real-time analytics, thus coupling decision and evaluation design.

Customer 360 representation learning can also be embedded into this formulation [51]. Suppose that the state $s_t$ is generated by an encoder function with parameters $\psi$, acting on recent event sequences. The joint optimization problem then involves both $\theta$ and $\psi$. One may consider objectives that combine prediction accuracy for outcomes and auxiliary tasks, regularization for stability, and compatibility with policy evaluation constraints. A stylized objective is

$$\min_{\theta,\psi} L(\theta, \psi),$$

with $L$ including empirical loss terms and penalties representing constraints or robustness requirements. Although such joint optimization is computationally intensive, it captures the interaction between how customers are represented and how actions are chosen. [52]

Latency constraints enter mathematically through restrictions on which events can be used in $s_t$ at decision time. Let $\Delta$ denote maximal acceptable decision latency. If raw events have arrival delay distribution and processing time, only a subset of them are available before $\Delta$. Hence, define $H_t^{(\Delta)}$ as the history truncated to events effectively observable within $\Delta$, and construct

$$s_t = G(H_t^{(\Delta)}),$$

for some function $G$. Any modeling strategy that implicitly uses future or delayed information in training must reconcile this with the feasible online construction of $s_t$. Otherwise, performance estimates will be optimistic relative to deployable configurations. Real-time analytics that accurately track delays and arrival patterns provide the empirical basis for calibrating $G$ and

| Component | Symbol | Function | Governance Role |
|---|---|---|---|
| State Construction | $s_t$ | Derived from 360 pipeline | Data stability |
| Action Policy | $\pi_{\theta_t}$ | Maps state to actions | Controlled adaptation |
| Metric Estimation | $\hat{m}_t$ | Real-time feedback | Monitoring |
| Parameter Update | $\theta_{t+1}$ | $\theta_t - \alpha_t \nabla_\theta \ell_t$ | Learning governance |

**Table 3** Dynamics in Integrated Decisioning Systems

| Risk Type | Source | Manifestation | Mitigation |
|---|---|---|---|
| Data Risk | Stream delays, ID errors | Missing features | Monitoring, validation |
| Model Risk | Drift, bias | Unstable predictions | Retraining governance |
| Feedback Risk | Action attribution | Oscillation | Causal modeling |
| Operational Risk | Capacity misalignment | System overloads | Simulation checks |

**Table 4** Risk Classes in Integrated Decisioning

for validating that deployed pipelines conform to the assumptions made during modeling. [53]

Finally, fairness and other distributional constraints can be expressed as conditions on outcome distributions across subgroups defined over components of $s_t$. For example, one may require that the disparity in a selected metric between groups does not exceed a given bound. Let $m_g$ denote a metric estimate for group $g$, then a simple constraint is

$$|m_{g_1} - m_{g_2}| \le \epsilon$$

for specified groups and tolerance $\epsilon$. Enforcement involves both monitoring these quantities in real time and incorporating them into the policy updating mechanism, potentially via constrained optimization. This further couples the analytics layer, which estimates $m_g$, with policy design, highlighting the importance of integrated formulation.

## Experimental Evaluation and Implementation Considerations

A practical assessment of an integrated Customer 360, next-best-action, and real-time analytics system can be structured around several dimensions: quality of state representation, decision performance under multiple objectives, robustness to drift and delays, and operational reliability. While experimental protocols vary across organizations and domains, a common pattern involves a combination of historical replay, controlled online experiments, and ongoing monitoring of key indicators [54]. The interpretation of results is inherently conditional on the chosen metrics, traffic allocation strategies, and environmental conditions during the evaluation period.

Historical replay methods construct simulated decision trajectories by applying candidate policies to recorded event streams. Given logs containing states, actions, and outcomes under some behavior policy, one can emulate how alternative policies would have acted, subject to logged constraints. However, since actions in the logs were not chosen by the candidate policies, importance weighting or model-based estimators are needed to correct for distribution mismatches. As mentioned earlier, the variance and bias properties of such estimators depend strongly on how behavior policies evolved and how thoroughly they were recorded [55]. When behavior policies are under explicit control of the same system, maintaining structured exploration and logging can support more stable evaluation.

Controlled online experiments, such as randomized policy assignments across segments or time, provide direct evidence about relative performance and side effects. The real-time analytics layer is essential for implementing such experiments, ensuring correct randomization, and recording outcomes at appropriate granularities. In these experiments, the Customer 360 representation influences how treatment groups are stratified and how heterogeneity is analyzed [56]. If the representation captures relevant dimensions of customer behavior and risk, experimental results can be decomposed along these axes to identify stable patterns and interaction effects. Conversely, if the representation is unstable or excessively complex, interpretation of experimental results becomes more difficult.

Robustness to distribution shifts and data quality issues can be studied by introducing controlled perturbations in offline or shadow environments and observing their impact on state construction, policy outputs, and metrics. For example, simulating increased event delays or partial data loss in upstream feeds reveals how sensitive the Customer 360 and next-best-action layers are to such disruptions. Real-time analytics that expose discrepancies between expected and observed distributions, as described earlier, are part of the mitigation approach [57]. A neutral evaluation does not assume immunity to such issues but characterizes the range of conditions under which system behavior remains within acceptable limits.

Implementation considerations extend to schema design, storage models, and deployment topologies. Columnar storage and key-value access patterns are often combined to support both analytical scans and low-latency profile reads. Streaming pipelines transform raw events into incremental updates to the Customer 360 state. Decision services hosting next-best-action models must handle throughput peaks while respecting latency budgets [58]. Real-time analytics systems consume the same or parallel streams to avoid double counting and inconsistencies. Co-location or logical alignment of these components reduces the risk of divergence between the state used for decisions and the state used for evaluation.

From a modeling lifecycle perspective, governance mecha-

| Policy Component | Notation | Function | Constraint |
|---|---|---|---|
| Action Distribution | $\pi_\theta(a \mid \tilde{x}_t, c_t)$ | Stochastic policy | Eligibility rules |
| Reward Estimator | $\mu_\phi(a, \tilde{x}_t, c_t)$ | Expected outcome | Bias correction |
| Sequential Objective | $J(\pi)$ | Discounted reward sum | Horizon $T$, $\gamma$ |
| Arbitration Layer | $g(a_t^{(1)}, \ldots, a_t^{(K)})$ | Model coordination | Capacity, pacing |

**Table 5** Next-Best-Action Policy Components

| Objective | Variable | Relation | Interpretation |
|---|---|---|---|
| Fairness | $\lvert q_{g,t} - q_g^* \rvert \leq \delta$ | Group coverage | Equity constraint |
| Stability | $\sup_t \mathbb{E}\lVert \pi_{\theta_t}(s_t + \Delta s_t) - \pi_{\theta_t}(s_t) \rVert_1 < \kappa$ | Bounded response | Robustness |
| Adaptation | $\beta_{t+1} = \beta_t + \eta(\hat{m}_t - \bar{m})$ | Feedback regulation | Avoid oscillation |
| Capacity | $D_t \leq C_t$ | Resource feasibility | Load control |

**Table 6** Governance and Constraint Formulations

nisms define how new models are introduced, modified, or retired. Release processes may involve shadow deployment, where a candidate policy runs in parallel and its decisions are logged but not executed, enabling comparative analysis without customer impact. Real-time analytics support such patterns by keeping aligned metrics for multiple synthetic policies [59]. These procedures reduce the risk of abrupt degradations following model changes, especially when next-best-action policies are high leverage elements in customer interactions.

Finally, monitoring and observability are integral. Meaningful metrics include response distributions, contact frequencies, capacity utilization, model score distributions, calibration curves, and constraint adherence. Deviations or saturations in these metrics can indicate conditions such as overfitting to short-term responses, under-utilization of eligible actions, or systematic exclusion of certain segments [60]. The measurement of these phenomena relies on both the Customer 360 representation, which defines segment and feature spaces, and the analytics infrastructure, which computes and exposes aggregates with low latency. Therefore, implementation choices that facilitate transparent and consistent monitoring contribute to maintaining predictable and analyzable system behavior.

## Conclusion

This paper has described a data-driven approach that integrates Customer 360 representations, next-best-action models, and real-time analytics into a unified decisioning framework without assuming idealized data, stable objectives, or static environments. The central elements of this formulation are the interpretation of Customer 360 as a dynamic and partially latent state space subject to governance, latency, and identity resolution constraints; the framing of next-best-action as a policy defined over that state under multiple, and at times competing, objectives; and the treatment of real-time analytics as a structured feedback mechanism that supplies measurements, diagnostics, and constraint indicators back into both representation learning and policy adaptation. Viewed together, these components form a coupled system whose behavior depends on how data is ingested and transformed, how decisions are parameterized and constrained, and how outcomes are observed and aggregated [61]. The resulting perspective underscores the value of explicit assumptions, coherent definitions of state, policy, and feedback operators, and alignment between conceptual models and deployed infrastructures, rather than relying on loosely connected descriptions of isolated capabilities.

The mathematical constructs discussed throughout, including controlled stochastic processes, constrained and stochastic policies, streaming estimators, and robustness considerations, are intended to render this integrated system analyzable without implying guarantees beyond what the data, architecture, and operating conditions can support. Under realistic conditions, responses to actions are delayed, outcomes are only partially observed, and both customer populations and applicable constraints evolve over time. In such settings, formal guarantees become conditional statements that depend on aspects such as the adequacy of the state representation, the calibration of models under deployment constraints, and the stability of feedback loops that connect behavior, measurement, and adaptation. It is therefore appropriate that performance claims for integrated Customer 360 and next-best-action architectures be interpreted with respect to specific configurations, objective definitions, evaluation windows, and data quality regimes, rather than as universal properties of a methodological template [62]. The unified treatment proposed here provides a language for articulating those conditions more transparently.

A consistent theme in the analysis is that observable behavior in these systems emerges from interactions among elements that are often specified and owned by different technical and organizational stakeholders. Customer 360 pipelines encode decisions about identity resolution, attribution windows, consent enforcement, and feature construction. Next-best-action policies encode modeling choices about objectives, exploration, eligibility, and arbitration. Real-time analytics encode choices about time windows, aggregation rules, baselines, and alerting logic [63]. When these elements are defined independently, implicit assumptions can accumulate at their interfaces, introducing silent dependencies that complicate diagnosis when behaviors drift or constraints are breached. By embedding all three within a shared analytical structure, it becomes possible to trace how changes in one layer affect the others, to design compatibility conditions in advance, and to distinguish between deviations attributable to environment shifts and those arising from internal configuration changes. This does not eliminate complexity, but it narrows the gap between informal reasoning about the system and the mechanisms that actually govern its operation.

The discussion of stability, governance, and risk management further reinforces the view that integrated decisioning systems should be approached as evolving artifacts rather than fixed solutions. Because decision policies shape the distribution of observed data, which in turn shapes subsequent model updates and monitoring thresholds, these systems naturally form feedback loops [64]. Without deliberate controls on update frequency, learning rates, feature revisions, and constraint logic, such loops may exhibit oscillatory behavior, rapid shifts in segment treatment, or sensitivity to short-term fluctuations. Similarly, modest changes in Customer 360 construction rules, such as adjustments to identifier matching or handling of delayed events, can materially alter the state space on which policies operate. Real-time analytics can mitigate some of these risks by revealing anomalies, drift, or constraint violations at short latency, but they can also introduce additional feedback channels when their outputs drive automated policy adjustments. A structured, data-driven approach treats these interactions explicitly, using modeling and simulation where feasible, and employing governance processes that require that modifications to one part of the system be assessed in the context of its couplings to others. [65]

The neutral stance adopted here is intentional. Integrated architectures combining Customer 360, next-best-action, and real-time analytics offer a broad range of potential configurations, and their realized properties depend on concrete decisions about objectives, constraints, and technical implementation. Rather than characterizing such systems as inherently superior, the analysis focuses on how they can be expressed in terms that admit scrutiny: how states are defined and updated, how actions are selected and evaluated, how metrics are computed and acted upon, and how constraints are encoded and enforced. From this vantage point, questions about fairness, regulatory compliance, contact governance, and capacity management become questions about the structure of policies and the design of feedback operators, not external considerations applied post hoc. Similarly, questions about robustness to drift, resilience to data quality issues, or interpretability of recommendations can be framed within the same set of constructs, allowing heterogeneous requirements to be assessed using a shared vocabulary of states, mappings, and trajectories. [66]

Future extensions of this work may explore richer models of sequential decision-making in which long-term outcomes are modeled more explicitly, along with methods for reconciling short-term response optimization with multi-period objectives defined at customer, portfolio, or system levels. More granular treatments of uncertainty and causal structure could refine how uplift models, contextual bandits, and approximate reinforcement learning methods are employed within constrained, high-volume environments. Additional analysis of multi-channel patterns, including asynchronous and cross-device sequences, could improve understanding of how actions in one channel influence state evolution and opportunity sets in others. Further, detailed modeling of hybrid humanalgorithm workflows may clarify how recommendations interact with human discretion in assisted channels, and how real-time analytics can separate policy effects from behavioral adaptations of agents and customers. These extensions would not seek to change the fundamental framing, but to populate it with more detailed mechanisms and empirical characterizations that align with specific domains and regulatory contexts. [67]

At the operational level, the framework outlined here may be used to structure questions about implementation choices without prescribing a particular technology stack. For example, decisions about whether to centralize or federate profile storage, how to partition features between offline and online computation, or how to schedule retraining and deployment of decision models can be evaluated in terms of their impact on state consistency, latency, estimation quality, and stability of feedback. Similarly, choices about experimentation strategies, such as policy-level randomized trials or shadow policies evaluated via off-policy estimation, can be examined using the same constructs that describe production decisioning. Real-time analytics, in turn, can be configured to monitor alignment between intended policies and realized behavior, using metrics that are derived from and consistent with the formal definitions of state and action spaces, rather than from a separate and potentially misaligned reporting layer. In this way, the integrated formulation acts less as a blueprint and more as a reference model against which concrete implementations can be mapped and interrogated. [68]

The analysis suggests that treating Customer 360 systems, next-best-action models, and real-time analytics as components of a single, explicitly defined decisioning process supports more transparent reasoning about their behavior, limitations, and configuration options. It encourages moving from informal narratives about personalization, orchestration, or intelligence toward specifications that can be checked against data flows, code, and observed outcomes. It does not assume that every organization or context requires maximal automation or complexity, nor that integrated architectures invariably dominate simpler alternatives. Instead, it provides tools for determining where the additional coordination and modeling effort is warranted, how to document and test assumptions, and how to interpret results in light of the intertwined nature of data, models, and feedback. Viewed in this manner, such systems are complex but analyzable, and their properties are emergent from design choices that can be surfaced, examined, and revised as objectives, constraints, and environments evolve. [69]

## References


[1] S. A. R. Khan and S. I. Sümer, “Modern perspectives in business applications,” 5 2020.

[2] O. Klisenko and E. S. Asensio, “Towards a maturity model for iot adoption by b2c companies,” *Applied Sciences*, vol. 12, pp. 982–982, 1 2022.

[3] S. Karim and E. Gide, “A study to analyse bangladeshi consumers e-commerce security and privacy satisfactions in small to mid-sized enterprises (smes),” *Global Journal of Computer Sciences: Theory and Research*, vol. 8, pp. 32–40, 4 2018.

[4] F. Ahmad, K. Mustafa, S. A. R. Hamid, K. F. Khawaja, S. Zada, S. Jamil, M. N. Qaisar, A. Vega-Muñoz, N. Contreras-Barraza, and N. Anwer, “Online customer experience leads to loyalty via customer engagement: Moderating role of value co-creation.,” *Frontiers in psychology*, vol. 13, pp. 897851–, 7 2022.

[5] L. Bruce, K. Bondy, R. Street, and H. Wilson, “Channel evolution: How new multichannel thinking can deliver com-

petitive advantage," *Journal of Direct, Data and Digital Marketing Practice*, vol. 10, pp. 329–335, 5 2009.

[6] M. Nilashi and O. Ibrahim, "A model for detecting customer level intentions to purchase in b2c websites using topsis and fuzzy logic rule-based system," *Arabian Journal for Science and Engineering*, vol. 39, pp. 1907–1922, 11 2013.

[7] A. D. Vittorio and D. G. V. Alessandro, "How lead generation can be the link between marketing and sales through the customer experience? a case study of talentia software italy," *European Scientific Journal, ESJ*, vol. 16, pp. 45–45, 12 2020.

[8] V. Nandavadekar, A. R. Wafgaonkar, and P. V. Kulkarni, "D-commerce- a way for business," *International Journal of Computer Applications*, vol. 1, pp. 60–68, 2 2010.

[9] W. T. W. Loon, S. K. Ariffin, and N. A. Wahid, "The influencing factors affecting consumers buying decision towards "made in malaysia" product: A proposed model," *International Journal of Education, Psychology and Counseling*, vol. 7, pp. 79–89, 9 2022.

[10] M. Aparicio, C. Costa, and R. Moises, "Gamification and reputation: key determinants of e-commerce usage and repurchase intention," *Heliyon*, vol. 7, pp. e06383–, 3 2021.

[11] S. N. Bipasha, "Promotional tools influencing service consumption of platform businesses: A study on dhaka city," *European Journal of Business and Management Research*, vol. 7, pp. 161–169, 11 2022.

[12] L. Mary and A. Antony, "Digital payment systems (dps) and its influence on impulsive buying behaviour of consumers," *International journal of health sciences*, pp. 2095–2102, 3 2022.

[13] N. A. Vences, J. Díaz-Campo, and D. F. G. Rosales, "Neuromarketing as an emotional connection tool between organizations and audiences in social networks. a theoretical review.," *Frontiers in psychology*, vol. 11, pp. 1787–, 7 2020.

[14] P. C. Addo, S. K. Akpatsa, P. Nukpe, A. A. Ohemeng, and N. B. Kulbo, "Digital analytics approach to understanding short video advertising in digital marketing," *Journal of Marketing Theory and Practice*, vol. 30, pp. 405–420, 4 2022.

[15] S. H. Elgazzar, "Investigating digital supply chain transformation drivers: An empirical study," , vol. 58, pp. 139–170, 9 2021.

[16] A. Ferreira and F. Antunes, "Essential functionalities for commercial internet presence: A portuguese study," *International Journal of E-Business Research*, vol. 11, pp. 56–83, 1 2015.

[17] S. Román, "Relational consequences of perceived deception in online shopping: The moderating roles of type of product, consumers attitude toward the internet and consumers demographics," *Journal of Business Ethics*, vol. 95, pp. 373–391, 2 2010.

[18] C. A. Anacleto, E. P. Paladini, and C. R. Vaz, "Customer requirements for the continuous improvement of product-service system: a conceptual approach," *Independent Journal of Management & Production*, vol. 9, pp. 526–544, 6 2018.

[19] N. T. Q. Nga and . H. Vân, "The impacts of customer satisfaction on loyalty of the youth (18-25 years old) towards e-commerce trading floors in ho chi minh city," *Science & Technology Development Journal - Economics - Law and Management*, vol. 4, pp. First–, 10 2020.

[20] Q. Zhou, C. Wang, and Y. Wang, "The change of supply chain in the transformation of traditional retail enterprises to new retail enterprises: A case study of chinese rt-mart," *BCP Business & Management*, vol. 16, pp. 139–146, 12 2021.

[21] L. Bitomsky, O. Bürger, B. Häckel, and J. Töppel, "Value of data meets it security  assessing it security risks in data-driven value chains," *Electronic Markets*, vol. 30, pp. 589–605, 2 2020.

[22] Y. Wang, M. Agyemang, and F. Jia, "Resource orchestration in supply chain service-based business model: The case of a cross-border e-commerce company," *Sustainability*, vol. 13, pp. 11820–, 10 2021.

[23] W. Rachbini, D. Anggraeni, and H. Wulanjani, "The influence of electronic service quality and electronic word of mouth (ewom) toward repurchase intention (study on e-commerce in indonesia)," *Jurnal Komunikasi: Malaysian Journal of Communication*, vol. 37, pp. 42–58, 3 2021.

[24] S. Sleep, J. Hulland, and R. A. Gooner, "The data hierarchy: factors influencing the adoption and implementation of data-driven decision making," *AMS Review*, vol. 9, pp. 230–248, 6 2019.

[25] A. Stancu, A. Filip, M. I. Roca, D. Ioni, R. Caplescu, A. Cânda, and L. D. Roca, "Value creation attributesclustering strategic options for romanian smes," *Sustainability*, vol. 12, pp. 7007–, 8 2020.

[26] X. Zhang, H. Liu, and P. Yao, "Research jungle on online consumer behaviour in the context of web 2.0: Traceability, frontiers and perspectives in the post-pandemic era," *Journal of Theoretical and Applied Electronic Commerce Research*, vol. 16, pp. 1740–1767, 6 2021.

[27] J. Becker, D. Beverungen, and R. Knackstedt, "The challenge of conceptual modeling for product-service systems status-quo and perspectives for reference models and modeling languages," *Information Systems and e-Business Management*, vol. 8, pp. 33–66, 2 2009.

[28] A. Islam, M. A. Khan, T. Ramayah, and M. M. Hossain, "The adoption of mobile commerce service among employed mobile phone users in bangladesh: Self-efficacy as a moderator," *International Business Research*, vol. 4, pp. 80–, 3 2011.

[29] G. Singh, "A research paper on g-mart: The e-commerce website," *International Journal for Research in Applied Science and Engineering Technology*, vol. 9, pp. 1949–1956, 11 2021.

[30] B. Ocicka and M. Raniewska, "In search of excellence in e-customer logistics service," *International Journal of Management and Economics*, vol. 49, pp. 135–155, 3 2016.

[31] Y. W. Park and P. Hong, "A research framework for sustainable digital innovation: Case studies of japanese firms," *Sustainability*, vol. 14, pp. 9218–9218, 7 2022.

[32] S. Karim and E. Gide, "Barriers to adopting e-commerce with small to mid-sized enterprises-smes in developed countries: an exploratory study in australia," *Global Journal of Information Technology: Emerging Technologies*, vol. 8, pp. 43–54, 8 2018.

[33] D. V. Arianna and D. G. V. Alessandro, "How crm performs inside sales activities in itc companies: A case study of fincons group, italy," *European Scientific Journal, ESJ*, vol. 12, pp. 312–312, 7 2016.

[34] 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.

[35] G. J. Heyns and P. J. Kilbourn, "Online shopping behaviour and service quality perceptions of young people in south africa: A covid-19 perspective," *Journal of Transport and*

*Supply Chain Management*, vol. 16, 8 2022.

[36] J. W. Lee and M. Kwag, "Corporate marketing strategy using social media: A case study of the ritz-carlton seoul," *The Journal of Asian Finance, Economics and Business*, vol. 4, pp. 79–86, 2 2017.

[37] A. K. Swain and R. Q. Cao, "Using sentiment analysis to improve supply chain intelligence," *Information Systems Frontiers*, vol. 21, pp. 469–484, 5 2017.

[38] A. P. C. Chan and E. K. Owusu, "Evolution of electronic procurement: Contemporary review of adoption and implementation strategies," *Buildings*, vol. 12, pp. 198–198, 2 2022.

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

[40] K. Fletcher, "Privacy: The achilles' heel of the new marketing," *Interactive Marketing*, vol. 3, pp. 141–153, 10 2001.

[41] A. Amatus and I. A. Gisip, "Effects of website appearance, security and electronic word-of-mouth (ewom) on online customer loyalty: Trust as mediating factor," *International Journal of Academic Research in Business and Social Sciences*, vol. 12, 12 2022.

[42] E. C. Mendoza, "A study of online customers repurchase intention using the 4rs of marketing framework," *International Review of Management and Marketing*, vol. 11, pp. 1–10, 3 2021.

[43] C. Galera-Zarco and J. A. Campos, "Exploring servitization in industrial construction: A sustainable approach," *Sustainability*, vol. 13, pp. 8002–, 7 2021.

[44] null Krzysztof, "New trends in consumer behaviour in 2020-2021: A comparative analysis of marketing strategies during the covid-19 pandemic in poland and south korea," *EUROPEAN RESEARCH STUDIES JOURNAL*, vol. XXIV, pp. 596–608, 9 2021.

[45] S. Shorman, M. H. Allaymoun, and O. A. H. Hamid, "Developing the e-commerce model a consumer to consumer using blockchain network technique," *International Journal of Managing Information Technology*, vol. 11, pp. 55–64, 5 2019.

[46] N. Larasati, I. Fahmi, and A. W. Palupi, "The design of minimum viable marketing strategy for research-based innovation product (case study pt rumah rumput laut)," *Business Review and Case Studies*, vol. 2, pp. 57–57, 8 2021.

[47] V. Mazilescu, "Tourism industry must benefit from it&c hyperautomation," *Annals of Dunarea de Jos University of Galati. Fascicle I. Economics and Applied Informatics*, vol. 27, pp. 97–106, 4 2021.

[48] E. A. Krasilnikova, S. Panasenko, V. P. Cheglov, A. F. Nikishin, and O. A. Boris, "Transformation of goods supply model under conditions of digital economy," *Revista Gestão Inovação e Tecnologias*, vol. 11, pp. 480–493, 7 2021.

[49] X. Zhang, J. Luo, and Q. Li, "Do different reputation systems provide consistent signals of seller quality: a canonical correlation investigation of chinese c2c marketplaces," *Electronic Markets*, vol. 22, pp. 155–168, 5 2012.

[50] G. Xie, L. Huang, C. Apostolidis, Z. Huang, W. Cai, and G. Li, "Assessing consumer preference for overpackaging solutions in e-commerce.," *International journal of environmental research and public health*, vol. 18, pp. 7951–, 7 2021.

[51] R. Anindita and D. P. Perdana, "Website aesthetic in millenial's customer experience on satisfaction and impulse buying: Ewom as moderator," *Media Ekonomi dan Manajemen*, vol. 37, pp. 210–210, 7 2022.

[52] A. M. Kurnianto, T. Y. R. Syah, S. Pusaka, and D. Ramdhani, "Marketing strategy on the project planning of retail business for garage shop," *International Journal of Multicultural and Multireligious Understanding*, vol. 6, pp. 217–228, 2 2019.

[53] M. Dehnert, "Sustaining the current or pursuing the new: incumbent digital transformation strategies in the financial service industry," *Business Research*, vol. 13, pp. 1071–1113, 12 2020.

[54] U. B. C. Nyamanza, "The dynamics of smes growth and electronic marketing potential in developing economies: Evidence from the manufacturing firms in tanzania," *Journal of Economics, Finance And Management Studies*, vol. 04, 8 2021.

[55] K. Madhura and N. Panakaje, "Customer perception in relationship between social-media and purchasing behavior of fashion products," *Zenodo (CERN European Organization for Nuclear Research)*, 8 2022.

[56] L. Ellitan, "The information technology industrial revolution and its role in building business strategy of global retail," *Jurnal Manajemen Maranatha*, vol. 19, pp. 151–158, 5 2020.

[57] G. Baranauskas, "Application of customisation and personalisation in digital solutions of the non-life insurance market: a case study of lithuanian, latvian and estonian e-sales platforms," *Engineering Management in Production and Services*, vol. 13, pp. 68–82, 6 2021.

[58] J. Zhao, M. Ji, and B. Feng, "Smarter supply chain: a literature review and practices," *Journal of Data, Information and Management*, vol. 2, pp. 95–110, 3 2020.

[59] T. Y. Kim, R. Dekker, and C. Heij, "Cross-border electronic commerce: Distance effects and express delivery in european union markets," *International Journal of Electronic Commerce*, vol. 21, pp. 184–218, 1 2017.

[60] S. Michel, A. Michaud-Trévinal, and F. Cocula, "Net impacts in front office is: a first operationalization of delone and mclean model in the banking sector," *Electronic Journal of Information Systems Evaluation*, vol. 22, pp. 92–112, 11 2019.

[61] A. R. C. Perinotto, S. B. Simões, S. M. de Sousa, and S. de Souza Braga, "Impacto das plataformas digitais na oferta de meios de hospedagem: Comparativo de vendas entre booking.com e airbnb," *Marketing & Tourism Review*, vol. 5, 9 2020.

[62] W. Tang, "Solution to the trust gap of cross-border e-commerce under blockchain technology," *Financial Forum*, vol. 9, pp. 161–, 11 2020.

[63] E. Skawiska and R. I. Zalewski, "Success factors of startups in the eua comparative study," *Sustainability*, vol. 12, pp. 8200–, 10 2020.

[64] J. Kwak, Y. Zhang, and J. Yu, "Legitimacy building and e-commerce platform development in china: The experience of alibaba," *Technological forecasting and social change*, vol. 139, pp. 115–124, 6 2018.

[65] A. D. Keyser, K. Verleye, K. N. Lemon, T. L. Keiningham, and P. Klaus, "Moving the customer experience field forward : introducing the touchpoints, context, qualities (tcq) nomenclature," *Journal of Service Research*, vol. 23, pp. 433–455, 6 2020.

[66] 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.

[67] A. ChromÄkov and H. Starzyczn, "zen vztah se zkaznky v malch a stednch podnicch v moravskoslezskm kraji: kvalitativn vzkum," *Acta academica karviniensia*, vol. 19, pp. 42–54, 7 2019.

[68] K. Wesbroom, "Advantages and disadvantages of use of the internet for occupational pension schemes: An industry survey," *Pensions: An International Journal*, vol. 10, pp. 152–166, 3 2005.

[69] Y. Chu, "The characteristics and mode selection of agricultural product e-commerce distribution marketing," *WSEAS TRANSACTIONS ON SYSTEMS AND CONTROL*, vol. 17, pp. 535–548, 12 2022.