A Work Project, presented as part of the requirements for the Award of a Master's degree in Business Analytics from the Nova School of Business and Economics.

Agent-Based Model for a Token-Based Coalition Loyalty Program: Evaluating the Impact of Campaign Review Frequency on Merchant Utility and Retention in a Coalition Loyalty Program

PAULO AUGUSTO DE LIMA
Student ID: 44325

Work project carried out under the supervision of:

Ronnie Figueiredo - ADVISOR

21/01/2026

This work was financially supported by Project Blockchain.PT—Decentralize Portugal with Blockchain Agenda, (Project no 51), WP5: New Data and Knowledge Economy, Call no 02/C05-i01.01/2022, funded by the Portuguese Recovery and Resilience Program (PPR), The Portuguese Republic, and The European Union (EU) under the framework of the Next Generation EU Program.

Abstract

This work project develops an agent-based model to examine how behavioural feedback, strategic adaptation, and governance rules interact in a token-based coalition loyalty platform. The model simulates heterogeneous consumers, adaptive merchants, and a shared points economy, allowing participation and performance to emerge endogenously over time. Using the Compro no Fundão platform as a reference context, the study investigates how incentive design and decision-making structures influence merchant retention, consumer engagement, and sustainability. The findings demonstrate the value of simulation-based analysis for anticipating coordination challenges in digital loyalty ecosystems.

Keywords

Agent-Based Model; Coalition Loyalty Program; Digital Platform; Stakeholder Participation; Token Economy

Acknowledgements

We would like to express our heartfelt gratitude to **Ronnie Figueiredo** (NOVA SBE Data, Operations & Technology Knowledge Center - DOTKC), **Deonilde Clérigo**, and **Marco Carvalho** (Celfocus) for their insights, perspectives, and support during this work project. We also acknowledge the use of AI for streamlining processes.

This work used infrastructure and resources funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019 and Social Sciences Data Lab, Project 22209), POR Lisboa (LISBOA-01-0145-FEDER-007722 and Social Sciences Data Lab, Project 22209) and POR Norte (Social Sciences Data Lab, Project 22209).

Table of Contents

1 Introduction 3
2 Objectives and Relevance of the Study 5
2.1 Platform Sustainability and Token Economy 6
2.2 Customer Adoption and Engagement Dynamics 7
2.3 Merchant Strategies and Participation 7
2.4 Financial Viability and ROI Analysis 8
3 Literature Review 9
3.1 Loyalty Market and Platform Models 10
3.2 Customer Adoption and Engagement Dynamics 14
3.3 Merchant Strategies and Participation 17
3.4 Financial Viability and Return-on-Investment (ROI) Analysis 19
3.5 Methodological Foundations: Agent-Based Model in Business Research 21
4 Methodology 27
4.1 Problem Definition 30
4.2 Objectives of the Solution 30
4.3 Artefact Design and Development 31
4.4 Demonstration and Experiment Design 34
4.5 Evaluation Strategy 34
4.6 Expected Contributions 36
5 Evaluating the Impact of Campaign Review Frequency on Merchant Utility and Retention in a Coalition Loyalty Program 37
5.1 Introduction 37
5.2 Experimental Design and Methodology 38
5.3 Results: 30% Initial Enrolment 40
5.4 Results: 60% Initial Enrolment 43
5.5 Mechanism Analysis 46
5.6 Theoretical Implications 49
5.7 Practical Recommendations 50
5.8 Limitations 51
5.9 Future Research Directions 51
6 Agent-Based Insights in a Coalition Loyalty Program Context 52
7 Conclusion and Future Directions 53
References 56
Appendix 67

List of Abbreviations

ABM – Agent-Based Model

AI – Artificial Intelligence

CAC – Customer Acquisition Cost

CLP – Coalition Loyalty Program

CLV – Customer Lifetime Value

CRM – Customer Relationship Management

DSR – Design Scientific Research

DSRM – Design Science Research Methodology

EMA – Exponential Moving Average

EU – European Union

GenAI – Generative Artificial Intelligence

GMV – Gross Merchandise Value

GSR – Global Sensitivity Analysis

KPI – Key Performance Indicator

LLMS – Large Language Models

LP – Loyalty Program

LPE – Loyalty Program Engagement

ODD – Overview, Design Concepts, Details

PRR – Portuguese Recovery and Resilience Program

ROI – Return on Investment

SD – Standard Deviation

TRACE – Transparent and Comprehensive Model Evaluation

1 Introduction

Loyalty programs have evolved over several decades from simple transactional schemes into interconnected ecosystems where customers, merchants, and platform operators co-create value, marking a shift from isolated brand initiatives to integrated networks of engagement (Bijmolt, Dorotic, and Verhoef 2011). This historical progression sets the stage for understanding how digitalisation has transformed loyalty into a more complex strategic function. Initially designed to reward repeat purchases, these programs now leverage digital infrastructures and data analytics to enable personalised experiences and cross-brand engagement (Obiegbu and Larsen 2024). Industry research similarly reports that emotional connection and data-driven personalisation have become dominant loyalty drivers in recent years (EY 2025). The emergence of multi-vendor and coalition models has expanded their scope, allowing multiple firms to share customer data, marketing efforts, and reward mechanisms (Bijmolt, Dorotic, and Verhoef 2011). These ecosystems blend competition and collaboration, with success depending on the alignment of incentives across heterogeneous participants (Breugelmans et al. 2015).

Building on this technological evolution, digitalisation has further transformed loyalty into a strategic layer of customer relationship management (Lin and Bowman 2022). Advanced analytics and AI-driven personalisation not only enhance engagement but also create new dependencies on data sharing, governance, and technological coordination. Furthermore, the integration of blockchain technologies marks another turning point, providing secure and transparent infrastructures for managing rewards and transactions (Anudeep et al. 2024). Early consultancy insights anticipated these developments, suggesting that blockchain could improve interoperability, reduce administrative costs, and increase trust among ecosystem partners (Deloitte 2016). Blockchain-enabled systems have shown potential to reduce administrative friction, prevent fraud, and enable interoperability across brands (Utz et al. 2023). Within these architectures, tokenisation, defined as the process of converting assets or rewards into blockchain-based units of value, emerges as a foundational design element that enables flexible incentive structures and peer-to-peer value exchange (Boukis 2024; Chainlink 2024).

However, as digital loyalty systems become more sophisticated, they introduce new managerial and economic challenges. For instance, token-based programs may face the challenge of balancing engagement with stability: excessive issuance can cause inflationary effects, while restrictive redemption policies can discourage participation (Tanveer, Ishaq, and Hoang 2025). Governance and incentive alignment also become critical, as decision power and value distribution shift among decentralised actors (Abadi and Brunnermeier 2024; Nguyen and Nguyen 2025). These shifts are also reflected in consumer behaviour, as participation in digital loyalty ecosystems continues to rise across Europe, with users increasingly engaging in cross-brand and digital-first programs (Mando-Connect and YouGov 2024).

At the same time, financial sustainability presents a separate but related set of challenges. Empirical evidence from coalition programs shows that variations in reward value can substantially alter customer and merchant behaviour, affecting both profitability and retention (Stourm and Bradlow 2023). In addition, unredeemed currency can become a significant liability on balance sheets, leading many programs to adopt expiration policies (Breugelmans et al. 2015). Beyond these behavioural effects, maintaining liquidity, managing redemption liabilities, and ensuring a positive return on investment across stakeholders are essential to the long-term viability of loyalty ecosystems (Cong, Li, and Wang 2021). Together, these findings indicate that sustaining engagement, profitability, and value creation in such ecosystems requires a systemic understanding of behavioural, strategic, and financial interdependencies.

Addressing these complexities of digital loyalty ecosystems often exceeds the capacity of traditional analytical models, which are typically limited by equilibrium assumptions and lack agent-level interaction detail (Rand and Rust 2011; Macal and North 2009). These ecosystems are inherently adaptive, shaped by heterogeneous agents whose decisions evolve in response to others. Agent-Based Model (ABM), therefore, provides a suitable approach for analysing these phenomena. By representing consumers, merchants, and platforms as autonomous agents with decision rules, ABM enables researchers to study how local interactions generate emergent system-level outcomes (Rand and Stummer 2021; Waldherr, Hilbert, and González-Bailón 2021). This makes it especially useful for exploring how digital incentives, feedback loops, governance mechanisms, and financial flows influence long-term stability and engagement.

Against this backdrop, the present study investigates how digital loyalty platforms can achieve sustainable participation, profitability, and token value stability through well-calibrated incentive structures. The research uses an Agent-Based Model to simulate interactions between customers, merchants, and platform operators, examining how different configurations affect engagement, financial performance, and ecosystem resilience. By integrating behavioural feedback, economic design, and ROI analysis, this work contributes to both the academic understanding of loyalty ecosystems and the managerial design of data-driven, token-based loyalty programs.

2 Objectives and Relevance of the Study

Following the contextual overview, this section outlines the main objectives and relevance of the study. Developed in collaboration with Celfocus under the Blockchain.PT initiative, the project applies Agent-Based Model (ABM) to simulate how customers, merchants, and platform operators interact within digital loyalty ecosystems. Through this approach, the study explores how token-based incentives and behavioural feedback mechanisms can be designed to foster sustainable, data-driven loyalty platforms that balance engagement, stability, and profitability over time.

The study is guided by four interrelated objectives: (i) platform sustainability and token economy, (ii) customer adoption and engagement, (iii) merchant strategies and participation, and (iv) financial viability. These objectives were defined collectively to reflect the multi-stakeholder nature of loyalty ecosystems and the dynamic feedback loops that ABM can represent. Each objective corresponds to one of the platform's main dimensions: economic design, user behaviour, merchant participation, and financial sustainability. This ensures that both micro-level decisions and macro-level outcomes are captured across the different dimensions of the loyalty platform.

2.1 Platform Sustainability and Token Economy

This objective examines the sustainability of token and reward mechanisms within digital loyalty platforms. It focuses on how reward rates, redemption rules, and inflation controls influence long-term platform stability. The aim is to design token economies that maintain value over time, preventing inflationary pressures while still providing meaningful incentives for user adoption and participation.

From a modelling perspective, this objective defines the agent behaviours and parameters that govern token issuance, redemption, and inflation dynamics within the ABM environment. It evaluates how variations in reward rates and redemption rules affect token value, user engagement, and overall platform stability across different simulated scenarios. The results are expected to identify the conditions and design principles under which loyalty platforms can remain economically sustainable without devaluing rewards, contributing both to the academic understanding of token sustainability and to the practical design of long-term loyalty programs.

2.2 Customer Adoption and Engagement Dynamics

This objective focuses on how customers join a coalition loyalty program, how engagement evolves over time, and which design choices sustain participation. It examines how different incentive schemes, such as reward generosity through tiered reward structures or sign-up incentives shape customers' decisions to register, stay active, or churn. Ultimately, we aim to understand the full engagement lifecycle: From initial unawareness to awareness of the program, the decision process to registration, redemption behaviour, and eventual disengagement.

Understanding how consumers adopt and engage with loyalty programs is essential for platform designers seeking to balance acquisition costs against long-term retention. By simulating heterogeneous consumer populations under varying incentive configurations, the ABM reveals which design choices produce sustainable engagement patterns and which foster disengagement in a coalition loyalty setting.

2.3 Merchant Strategies and Participation

Our third objective is to understand how merchants decide to join, remain, or exit the loyalty platform and to identify any strategic behaviours that emerge during the simulation. This includes analysing how participation decisions are influenced by perceived profitability, data-sharing benefits, and cost distribution fairness, as well as how merchants adapt their strategies through pricing, promotions, or partnerships, to maximise value within the ecosystem.

This objective is particularly relevant because it helps identify the design principles that support sustainable merchant participation and cooperative behaviour. By anticipating both beneficial and harmful strategic patterns, the model contributes to understanding how governance mechanisms, incentive structures, and token-based rewards can be adjusted to promote platform stability and long-term value creation.

2.4 Financial Viability and ROI Analysis

Our goal is to evaluate the loyalty platform's financial viability from the viewpoints of all parties involved (platform, merchants, and customers) and to calculate the return on investment that comes from joining the coalition. This entails examining how platform revenue changes in relation to point liabilities, how merchants' investments in point purchases and discounts result in quantifiable increases in sales, and how customers obtain value through redemption and savings practices.

This goal is especially important since it lays the financial groundwork required for the platform's long-term viability. The analysis determines the circumstances in which the platform reaches self-sustaining economics by simulating the movement of points through the ecosystem from issuance to redemption and monitoring important financial metrics like redemption rates, breakage estimates, and liability exposure. Sensitivity analysis also identifies the factors that have the biggest effects on stakeholder outcomes, allowing evidence-based suggestions for commission rates, pricing schemes, and incentive calibration. Designing loyalty coalitions that provide equitable value to all members while preserving operational sustainability requires an understanding of these financial dynamics.

Taken together, these objectives integrate behavioural and economic mechanisms into a unified analytical framework for the loyalty ecosystem. The ABM approach enables the team to test multiple scenarios and examine how changes in incentives and parameters shape participation, cooperation, and financial outcomes. This integrated perspective supports the translation of complex system dynamics into actionable insights for platform design and management.

The relevance of this study spans academic, practical, and societal domains. Academically, it contributes to the growing application of ABM in digital platform research by examining how decentralised incentive systems influence engagement and sustainability. Practically, the industry collaboration ensures that the findings directly inform innovation initiatives within the Blockchain.PT framework, offering guidance on reward design, stakeholder alignment, and long-term platform viability. More broadly, the project supports sustainable digital transformation by illustrating how responsibly designed token systems can balance economic performance with long-term value creation.

This section establishes the foundation for the subsequent parts of the report, which present the literature review, methodological framework, and simulation design that operationalise these objectives within the Celfocus–Blockchain.PT field lab.

3 Literature Review

The purpose of this literature review is to build a coherent theoretical foundation for understanding how digital loyalty platforms function as multi-sided ecosystems connecting consumers, merchants, and platform providers. It aims to consolidate insights from marketing, behavioural economics, and computational modelling to identify the key factors influencing loyalty program success and to explore how agent-based models (ABM) can be applied to simulate and analyse these dynamics.

The review focuses on digital and token-based loyalty programs, reflecting the increasing use of gamification, blockchain, and reward-based incentives in modern customer engagement systems. It examines how consumer adoption, merchant participation, and platform strategy interact to create or hinder value within these ecosystems. While the discussion touches on technological and financial considerations, the emphasis remains on behavioural, strategic, and systemic mechanisms rather than technical implementation details.

To maintain a clear narrative, the review is organised into five interconnected sections. It begins by outlining the loyalty market and platform models, followed by an analysis of customer adoption and engagement dynamics, and merchant strategies and participation. Next, it considers financial viability and return-on-investment perspectives before introducing the relevance of agent-based models in business research. The review concludes by synthesising these literature streams to motivate the group's agent-based model, illustrating how simulation-based approaches can complement traditional analytical methods in studying dynamic loyalty ecosystems.

By structuring the review in this way, the study therefore aims to bridge theoretical and methodological perspectives, linking established marketing and behavioural theories with emerging computational tools. The outcome is an integrative understanding of how ABM can be leveraged to model complex interactions and assess the long-term sustainability of loyalty platforms.

3.1 Loyalty Market and Platform Models

Loyalty programs have increasingly evolved from firm-centric schemes into multi-sided digital platform ecosystems in which consumers, merchants, and platform operators jointly create and capture value. Drawing on insights from marketing, economics, and information systems, this section synthesises the theoretical and empirical foundations of loyalty markets to explain how these interactions are structured and governed. By tracing the transition from traditional loyalty schemes to data-driven and tokenised models, it establishes the system-level foundations for analysing customer adoption and engagement dynamics discussed in the following section.

3.1.1 Evolution of Loyalty Platforms

Over the past three decades, loyalty programs have developed from simple transactional tools into complex, data-driven ecosystems that support long-term customer relationships. Loyalty program research has expanded from status-based and inertia-based mechanisms toward relationship-based and technology-enabled frameworks (Chen, Mandler, and Meyer-Waarden 2021). While early programs primarily rewarded repeat purchases through points or discounts, contemporary approaches emphasise personalisation, experience, and emotional connection (Obiegbu and Larsen 2024). Digitalisation and artificial intelligence have enabled richer data collection and individualised offers, transforming loyalty programs into strategic components of customer relationship management. Yet, Lin and Bowman (2022) demonstrate empirically that the introduction of a loyalty program does not always guarantee sustained performance; sales and profitability often rise initially but can decline without continuous adaptation and data-driven refinement. This evolution highlights a key insight: loyalty programs must now operate as adaptive, learning systems that align incentives across all participants rather than static reward schemes.

This diversification has occurred not only in the technological sophistication of loyalty mechanisms but also in their organisational structures. Several noteworthy configurations have gained considerable recognition in both research and practice. Single-firm programs, such as airline frequent-flyer schemes or retailer discount cards, operate within a single organisation and primarily leverage same-side effects like customer retention and habit formation (Dowling and Uncles 1997; Sharp and Sharp 1997). Coalition or multi-vendor programs connect multiple firms through a shared platform, managed either by a neutral third party or a leading coalition member, that facilitates cross-brand earning and redemption while relying on cross-side network effects between consumers and merchants (Berman 2006; Capizzi and Ferguson 2005). More recently, tokenised or digital platforms extend this logic to blockchain-based environments, introducing transferable and interoperable assets that can circulate across multiple actors and markets (Abooleet and Kinnett 2023; Boukis 2024; Wang, Luo, and Lee 2019). These models differ not only in scope and governance but also in how they generate and distribute value across participants.

3.1.2 Platform Economics and Network Effects

Shifting towards digital and multi-brand loyalty ecosystems has brought these programs under the theoretical umbrella of two-sided or multi-sided platform economics. Successful platforms create value by coordinating interactions between two distinct but interdependent user groups (Rochet and Tirole 2003). Applied to loyalty contexts, consumers represent one side, and merchants or partner brands form the other (Chen, Mandler, and Meyer-Waarden 2021; Lin and Bowman 2022). The platform's role is to balance incentives, often subsidising one side (for instance, customers through free points or bonuses) to stimulate participation on the other. This framework clarifies why loyalty programs thrive on network effects: as more consumers enrol, merchants perceive higher marketing value; as more merchants participate, consumers find greater redemption opportunities. Pricing, reward structure, and governance decisions, therefore, mirror the strategic trade-offs identified in platform theory. The economic interdependence outlined by Rochet and Tirole (2003) provides a foundation for understanding how digital loyalty markets function as ecosystems rather than linear firm–customer relationships.

3.1.3 Data Driven Value Creation

Contemporary loyalty platforms increasingly rely on data analytics and algorithmic personalisation to sustain engagement and differentiate their offerings (Chen, Mandler, and Meyer-Waarden 2021). This concept has been framed through the lens of experiential brand loyalty, which shows how personalised, AI-driven interactions generate positive emotions and a sense of being "understood" (Obiegbu and Larsen 2024). Conversely, inadequate personalisation can create a sense of de-personalisation, leading to user disengagement. This insight reframes data as a tool for segmentation and as a source of experiential value within the loyalty exchange. Data flows between consumers and merchants also reinforce cross-side network effects: improved targeting enhances merchant returns, while relevant offers increase consumer utility, making data the connective tissue of loyalty markets (Lin and Bowman 2022). The integration of algorithmic learning into loyalty design thus represents both a marketing and an operational capability that sustains long-term participation.

3.1.4 Emerging Tokenised Models

A growing body of research explores blockchain-enabled or tokenised loyalty programs as the next evolutionary stage in loyalty ecosystems. Early studies investigated how blockchain can enhance intrinsic and extrinsic motivations in loyalty participation by improving transparency, autonomy, and peer-to-peer exchange (Wang, Luo, and Lee 2019). Building on this foundation, Utz et al. (2023) demonstrated through design research how blockchain architecture can be adapted to manage multi-merchant programs efficiently by ensuring security, real-time transactions, increased transparency, and trust. Extending this perspective, Abooleet and Kinnett (2023) provided a systematic review identifying twenty-one active blockchain-based loyalty projects, concluding that while the field is nascent, real-world implementation is progressing faster than expected. Empirical evidence further supports these conceptual advances, showing that tokenised rewards increase perceived novelty, psychological ownership, and program attractiveness, particularly in premium service contexts (Boukis 2024). Collectively, these studies indicate that tokenisation enhances value creation through interoperability, transparency, and user control, positioning blockchain as an infrastructure for next-generation loyalty ecosystems.

Integrating the sources, loyalty markets can now be viewed as dynamic, multi-agent systems where value emerges from the continuous interaction of consumers, merchants, and platform design choices. The theoretical evolution from transactional programs to tokenised ecosystems underscores three key shifts: from isolated rewards to networked value creation, from static design to adaptive, data-driven personalisation, and from firm-controlled incentives to shared governance through digital tokens. These dynamics directly influence how customers perceive, adopt, and engage with loyalty programs, a theme that the next section explores in greater depth.

3.2 Customer Adoption and Engagement Dynamics

This section identifies the behaviours and mechanisms that move customers to join a digital loyalty program, engage with it, and persist or disengage. It aims to provide an overview of the relevant literature on the engagement lifecycle, from early adoption decisions to repeat engagement with the platform and potential churn.

3.2.1 Conceptual Foundations

Decades of loyalty research show that these mechanisms are not monolithic: Adoption and ongoing participation dynamics are driven by status based (through tiers and social signalling), inertia based (through habit, switching costs or goal progress), and relationship based (through perceived value and reciprocity) processes, with substantial heterogeneity across customers and contexts (Chen, Mandler, and Meyer-Waarden 2021). Small shifts in perceived value, status, reward timing, communication, or incentive setting can affect enrolment, redemption timing, breakage, and long-run loyalty.

Outcomes can be divided into behavioural loyalty like frequency of redemptions, share of wallet, retention and attitudinal loyalty like affect, preference, or advocacy (Belli et al. 2022). Meta-analytic evidence shows LPs move behavioural more reliably than attitudinal loyalty, implying that programs are often effective at changing near-term behaviour but having difficulties in altering deeper attitudes (Belli et al. 2022).

For the customer agent in our model, this suggests focusing on modelling behavioural manifestations like redemption rate and timing, breakage, churn, and CLV rather than attitudinal changes.

3.2.2 Adoption and Enrolment Design

At the point of adoption, customers perform a value assessment, weighing if the expected benefits exceed the perceived costs, with trade-offs between privacy and personalisation in digital programs becoming increasingly important (Chen, Mandler, and Meyer-Waarden 2021). Closed or fee-based memberships can strengthen commitment, but at the cost of smaller eligible pools while open programs maximise reach but risk shallow engagement. Coalition programs raise perceived value yet can dilute firm specific loyalty (Belli et al. 2022). These design contingencies matter for adoption likelihood and the speed of activation.

Empirically, adoption focused outcomes like join/cancel decisions or time to enrol have received less attention in recent years relative to redemption and affective responses (Chen, Mandler, and Meyer-Waarden 2021), a reason to model activation explicitly as a leading indicator in our simulations.

3.2.3 Customer Engagement

Ensuring long-term customer engagement in Loyalty platforms is one of the most difficult obstacles to overcome and loyalty platform managers primary aim should be to establish personal, deep and long-term relationships to consumers (Meyer-Waarden, Bruwer, and Galan 2023).

The literature distinguishes “being a member” from “being engaged”. Loyalty program engagement (LPE) is a multidimensional and hierarchically ordered series of behaviours that require increasing effort by the customer: proactive use of the card/app, redeeming points, adjusting the time of purchase, receiving or searching for LP information, and sharing or recommending. Measurements show that these behaviours vary, with actions that require greater effort occurring less frequently but being more indicative of engagement (Bruneau, Swaen, and Zidda 2018). Critically, perceived value raises engagement, which in turn lifts loyalty and brand loyalty, and thus customer engagement with the company (Meyer-Waarden, Bruwer, and Galan 2023). Evidence across grocery, department store, and airline loyalty programs supports this pathway and identifies "points pressure", short-term acceleration toward rewards, and "rewarded behaviour" which functions like post-redemption stickiness, as mechanisms shaping timing (Taylor and Neslin 2005; Meyer-Waarden, Bruwer, and Galan 2023).

Furthermore, experimental evidence shows that consumers infer their progress from the points medium itself. Manipulating the magnitude of the point-currency and the precision of points-per-dollar information creates "illusionary progress," which then affects perceived time to reward, loyalty, and recommendation, even when actual spending and reward value are held constant (Bagchi and Li 2011).

3.2.4 Design Levers That Shape Engagement

Tiering is a common strategic structure in loyalty programs. Tiers are designed with hierarchical "status levels" that can enhance members' sense of status by offering higher rewards for higher tiers, but demotion from a higher to a lower tier can trigger frustration and reduced loyalty. Meta-analytic evidence shows that on average, tiered programs do not have stronger loyalty effects than non-tiered programs, and tiers themselves do not significantly moderate a loyalty program's effectiveness (Belli et al. 2022).

Belli et al. (2022) additionally find that loyalty programs are more effective when they offer rewards through savings of accumulated points and exclusive benefits, whereas immediate discounts and indirect rewards weaken program impact, underscoring that reward formats framed closely aligned with a firm's core products or services foster stronger loyalty behaviour than generic price cuts or unrelated rewards. Additionally, in continuous, non-expiring points programs, customers tend to stockpile large point balances and redeem often only a small fraction of their points once internal redemption thresholds are reached (Dorotic et al. 2014). Taken together, these findings indicate that strong pre- and post-reward purchase effects can arise even in the absence of point-expiry deadlines, suggesting that expiry policies are not a necessary condition for loyalty programs to generate substantial behavioural lift (Dorotic et al. 2014).

A sign-up (enrolment) bonus that gives new members an immediate "head start" through pre-loaded points/credits can increase persistence toward earning the first reward and thereby support early engagement (Nunes and Drèze 2006). However, Nishio and Hoshino (2024) observed that one-off loyalty incentives, such as sign-up bonuses, had no positive effect on long-term customer value, suggesting minimal impact on true loyalty.

3.3 Merchant Strategies and Participation

Understanding merchant behaviour within loyalty ecosystems has become increasingly important as research evolves beyond a consumer-centric perspective. Early studies focused primarily on customer behaviours and responses, treating merchants as contextual variables rather than strategic actors. This reflected the prevailing market landscape at that time: most loyalty initiatives were single-firm programs such as airline or grocery schemes, where inter-firm collaboration was rare (Bijmolt, Dorotic, and Verhoef 2011). Coalition programs like Air Miles (Canada) or Nectar (UK) were exceptional cases, and as a result, research initially emphasised firm–customer interactions rather than firm–firm relationships. Recognising this gap, Bijmolt, Dorotic, and Verhoef (2011) urged a shift toward examining vendor behaviour, and Dorotic et al. (2011) provided one of the first empirical analyses of multi-vendor systems, demonstrating that merchant strategies significantly shape coalition performance. This marked a turning point toward treating merchants as active participants in value co-creation rather than passive adopters of a given platform design.

3.3.1 From Firm-Centric to Network-Oriented Perspectives

As research matured, attention expanded to inter-firm dynamics within these coalitions. Breugelmans et al. (2015) showed that as loyalty networks grew, partner coordination, incentive alignment, and cost distribution became central to sustaining program viability. Similarly, Bijmolt and Verhoef (2017) argued that modern loyalty ecosystems operate as data-driven, platform-based collaborations in which governance and cooperation are as critical as consumer engagement. These studies collectively reveal a conceptual reorientation: merchant strategy and platform design are interdependent, shaping each other through continuous feedback.

3.3.2 Merchant Behaviour and Competitive Dynamics

Recent findings highlight merchants' adaptive behaviour within these systems. Dorotic et al. (2021) demonstrated that competitive tensions, especially among dominant or overlapping firms, can erode coalition value, while Stourm and Bradlow (2023) revealed asymmetries in how reward policy changes affect different merchants. Such evidence portrays merchants as heterogeneous, strategic agents balancing cooperation and competition. This heterogeneity is crucial for modelling purposes, as it implies that equilibrium behaviour cannot be assumed but must emerge from interactions between differently positioned firms.

3.3.3 Motivations and Constraints for Participation

Extending this behavioural perspective, De Noni, Orsi, and Zanderighi (2014) identified dual participation drivers, economic (customer acquisition, shared marketing) and relational (reputation, data access), while emphasising that unequal costs or data asymmetry can undermine trust and lead to exit. Cost-sharing and governance mechanisms thus emerge as structural levers of stability (Cao, Nsakanda, and Diaby 2015; Breugelmans et al. 2015). Smaller firms face distinctive barriers, including high entry costs and uncertain returns, which link participation decisions to perceptions of fairness and coalition design quality.

3.3.4 Implications for Modelling Merchant Agents

Merchants play a critical role in coalition systems, where their decisions to join, invest, or exit can significantly influence the system's trajectory. Traditionally, studies on loyalty programs relied on econometric methods and surveys to analyse consumer and firm behaviour. However, as platforms and multi-brand loyalty programs gained traction, those approaches failed to capture dynamic feedback loops and adaptive behaviours that are critical in these more complex network structures.

Agent-based model offers a more effective approach to representing these interdependencies and examining how merchants adapt their strategies within coalition ecosystems. By simulating varied strategies and adaptive learning, ABM illustrates the evolution of cooperation, competition, and cost-sharing over time, affecting coalition sustainability. This evolution clarifies merchant incentives and strengthens the link between behavioural economics and computational modelling in loyalty research.

In the context of our model, merchant-level KPIs primarily include participation decisions (whether to enter, remain, or exit the coalition) and campaign participation intensity (the decision to engage in joint promotions or token-based incentive events promoted by the platform). These outcomes reflect both strategic commitment and adaptive responsiveness, serving as proxies for merchant-side performance and stability within the coalition.

3.4 Financial Viability and Return-on-Investment (ROI) Analysis

Financial viability and sustainable return-on-investment (ROI) are central challenges for tokenised business models, given their reliance on incentive-driven adoption, novel financing structures, and complex ecosystem dynamics. This section synthesises prior literature examining these challenges in the context of value creation and platform viability within tokenised ecosystems.

3.4.1 Theoretical Perspectives on Financial Sustainability in Tokenised Business Models

The financial sustainability of tokenised business models has emerged as a critical area of inquiry in recent literature. Tokenomics research demonstrates that dynamic adoption mechanisms and proper token distribution are essential for platform valuation and long-term viability (Cong, Li, and Wang 2021; Abadi and Brunnermeier 2024). Token-based platforms create unique financing structures that differ fundamentally from traditional business models, enabling platforms to bootstrap network effects through carefully designed incentive systems (Cong, Li, and Wang 2022; Li and Mann 2025).

Business model sustainability requires alignment between product market strategy and the underlying organisational structure, with system dynamics approaches offering valuable frameworks for understanding feedback loops and long-term equilibrium states (Zott and Amit 2008; Abdelkafi and Täuscher 2016). From a complexity economics perspective, tokenised ecosystems exhibit emergent properties where agent interactions create non-linear outcomes that cannot be predicted through traditional financial modelling alone (Arthur 2021; Van der Merwe et al. 2019).

3.4.2 Design Mechanisms for Sustainable Tokenised Ecosystems

The integration of blockchain technology into loyalty programs presents significant opportunities for reducing liability management costs while enhancing program value through improved liquidity and transparency (Chun, Iancu, and Trichakis 2020; Luo 2024; Deloitte 2016). Self-sustaining ecosystem design requires careful consideration of token utility, governance structures, and value capture mechanisms to ensure financial viability across multiple stakeholder groups (Shah and Jansen 2021; Freni, Ferro, and Moncada 2022).

3.4.3 KPI’s and Relevance to Platform Viability

Agent-based model provides powerful tools for simulating market diffusion patterns and understanding how individual agent behaviours aggregate into observable system-level KPIs, particularly relevant for assessing token adoption trajectories and network growth (Rand and Stummer 2021; Axtell and Farmer 2025).

3.4.4 Platform Economics

The economics of multi-sided platforms provide essential theoretical foundations for understanding tokenised ecosystems, where network effects and coordination mechanisms determine value creation and distribution (Hagiu and Wright 2015; Evans and Schmalensee 2016). Understanding these complex platform dynamics reveals that successful tokenised platforms emerge from properly calibrated feedback mechanisms between supply-side and demand-side participants, with agent-based approaches offering computational tools for exploring design alternatives before deployment (Arthur 2021; Axtell and Farmer 2025).

3.5 Methodological Foundations: Agent-Based Model in Business Research

This section presents the methodological foundations of the study, detailing how the agent-based model (ABM) was applied to analyse interactions within the loyalty-platform ecosystem. It explains the theoretical grounding, model structure, and integration process adopted by the group, outlining how individual agents, decision rules, and feedback mechanisms collectively support the study’s research objectives.

3.5.1 Foundations of Agent-Based Model

Agent-Based Model (ABM) has become an increasingly valuable approach for examining complex systems in the social sciences and management research. Rather than assuming centralised control or equilibrium, ABM studies how local interactions among autonomous agents generate collective outcomes that are often nonlinear and unexpected (Bonabeau 2002). Each agent follows a set of behavioural rules and goals, interacting with others and the surrounding environment in ways that evolve over time. When simulated collectively, these micro-level behaviours reveal macro-level patterns such as market growth, stability, or decline (Tesfatsion 2006).

This bottom-up approach makes ABM particularly suited to the study of business ecosystems, where customers, merchants, and platforms behave independently yet remain interdependent. Traditional analytical or econometric models usually rely on aggregate assumptions that overlook heterogeneity and feedback effects. ABM, by contrast, allows bounded rationality, adaptive learning, and behavioural diversity to be represented explicitly (Rand and Rust 2011). In practice, this means that even minor changes in incentives or information flows can lead to major systemic shifts, outcomes that linear models often miss.

More recent overviews show that ABM has matured into a mainstream methodology for analysing socio-economic and organizational systems (Macal 2016; Steinbacher et al. 2021). Building on this progress, Onggo and Foramitti (2021) provide a comprehensive review of ABM applications in business and management, demonstrating its versatility in modelling decision-making, market dynamics, and adaptive organizational processes. Its flexibility enables researchers to explore "what if" scenarios that combine qualitative reasoning about behaviour with quantitative simulation. For instance, an ABM can test how different reward structures or token-liquidity policies influence consumer retention or merchant participation over time. Such experiments help identify conditions that sustain equilibrium between growth and profitability, an essential concern in loyalty-platform sustainability research.

Within innovation and sustainability studies, ABM has also been applied to model adaptive market systems in which coordination emerges from decentralised decision-making (Garcia and Jager 2011). As Secchi (2015) emphasises, simulation enables organisations to replicate behavioural rules, feedback mechanisms, and environmental constraints that shape real performance. This capacity to reproduce and observe adaptive behaviour makes ABM a powerful tool for investigating how loyalty platforms evolve under tokenised incentive systems.

3.5.2 Philosophical Grounding

The following section examines the philosophical assumptions that position ABM as compatible with a critical-realist perspective.

ABM is conceptualised as a mechanism-seeking, process-theory methodology grounded in ontological realism (that is, phenomena are presumed to exist independently of models) and epistemological fallibilism, which holds that knowledge is inherently partial and subject to revision. This perspective prioritises mechanism-based explanations over black-box predictions by specifying plausible micro-level rules and demonstrating how they produce macro-level patterns, rather than merely fitting empirical correlations. Methodologically, research proceeds through an abduction-to-deduction cycle: candidate mechanisms are inferred from theory and evidence, then tested via simulation; abduction proposes possible explanations, while deduction elucidates their implications. The emphasis is on the epistemological aspect of emergence, describing macro-outcomes that are difficult to anticipate from micro-level components but become intelligible when underlying rules and interactions are made explicit, rather than asserting strong "ontological" emergence. Models are intentionally selective, incorporating only causally relevant elements for the research objective, and are evaluated for selective realism and robustness, meaning results are expected to hold across plausible parameter variations. Recognising inherent limitations such as underdetermination, opacity, implementation error, and the gap between models and real-world agency, motivates transparency, replication, and triangulation with analytical and empirical approaches (Miller 2015).

3.5.3 Model Validity and Documentation

Validity as defined by Troost et al. (2023) entails adequacy for purpose: a model is considered valid when its design and inferences are suited to the specific research context, and when uncertainties are appropriately linked to final claims. In practice, this entails: (1) clearly defining the modelling context and objectives; (2) selecting components and methods that are theoretically and empirically justified; and (3) evaluating outcomes through sensitivity and uncertainty analyses. The KIA 12-step protocol operationalises these principles, supporting coherence throughout model development.

To ensure transparency, we adopt the ODD protocol (Grimm et al. 2010) to document the model's structure and processes and draw on the TRACE framework (Grimm et al. 2014) and KIA protocol (Troost et al. 2023) to guide our modelling, validation, and reporting. ODD specifies what the model is, including its purpose, entities, variables, processes, and design concepts, while TRACE records model construction and testing, including design decisions, calibration, validation, and code provenance. ODD standardises communication, and TRACE clarifies the modelling process, reinforcing the adequacy-for-purpose principle. Concise ODD and TRACE-based documentation tables are provided in the Appendix (Table 5-6).

The proposed ABM represents the loyalty ecosystem as a set of interacting consumers, merchants, and a platform, where collective decision-making shapes overall program performance. The aim is to interpret how micro-level factors, such as consumer engagement, merchant participation, and platform design, drive macro-level outcomes, including adoption, retention, and sustainability. By modelling these interdependencies, it is possible to examine how different incentive structures, governance strategies, and behavioural assumptions affect value creation and system stability in digital and token-based loyalty environments. In alignment with the adequacy-for-purpose principle, the model will be iteratively refined based on empirical evidence and will be documented according to ODD and TRACE standards to ensure clarity and reproducibility.

3.5.4 Integrating ABM into the Group's Framework

After defining the documentation standards via ODD and TRACE-based validation procedures, the next step is to translate these methodological foundations into the design of our group's simulation model. This subsection focuses on how the agent-based model approach is integrated within our loyalty-platform framework. Recent reviews emphasise that ABM has evolved from conceptual experimentation to an applied analytical tool capable of linking micro-level behaviour (Ribeiro-Rodrigues and Bortoleto 2024). This integrative capacity makes ABM particularly relevant for analysing loyalty-platform ecosystems in which consumers, merchants, and the platform itself interact through incentive and feedback mechanisms.

Rizzati and Landoni (2024) show that ABM captures interdependencies among heterogeneous agents pursuing distinct yet complementary objectives within circular-economy transitions. Their review underlines that adaptive decision rules and local learning can generate emergent patterns of cooperation, competition, or instability, dynamics directly applicable to digital ecosystems. Extending this perspective, Koide et al. (2023) demonstrate how agent-based experimentation can assess combinations of circular-economy strategies and quantify their effects on diffusion, circularity, and sustainability. Similarly, by iterating behavioural and policy adjustments, loyalty-platform simulations therefore help identify token-reward designs that balance user engagement and long-term viability.

At the micro level, Liu, Onggo, and Busby (2024) model user participation in new-product development, helps to illustrate how feedback loops between individual engagement and collective innovation foster sustained loyalty. This dynamic closely reflects how consumer agents behave in our framework. Ghanem, Leitner, and Jannach (2022) likewise apply ABM to recommender-system contexts, revealing how firms can balance consumer satisfaction and business profitability through adaptive decision rules. These insights inform our representation of merchants as strategic agents whose promotional intensity and partnership choices evolve in response to platform incentives and consumer patterns.

Integrating these behavioural layers, Secchi et al. (2024) argue that ABM can bridge organisational behaviour and sustainable-development theory by connecting individual actions with broader environmental objectives. In our design, this translates into modelling how token governance and participation rules affect both economic performance and ecological alignment. The integrated framework, therefore, represents a co-evolving system in which agents learn, adapt, and generate emergent outcomes consistent with real-world sustainability transitions.

Ultimately, the model's utility depends on how its results are communicated to decision-makers. Arnejo et al. (2025) highlight that the effectiveness of ABM research increases when simulation outcomes are expressed in forms accessible to stakeholders. Following this insight, our framework incorporates visualisation and narrative-reporting components that link agent interactions to managerial interpretation, ensuring that the simulation functions as both an analytical experiment and a decision-support instrument. Through this integration of methodological rigour, behavioural realism, and communicative clarity, the ABM provides a coherent foundation for exploring the long-term sustainability of digital-loyalty ecosystems.

4 Methodology

This study follows the Design Science Research (DSR) methodology (Hevner et al., 2004), which combines scientific rigour with practical problem-solving. DSR is particularly suited to this project because it provides a structured process for designing, building, and evaluating innovative artefacts that address real-world challenges and contribute to theoretical understanding. In this project, the artefact is an Agent-Based Model (ABM) developed to simulate the behavioural, strategic, and financial dynamics of token-based loyalty ecosystems. The model will be grounded on the "Compro no Fundão" pilot platform developed by Celfocus, which operates as a blockchain-enabled multi-partner loyalty system designed to connect local merchants and consumers within a shared digital rewards network. This real-world case provides a reference to testing a benchmark of how merchant, consumer, and platform decisions interact within a shared token economy.

Building on Hevner et al. (2004) and Peffers et al. (2007), this research follows the DSR logic that connects the environment, design, and knowledge base through iterative build-and-evaluate cycles (Figure 1 and Figure 2).

This logic emphasises the co-evolution of problem understanding and artefact development, ensuring that design decisions remain grounded in both empirical context and established theoretical foundations. Within this framework, the ABM is developed and refined iteratively as a design artefact that responds to real-world constraints while contributing generalisable design knowledge.

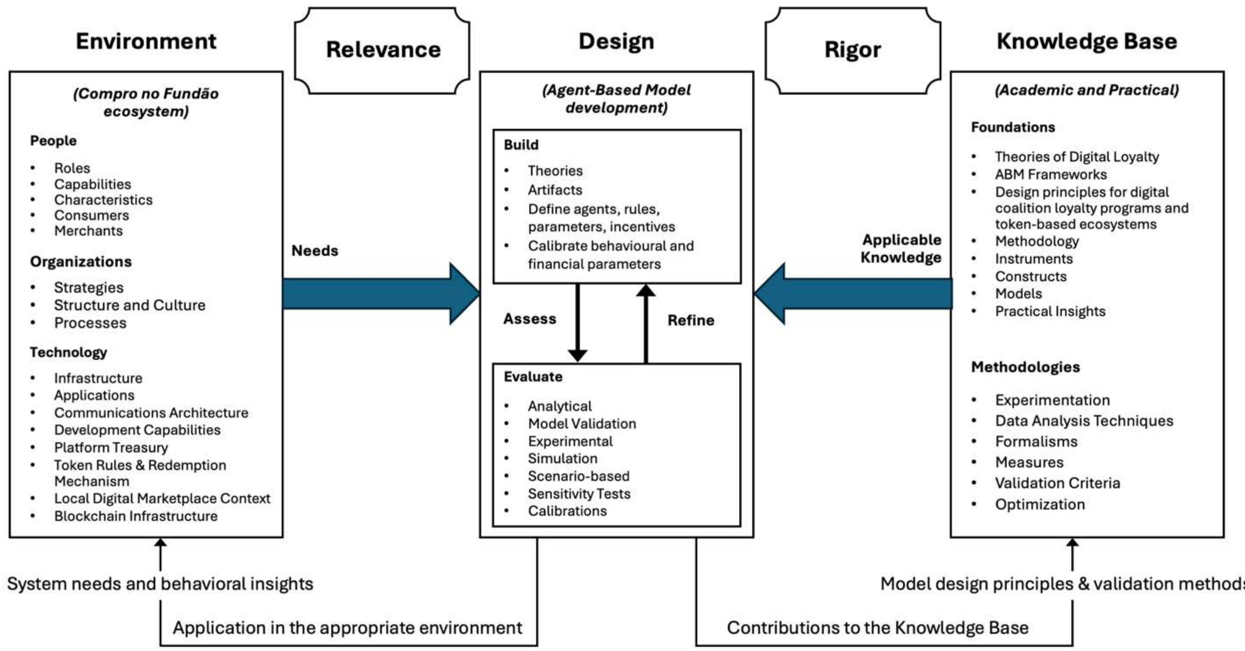

Figure 1. Design Science Research (DSR) Framework Illustrating the Relationship Between Environment, Design Activities, and the Knowledge Base. Adapted from Hevner et al. (2004).

Figure 1 presents the Design Science Research (DSR) framework adapted for this project. The environment represents the Compro no Fundão blockchain-enabled loyalty ecosystem, which connects consumers, merchants, and a central platform treasury. The design cycle captures the iterative development of the Agent-Based Model (ABM) that simulates token issuance, redemption, behavioural interactions, and financial viability, serving as the project's artefact. The knowledge base combines theoretical foundations on digital loyalty and ABM methodologies with practical insights derived from the project's collaboration with Celfocus within the Blockchain.PT initiative, producing design principles for sustainable digital coalition loyalty programs. This cyclical interaction between relevance and rigour ensures that the artefact is both scientifically robust and practically valuable to Celfocus and the broader loyalty-platform literature.

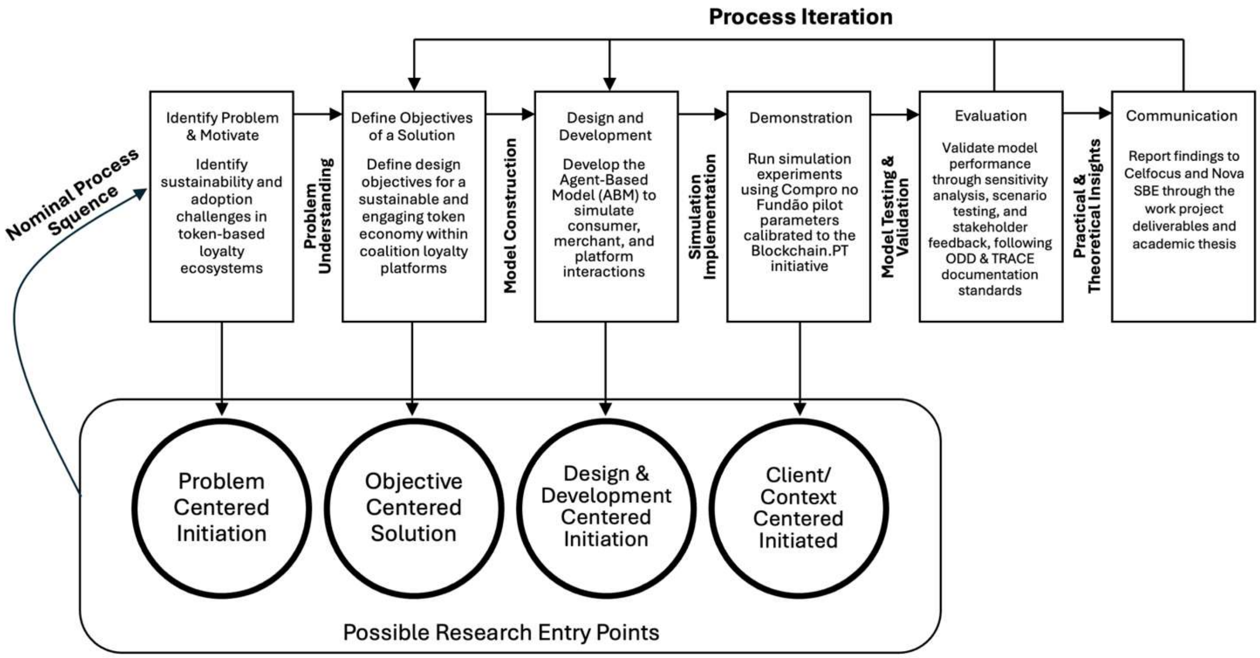

Figure 2. Design Science Research (DSR) Process Model Illustrating the Iterative Sequence from Problem Identification to Artefact Design, Demonstration, Evaluation, and Communication. Adapted from Peffers et al. (2007).

To operationalise this framework, the research follows the six-step Design Science Research Methodology (DSRM) proposed by Peffers et al. (2007), as illustrated in Figure 2, which outlines the Design Science Research (DSR) process applied in this project. The process begins by identifying the core challenges of designing and sustaining value creation within complex, blockchain-enabled coalition loyalty ecosystems. The design objectives are translated into the development of an Agent-Based Model (ABM) that simulates the behavioural and financial interactions among consumers, merchants, and the platform treasury. The artefact is demonstrated through simulation experiments using parameters calibrated to the Compro no Fundão pilot within the Blockchain.PT initiative and evaluated through sensitivity analysis, scenario testing, and TRACE-based documentation. The outcomes are communicated to both Celfocus and Nova SBE, ensuring relevance for industry practice and academic contribution.

4.1 Problem Definition

The core problem addressed in this project is how to design sustainable digital loyalty ecosystems capable of balancing participation, profitability, and platform stability. Traditional analytical approaches overlook the dynamic and interdependent behaviour of multiple stakeholders. The Compro no Fundão pilot illustrates this complexity: merchants define their own reward campaigns (point-to-euro conversion, redemption limits, and maximum discounts), customers accumulate points usable across all merchants, and a central treasury manages point issuance and merchant top-ups. The research treats this environment as a complex adaptive system in which agents respond to incentives, liquidity constraints, and perceived fairness.

By applying ABM within the DSR framework, we can simulate emergent outcomes, such as token inflation, merchant exit, or imbalances between issuance and redemption, that would not be visible through equilibrium-based models. The objective is not only to replicate the Fundão system but also to derive design principles for resilient, scalable coalition-loyalty programs.

4.2 Objectives of the Solution

The methodological aim is to design, implement, and test an ABM that explores behavioural and economic dynamics within coalition loyalty ecosystems, using the Fundão pilot as its empirical foundation. The study will analyse four interconnected dimensions: platform sustainability, customer adoption and engagement, merchant strategies, and financial viability. Together, these dimensions represent the full life cycle of a loyalty platform from incentive design to economic sustainability.

They were selected because each captures a different but complementary layer of system behaviour. Platform sustainability focuses on token value stability, liquidity, and issuance/redemption ratios; customer adoption and engagement measure how reward structures influence engagement and redemption timing; merchant strategies assess how heterogeneous cost-sharing and promotional rules affect coalition stability; and financial viability evaluates whether all stakeholders (customers, merchants, and platform) can achieve positive net outcomes. The four lenses ensure the model integrates behavioural realism with financial and operational sustainability.

Although specific objectives and hypotheses will be finalised during model calibration, the initial purpose is to construct a simulation environment capable of jointly testing these dimensions, identifying trade-offs and reinforcing mechanisms among them.

4.3 Artefact Design and Development

The artefact will be built iteratively through DSR's build-and-evaluate cycles, informed by the operational rules of the Compro no Fundão program. The base scenario will replicate its key parameters:

- Earning rule: Points awarded based on the transaction amount (non-integers are rounded down).

Table 1. Tiered Point Award Rules

Transaction Value	Rule	Points
0,00€ - 4,99€	No points awarded	0
5,00€ - 20,00€	Transaction value / 5	1 to 4
20,01€ - 35,00€	Transaction value / 4	5 to 8
35,01€ - 60,00€	Transaction value / 3	11 to 20
>= 60,01€	20	20

Redemption logic: Store-specific campaigns define (i) minimum spend for redemption, (ii) point-to-euro conversion (e.g., 1 pt = €0.20), and (iii) maximum discount per purchase.

1. Cross-merchant use: Points earned in one merchant can be redeemed in any other participating merchant.

2. Merchant wallets: Merchants hold finite point balances and must top up when depleted (manual in pilot; automated in model).
3. Client wallets: No maximum accumulation limit; usage constrained only by per-transaction caps.
4. Operational context: No POS integration in pilot; transactions and points flows are recorded digitally via a web application.

The ABM will be implemented in Python (Mesa framework) for modularity and transparency. Agents will include consumers, merchants, and a platform treasury, each following decision rules derived from behavioural economics and marketing theory. Consumers aim to maximise perceived value under bounded rationality; merchants adjust campaign aggressiveness to balance customer acquisition against token costs; the platform ensures liquidity by issuing or reclaiming tokens based on system-wide conditions.

Development proceeds in three phases. First, a conceptual model defines entities, variables, and interaction logic grounded in both literature and Fundão's business rules. Second, a computational model formalises these rules into code, ensuring consistency with real pilot parameters. Third, simulation experiments test configurations under controlled conditions, allowing observation of how micro-level decisions (e.g., redemption timing, merchant exit) aggregate into macro-level outcomes (token velocity, ROI). Each iteration will be validated through stakeholder workshops with Celfocus and academic supervisors, ensuring fidelity to both theory and practice.

This iterative development approach reflects the dual structure of Design Science Research, connecting the problem and solution spaces described by Drechsler and Hevner (2018) and summarised in vom Brocke, Hevner, and Maedche (2020).

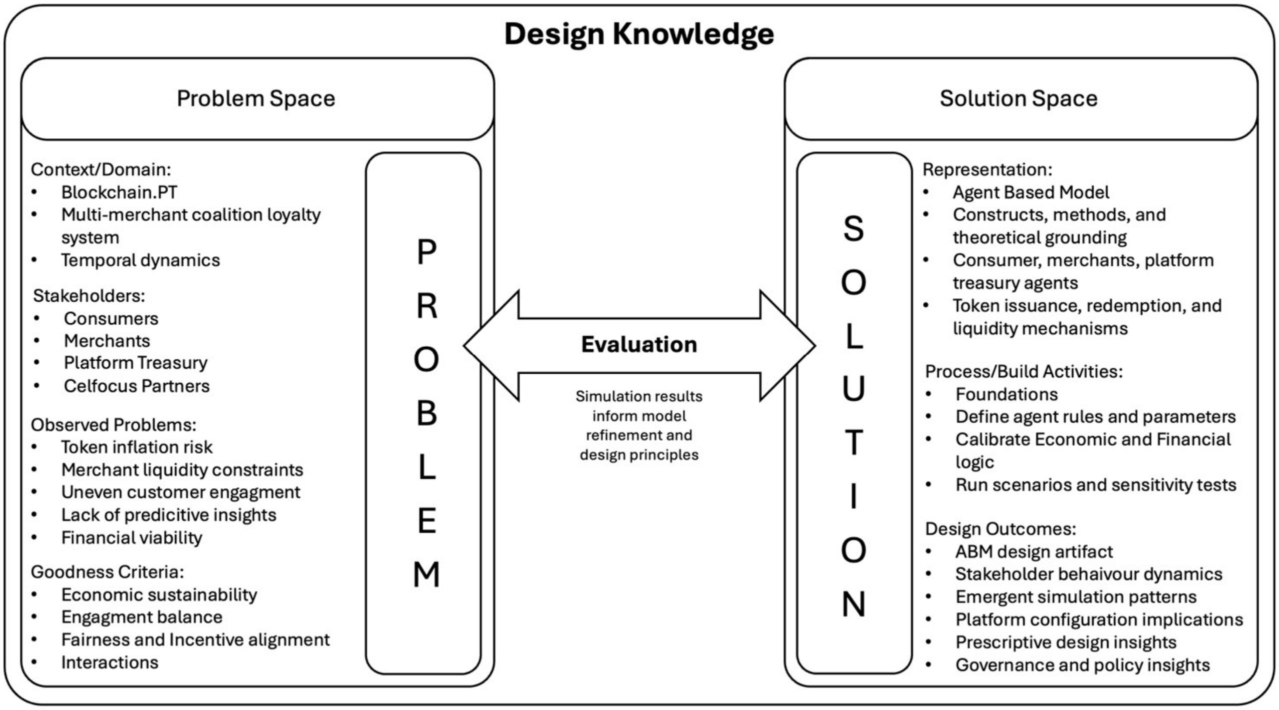

Figure 3. Problem–Solution Space Linking Challenges in Token-Based Loyalty Systems to the Agent-Based Model (ABM) Artefact as the Proposed Solution. Adapted from Drechsler and Hevner (2018), simplified by vom Brocke et al. (2020).

Figure 3 illustrates the connection between the problem and solution spaces within the Design Science Research (DSR) framework as applied in this project. The problem space represents the key challenges observed in blockchain-enabled coalition loyalty ecosystems, including token inflation risk, merchant liquidity constraints, uneven customer engagement, and limited predictive insights for long-term financial viability. The solution space corresponds to the development of the Agent-Based Model (ABM), which simulates consumer, merchant, and platform-treasury interactions to examine alternative incentive, redemption, and liquidity configurations. The iterative evaluation cycle links both spaces, allowing simulation results to inform model refinement and generate prescriptive design and governance insights for sustainable digital coalition loyalty platforms. This approach ensures that the artefact remains grounded in the real context of the Compro no Fundão pilot while producing insights relevant to Celfocus and the broader Design Science Research community.

4.4 Demonstration and Experiment Design

The demonstration phase will run controlled experiments on the ABM, varying platform and campaign parameters to assess their effects on participation, platform stability, and profitability. Scenarios will include both baseline Fundão settings and hypothetical stress tests (e.g., altered reward rates, increased redemption caps, or merchant exit shocks). Each configuration will be replicated across multiple runs to average out stochastic effects.

Simulation outputs will include:

- Agent-level data: activation timing, engagement frequency, redemption rates, wallet balances, and exit decisions for customers; campaign adjustments, liquidity levels, and top-up frequency for merchants.
- System-level indicators: token issuance-to-redemption ratio, token liquidity, active participation rates, merchant exit rate, ROI per stakeholder, and inequality measures (e.g., Gini coefficient)

These results will be visualised through dashboards to facilitate managerial interpretation, connecting agent behaviour to measurable platform KPIs.

4.5 Evaluation Strategy

The evaluation strategy integrates DSR's rigour and relevance balance with the TRACE-based documentation standard (Grimm et al. 2014), ensuring transparency and traceability between model design, calibration, and results interpretation. Evaluation will occur continuously throughout development.

The evaluation approach follows the ex-ante and ex-post structure proposed by Sonnenberg and vom Brocke (2012), as illustrated in Figure 4. Ex-ante evaluation establishes the artefact's theoretical and technical validity before experimentation, while ex-post evaluation assesses its empirical performance and stakeholder relevance after simulation. This iterative cycle ensures that the ABM's design principles are continually refined based on evidence.

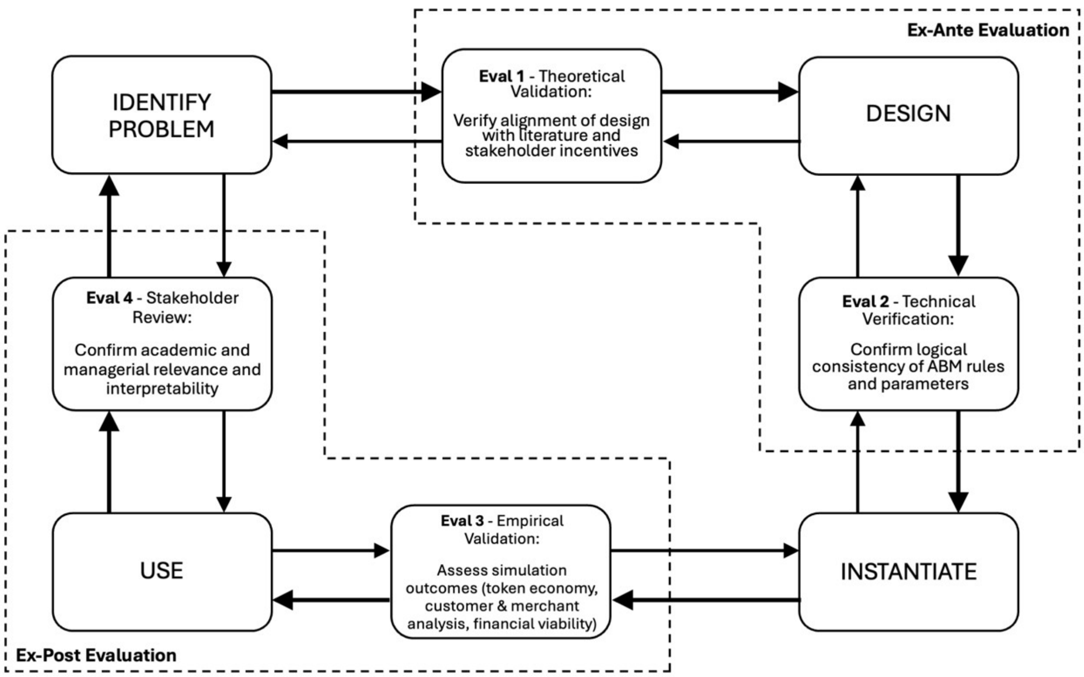

Figure 4. Ex-Ante/Ex-Post Evaluation Cycle Illustrating the Iterative Refinement of the ABM through Four Validation Stages (Eval 1–4) - Adapted from Sonnenberg and vom Brocke (2012)

Figure 4 illustrates the iterative evaluation logic within the DSR framework as applied in this project. The cycle distinguishes between ex-ante and ex-post evaluations, emphasising continuous validation of the ABM throughout its design, implementation, and use. During ex-ante evaluation, the model's conceptual and technical soundness is assessed through theoretical alignment and rule consistency checks. The ex-post evaluation phase assesses behavioural realism and managerial relevance using simulation results, sensitivity analyses, and stakeholder feedback. This cyclical process ensures that each evaluation stage refines the artefact, strengthening both its scientific rigour and its practical contribution within the Blockchain.PT initiative in collaboration with Celfocus.

- Design validity: Verified through iterative expert feedback from Celfocus and Nova SBE advisors, confirming that model assumptions mirror Fundão's operational reality.
- Behavioural realism: Tested using simple but reliable methods such as sensitivity analysis, boundary testing, and multi-run averaging to check that outcomes remain consistent across plausible variations in consumer and merchant behaviour.
- Utility and interpretability: Assessed by the model's ability to generate actionable insights, such as optimal campaign parameters, liquidity thresholds, or early warning indicators of token inflation, that inform real-world platform management.

This strategy ensures that the ABM is both analytically grounded and managerially useful, serving as a prototype decision-support tool for the future expansion of token-based loyalty systems within the Blockchain.PT initiative.

4.6 Expected Contributions

Methodologically, this project illustrates how combining Design Science Research with an Agent-Based Model enables a structured exploration of socio-technical ecosystems. In practice, it delivers a simulation framework that Celfocus can use to pilot and refine the "Compro no Fundão" platform. The ABM thus serves two purposes: (i) as an academic artefact contributing to the literature on digital-platform design, and (ii) as an operational tool for testing loyalty strategies before large-scale deployment.

5 Evaluating the Impact of Campaign Review Frequency on Merchant Utility and Retention in a Coalition Loyalty Program

5.1 Introduction

Coalition loyalty programs (CLPs) create an environment in which independent merchants collectively contribute to, and benefit from, a shared pool of consumer engagement. By participating in a cross-merchant ecosystem of earnable and redeemable points, businesses gain access to incremental footfall, enhanced customer loyalty, and peer-driven network effects that would be difficult to achieve individually. However, this collaborative structure also introduces challenges of coordination and governance. Because participating merchants retain complete discretion to adjust the generosity and structure of their campaigns, the platform becomes vulnerable to instabilities arising from poorly timed or excessively frequent campaign changes.

The ABM developed in this study captures this governance tension through a simplified representation of merchant behaviour. Each enrolled merchant maintains a single active campaign whose aggressiveness is encoded through discrete archetypes. Periodically, merchants evaluate their recent CLP performance and may transition to adjacent archetypes according to a utility-based decision rule. The primary design factor that influences this behaviour is the interval for reviewing campaigns, which determines the minimum number of simulation steps required between consecutive evaluations of merchants.

Merchant utility in the model evolves gradually because it aggregates multiple past transactions through rolling windows and EMAs. As a result, very frequent review cycles cause merchants to adjust campaigns before updated performance information has fully propagated into utility, creating repeated downgrades based on lagged signals. Conversely, very slow reviews keep merchants locked in underperforming campaigns for prolonged periods, increasing the likelihood of exit. Campaign changes also influence consumer behaviour indirectly through savings expectations and earn/redeem attractiveness, creating consumer-mediated interactions among merchants.

5.1.1 Research Question

Against this backdrop, the central research question for this section is:

How does campaign review frequency affect merchant utility and retention in a coalition loyalty program, and how do these effects depend on platform maturity?

To answer this, two controlled experiments were conducted: a baseline condition with 30% initial merchant enrolment, representing a developing coalition, and a robustness condition with 60% enrolment, representing a more mature, denser ecosystem.

5.2 Experimental Design and Methodology

5.2.1 Independent Variable: Review Interval

The experiment manipulates a single platform-level parameter: the number of simulation steps between merchant campaign reviews. Five review intervals were tested: 1, 7, 15, 30, and 60 steps, which correspond to daily, weekly, bi-weekly, monthly, and bi-monthly review cycles. These values were chosen to span realistic managerial planning horizons while also enabling the identification of potential nonlinearities or discontinuities in performance across very fast and very slow review regimes.

5.2.2 Robustness Conditions

To evaluate how coalition density affects review interval outcomes, two initial enrolment conditions were tested: 30% (baseline) and 60% (robustness). Higher enrolment increases CLP transaction volume per merchant, leading to more stable utility estimates and reduced noise in performance signals.

5.2.3 Stochasticity Control via Common Random Numbers

All conditions used the same set of 30 random seeds (common-random-numbers), ensuring each interval is tested in identical stochastic environments to improve statistical power.

5.2.4 Simulation Horizon and Replications

Each simulation runs for 720 steps (representing ~24 months), and each interval × enrolment combination is replicated 30 times, each run with a specific random seed.

5.2.5 Outcome Variables

The two primary outcomes examined are:

- **merchant utility,** a composite indicator capturing uplift, penetration, discount efficiency, wallet trend, and cash position.
- **merchant retention**, defined as the proportion of merchants that join the CLP and remain active at the end of the simulation.

These outcomes are evaluated using standard descriptive statistics, supplemented by one-way ANOVA and Tukey HSD tests to assess differences across review intervals; partial η^2 and Cohen's f are used to quantify effect sizes (Appendix Tables 7, 8, and 9).

Secondary outcomes are included in the appendices to contextualise the system-wide implications of review intervals. These include the Gini coefficient of merchant wallet balances, capturing inequality in CLP benefit distribution; the redemption rate, indicating consumer engagement with the program; the number of consumers registered in the CLP; and the number of consumers churned (Appendix Figures 8 to 12). These variables are interpreted qualitatively; their primary purpose is to aid understanding of the mechanisms at play. A summary of the experimental structure is presented in Table 2.

Table 2. Review interval experiment design summary

Element	Main Experiment	Robustness Check	Purpose
Focal factor	Review interval		Test governance effects
Levels	1, 7, 15, 30, 60		Daily to bi-monthly
Runs per interval	30		Stable estimates
Seeds	1 per run [43-72]		Common random numbers
Total runs	150		5 intervals × 30 runs
Initial enrolment	30%	60%	Test density effects
Outcome focus	Merchant utility, retention	Same + robustness	Platform performance

5.3 Results: 30% Initial Enrolment

5.3.1 Overall Patterns in Merchant Utility

The results for the baseline enrolment scenario reveal a distinct ranking among review intervals, although the noisy utility signal at low CLP density limits statistical significance. Table 3 summarises the mean utility and retention observed under each review interval, and Figure 5 displays the mean utility and corresponding confidence intervals.

Table 3. Merchant Utility and Retention (30% Enrolment)

Interval	*Mean Utility*	*95% CI Low*	*95% CI High*	*Retention*
1	0.1146	0.0884	0.1407	*49.6%*
7	0.1907	0.1594	0.2221	*54.9%*
15	0.1719	0.1382	0.2055	*51.7%*
30	0.1662	0.1396	0.1929	*46.3%*
60	0.1382	0.1035	0.1730	*31.8%*

Figure 5. Mean Merchant Utility by Review Interval (30% Enrolment)

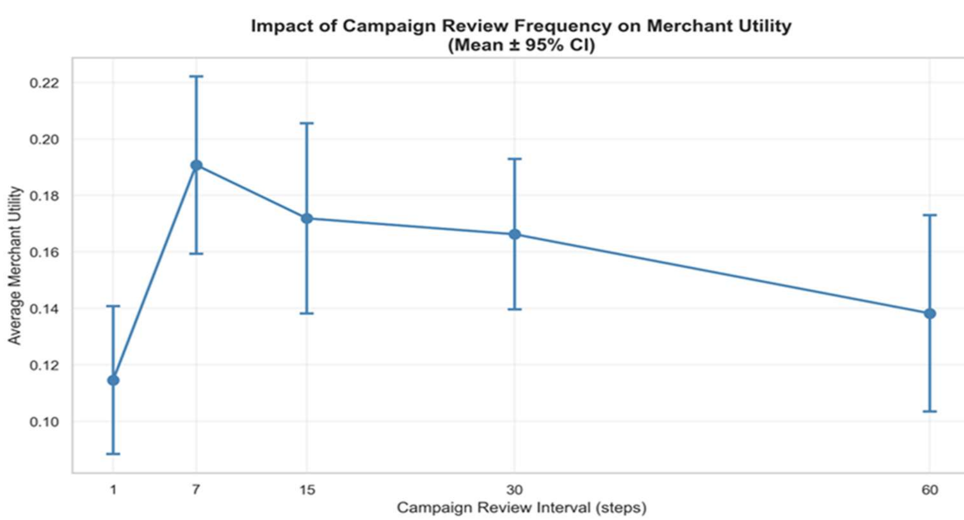

The weekly review (7 steps) achieves the highest performance on both metrics, while the daily review (1 step) produces the lowest utility and the second-lowest retention. Bi-weekly and monthly reviews form a middle-performance group. Notably, the bi-monthly review (60 steps) shows retention of only 31.8%, significantly below that of all other intervals.

5.3.2 Statistical Analysis (ANOVA and Tukey HSD)

A one-way ANOVA tested whether campaign review frequency significantly affects merchant utility: $\mathbf{F\,(4, 145) = 3.98,\ p = 0.0043,\ \eta^2_p = 0.099}$

This result is statistically significant, indicating that review frequency has a detectable effect on merchant utility in the base model condition. The partial eta-squared of 0.099 indicates that review frequency explains approximately 10% of the variance in merchant utility.

Tukey HSD post hoc pairwise comparisons reveal that only the comparison between intervals 1 and 7 is statistically significant. However, all pairwise comparisons involving interval 1 show consistent directional patterns: interval 1 performs worse than all others, with mean differences ranging from 0.023 to 0.076. (Appendix Table 8).

The confidence intervals for intervals 7, 15, and 30 show substantial overlap, reflecting relatively high variance within each group at this enrolment density.

5.3.3 Retention Outcomes

Merchant retention rates reveal more apparent differentiation across intervals than utility does. Retention is measured as the proportion of merchants that joined the CLP at any point during the simulation and remained enrolled at step 720. The model's exit logic evaluates merchants periodically based on accumulated "bad periods" (evaluation windows during which realised utility falls below a threshold), with the exit probability increasing as these bad periods accumulate.

The pattern organises into three tiers:

- Highest retention: Interval 7 (54.9%) significantly surpasses the average retention (46.9%) and achieves the highest retention among all tested intervals.
- Moderate retention: Intervals 1 (49.6%) and 15 (51.7%) constitute a middle tier, with retention rates near the overall average.
- Lowest retention: Intervals 30 (46.3%) and notably 60 (31.8%) show increasingly lower retention, with bi-monthly reviews resulting in exit rates nearly double those of weekly reviews.

An important observation emerges when comparing retention to utility: interval 7 achieves the highest retention (54.9%) and highest utility (0.191), while interval 60 achieves the lowest retention (31.8%) despite not having the lowest utility. This suggests that infrequent reviews primarily harm retention by prolonging exposure to underperforming campaigns, allowing merchants to accumulate enough consecutive bad periods to trigger exit before they have opportunities to adjust their strategies.

5.3.4 Summary for 30% Enrolment

The baseline experiment at 30% initial enrolment establishes three clear directional patterns: Weekly review outperforms all other intervals on both metrics. Daily review consistently underperforms, consistent with the hypothesis that very frequent reviews create self-amplifying adjustment loops on lagged signals. Bi-monthly review produces catastrophically low retention, demonstrating that extremely slow review intervals trap merchants in underperforming campaigns long enough to trigger exit.

While the ANOVA is statistically significant, only one pairwise comparison (interval 1 vs. 7) is significant in post hoc testing. Nevertheless, the consistent rank ordering and clear separation between extreme and moderate intervals suggest genuine governance effects that will become more pronounced under higher-density conditions.

The limited statistical power at 30% enrolment reflects sparse CLP activity: fewer transactions per merchant produce noisier utility windows and EMAs, wider confidence intervals, and reduced ability to discriminate between similar review frequencies. This noise obscures differences among moderate intervals (7, 15, 30) despite their directional consistency.

5.4 Results: 60% Initial Enrolment

5.4.1 Utility and Retention Patterns

The robustness scenario with 60% initial enrolment presents much clearer distinctions among review intervals. These results are visible in Table 4 and Figure 6.

Table 4. Merchant Utility and Retention (60% Enrolment)

Interval	*Mean Utility*	*95% CI Low*	*95% CI High*	*Retention*
1	0.1110	0.0893	0.1326	*48.1%*
7	0.2055	0.1835	0.2275	*70.9%*
15	0.2339	0.2157	0.2521	*62.4%*
30	0.2312	0.2118	0.2507	*55.1%*
60	0.1929	0.1694	0.2164	*54.7%*

Figure 6. Mean Merchant Utility by Review Interval (60% Enrolment)

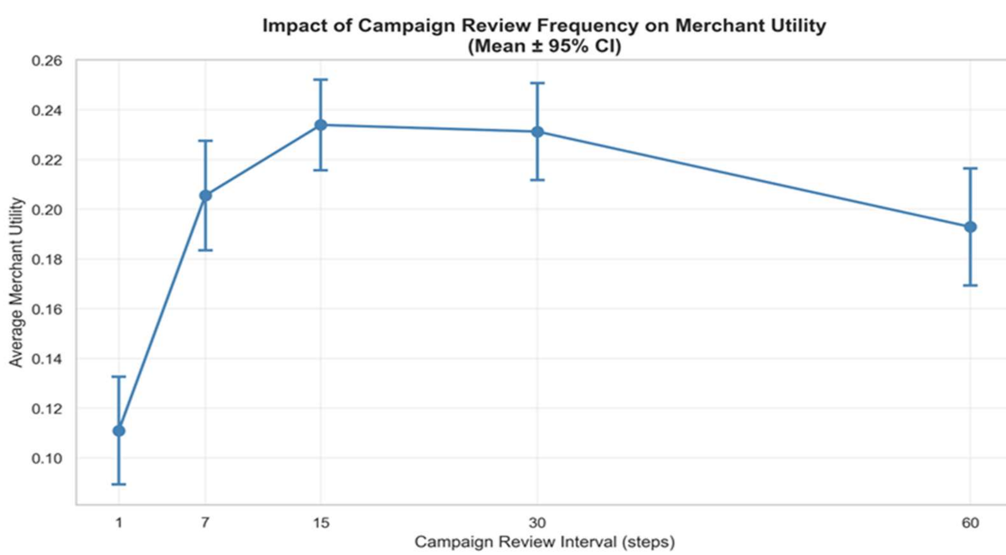

Bi-weekly and monthly reviews now achieve the highest utilities at 0.234 and 0.231, respectively. In comparison, the weekly review maintains the top retention rate of 70.9% but ranks third in utility. Daily review continues to perform the worst on both metrics. The bi-monthly review shows notable retention improvement from 30% enrolment (54.7% vs. 31.8%), but it still underperforms moderate intervals.

Two critical findings emerge: First, the rank ordering by utility is nearly preserved from 30% enrolment, with this consistency across density conditions strongly suggesting genuine effects rather than random variation. Second, the magnitude of effects increases substantially—the utility difference between worst and best intervals nearly doubles from 0.076 to 0.123, as seen in appendix tables 20 and 21, representing an improvement of 111% relative to daily review compared to 66% at lower density.

5.4.2 ANOVA and Pairwise Tests

At 60% enrolment, the ANOVA on merchant utility reveals dramatically stronger effects:

$$\mathbf{F\,(4,\,145) = 23.61,\ p < 10^{-15},\ \eta^2_p = 0.394}$$

Compared to the 30% enrolment condition ($F = 3.98$, $p = 0.0043$, $\eta^2_p = 0.099$), the higher-density scenario shows: (1) F-statistic increases by a factor of 5.9, (2) p-value decreases from 0.004 to effectively zero ($p < 10^{-15}$), and (3) effect size quadruples from $\eta^2_p = 0.099$ to $\eta^2_p = 0.394$. Campaign review frequency now explains approximately 39% of the variance in merchant utility, demonstrating that governance effectiveness becomes substantially stronger as coalition density increases.

Tukey post-hoc comparisons reveal that the daily review interval is significantly worse than all others, while the 15-step interval significantly outperforms the 60-step interval. Differences among the 7, 15, and 30-step intervals are not statistically significant, indicating a plateau of high performance across this mid-range region.

These findings support a simple but powerful conclusion: at higher merchant density, daily review is significantly worse than any other interval, while all intervals of 7 days or longer perform comparably well in terms of merchant utility. The emergence of a clear "safe zone" of moderate review frequencies (7-30 steps) demonstrates that governance policies have amplified consequences in mature coalitions.

5.4.3 Interpretation of Density Effects

The apparent emergence of statistical significance at 60% enrolment, despite the base model at 30% showing only moderate effects, can be attributed to three mechanisms operating simultaneously:

- Improved precision through higher transaction volume: Higher enrolment increases the frequency of CLP transactions per merchant, improving the reliability of utility estimates and narrowing confidence intervals.
- Amplified consumer-mediated feedback: Higher enrolment accelerates consumer CLP adoption through stronger savings signals, creating more pronounced feedback loops between campaign aggressiveness, consumer registration and usage, and merchant utility.
- Critical mass activation of platform dynamics: Below approximately 50-60% enrolment, the CLP behaves more like a collection of independent merchants with a shared currency. Consumers encounter CLP opportunities sporadically, wallet dynamics remain weak, and competitive effects are minimal. Once enrolment crosses this threshold, coalition effects fully activate: consumers find CLP opportunities in routine purchases, making loyalty behaviours more meaningful, and campaign decisions acquire strategic significance. In this regime,

campaign review frequency interacts with activated network effects rather than isolated merchant behaviours, magnifying the importance of governance choices.

Together, these mechanisms explain why the same parameter (review frequency) appears modestly significant at 30% enrolment but becomes highly significant at 60%. The effects are not just easier to measure, they actually become larger in magnitude as coalition density increases, demonstrating that governance effectiveness itself exhibits scale-dependent properties.

5.5 Mechanism Analysis

5.5.1 Why Daily Review Underperforms: Lagged Signals and Repeated Downgrades

Understanding these patterns involves examining the interaction between review timing and merchants' utility signals. Daily reviews consistently yield lower utility and retention across density conditions. This is a mechanical effect: merchant utility is calculated over moving windows that update slowly compared to the daily review cycle, creating a timing mismatch between decision and information cycles.

Mechanism 1: Slow-moving utility windows

Merchant performance is tracked through rolling windows of recent transactions. Utility is constructed from multiple weighted components including sales uplift, CLP penetration, discount efficiency, wallet trend, and cash position.

Campaign reviews occur at the specified interval and evaluate recent performance to determine whether to transition between archetypes. However, the utility signal itself is a slow-moving average over a fixed evaluation window. When a merchant changes campaigns, the utility signal in the next period remains dominated by the same transaction history—the window has not had time to incorporate enough activity under the new campaign configuration.

With daily reviews, merchants re-evaluate every step. If utility is below the transition threshold and they downgrade their campaign, they will likely observe "utility still low" in subsequent periods because the rolling window changes slowly. The decision rule applies repeatedly: low wallet balance and poor utility trigger further downgrades toward more conservative campaigns. This produces consecutive downgrades before the effects of any single adjustment can propagate into the performance data, creating a self-amplifying drift toward extreme conservative archetypes.

Mechanism 2: Early-phase structural underperformance

At simulation start, few consumers are registered, and CLP penetration is low. Merchants incur point-funding costs while not yet seeing large incremental CLP sales. The uplift and efficiency components of utility are, therefore, systematically low or negative during this startup period. If campaign evaluations are applied too early and too frequently during this phase, merchants accumulate consecutive periods of poor performance before the CLP reaches maturity. This increases both campaign downgrades and exit events during a transient phase that would eventually be resolved if merchants waited longer.

Very frequent reviews systematically expose merchants to negative utility windows before coalition dynamics stabilise, leading to early exits and permanent campaign degradation that persists even after the system matures.

5.5.2 Why Moderate Frequencies (7-30 steps) Succeed

Moderate review frequencies, particularly 7 and 15 steps, achieve the best outcomes by balancing three competing requirements: responsiveness, stability, and evidence accumulation.

- Sufficient time for signal propagation: Weekly and bi-weekly intervals provide enough steps for campaign changes to affect CLP sales, for those sales to feed into rolling windows, and for utility components (uplift, penetration, efficiency) to update before the following review.

Merchants act on performance data that reflects their current campaign, not stale historical windows.

- Adequate responsiveness: Intervals of 7-30 steps still allow merchants to escape persistently unprofitable configurations. If a campaign genuinely underperforms, merchants can adjust within 1-4 weeks rather than remaining trapped for months. This prevents prolonged accumulation of consecutive poor performance periods that would trigger exit.
- Reliable evidence accumulation: The model's decision rules rely on moving averages and windowed statistics. Weekly to monthly intervals ensure that evaluation windows contain sufficient CLP transactions to compute meaningful averages. In sparse coalitions (30% enrolment), this helps reduce noise; in dense coalitions (60% enrolment), it stabilises utility trajectories and prevents volatility-driven overreactions.

The 60% enrolment results show that intervals of 7 days or more are statistically indistinguishable in terms of utility. This suggests a threshold effect: once the review interval is long enough to dampen step-to-step noise (roughly one week), additional lengthening to 15, 30, or even 60 days yields diminishing returns in aggregate utility.

5.5.3 Why Very Infrequent Review (60 steps) Weakens Performance

Bi-monthly review (60 steps) produces intermediate utility but substantially lower retention, especially at 30% enrolment. The mechanism is the inverse of the daily review problem: instead of overreacting to lagged signals, merchants become trapped in underperforming campaigns for extended periods.

The exit mechanism evaluates merchants periodically based on accumulated periods of poor performance. Each evaluation window where performance falls below the exit threshold increments this counter, and exit becomes probable once the threshold is reached. When campaigns are reviewed infrequently (every 60 steps), merchants with overly conservative or poorly calibrated campaigns endure multiple consecutive evaluation windows below threshold before being allowed to adjust. This increases the likelihood of meeting the exit condition even if later CLP conditions would have supported recovery.

At 60% enrolment, the retention penalty of very slow review partially recovers (54.7% vs. 31.8% at 30% enrolment) because higher transaction volume provides more data points per evaluation window, making utility estimates more stable and reducing the frequency of "false negative" low-utility periods. However, even at high density, interval 60 underperforms moderate intervals (7-30) in both utility and retention, indicating that very infrequent reviews lead to avoidable merchant exits.

5.6 Theoretical Implications

Three theoretical implications emerge from these findings. First, governance effectiveness is fundamentally context-dependent: the impact of campaign review frequency depends critically on merchant density and platform maturity, with effect sizes quadrupling from 0.099 to 0.394 between enrolment conditions. These challenges simplified governance models that treat policy parameters as having fixed effects regardless of system state. Second, governance mechanisms themselves exhibit network-like properties that amplify with coalition density through consumer-mediated externalities, extending the concept of network effects from demand-side adoption to supply-side governance and suggesting that policy interventions can be subject to increasing returns and critical mass thresholds. Third, the analysis reveals an emergent "stability zone" of moderate review frequencies (7-30 steps) where performance remains robust across a 4-fold range, while extreme frequencies—whether too fast or too slow—degrade outcomes sharply. This pattern supports theories of bounded rationality in organisational decision-making, where moderate frequencies provide a balance between signal quality, responsiveness, and stability.

5.7 Practical Recommendations

For platform operators designing and managing coalition loyalty programs, the findings support several actionable recommendations:

5.7.1 Discourage Hyper-Frequent Review Cycles

Daily or near-daily campaign review cycles consistently produce the worst outcomes across all tested conditions. Platform designs should avoid features that implicitly encourage such behaviour and instead guide merchants toward weekly or longer review cycles.

5.7.2 Promote the 7-30 Step Review Zone

Weekly to monthly review cycles (7-30 steps) form a robust, high-performance zone where merchants achieve strong utility and retention while supporting healthy consumer adoption. Within this range, weekly review optimises retention, bi-weekly review optimises utility, and monthly review offers comparable performance. Platform operators can offer flexibility based on merchant-specific considerations and implement this guidance through recommended schedules or default settings.

5.7.3 Adapt Governance as Platform Matures

Governance strategies should evolve as coalition density increases. In early-stage coalitions (30% enrolment), operators should prioritise growing the coalition and maintaining stability rather than aggressively optimising campaign review policies. As the platform approaches 50-60% merchant participation, governance effectiveness amplifies, justifying investment in more sophisticated policy design. The key insight is that the same interventions have larger system-wide consequences once critical mass is achieved, making governance optimisation both more important and more cost-effective in mature coalitions.

5.8 Limitations

Several limitations must be acknowledged. The model employs significant simplifications in merchant and consumer behaviour, lacks empirical calibration, and tests a limited scenario range. These constraints affect both internal validity (behavioural assumptions) and external validity (generalisability). Despite these limitations, the experiment uses 30 replications per condition with common random numbers, and the main findings are consistent across both density levels, strengthening internal validity within the model's scope.

5.9 Future Research Directions

This analysis opens several promising directions for future research. First, systematic density sweeps across the full enrolment spectrum (10-90%) would identify critical mass thresholds and map how governance effectiveness scales with coalition maturity. Second, introducing strategic merchant behaviour through learning algorithms or game-theoretic reasoning would test whether the stability zone persists under competitive interactions. Third, examining interactions with other governance levers—such as fee structures, point accrual rates, or redemption restrictions—would support integrated policy optimisation. Fourth, extending simulation horizons to 5-10 years would capture slow dynamics, including multi-year. Future research could explore systematic density sweeps to identify critical mass thresholds, introduce strategic merchant behaviour to test stability under competition, examine interactions with other governance levers, extend simulation horizons to capture longer-term dynamics, and pursue empirical validation with operating coalitions.

6 Agent-Based Insights in a Coalition Loyalty Program Context

Viewed as a complex adaptive system, the coalition loyalty platform examined in this study exhibits outcomes that emerge from repeated interactions among consumers, merchants, and the platform rather than from isolated design parameters. Although Sections 5 through 8 examine token economics, consumer engagement, merchant strategies, and financial viability separately, their joint interpretation reveals system-level dynamics that are not visible within any single lens.

A central insight is that long-term sustainability depends on alignment across behavioural, strategic, and financial dimensions. High-incentive token designs stimulate rapid consumer adoption, frequent redemption, and strong circulation, reinforcing engagement and short-term performance. However, when issuance persistently exceeds redemption capacity, these same dynamics generate accelerating liabilities and inflationary pressure. Conversely, restrictive incentive regimes limit financial exposure but suppress participation, weaken circulation, and increase the risk of stagnation. Sustainability emerges from maintaining a balance between engagement intensity and the platform's financial absorption capacity and cannot be captured through any single KPI.

Consumer-side dynamics show that incentive policies primarily shape the quality and persistence of engagement, rather than as levers that reliably scale the member base. Tiered earning rules display an asymmetry: once a baseline structure makes the value proposition salient, additional generosity yields only modest gains in adoption and activity, whereas more restrictive tiering reduces participation and increases churn. One-off onboarding bonuses' leverage over adoption is questionable, but they can meaningfully shift engagement by increasing redemption frequency in the medium term. These consumer-facing promotions can redistribute value within the ecosystem, improving user-side outcomes while reducing the platform's ability to monetise point sales.

These consumer-side dynamics directly condition the economic environment faced by merchants. Merchant behaviour reinforces this trade-off observed at the system level: adaptive campaign strategies and participation decisions respond endogenously to realised utility, wallet dynamics, and local adoption patterns shaped by consumer engagement. When demand generated by the program remains sufficiently strong and campaign-related costs are manageable, merchants tend to remain enrolled and adjust strategies to stabilise performance. When cost burdens dominate perceived benefits, exit dynamics emerge even in otherwise active ecosystems. This highlights that merchant retention is not solely a function of aggregate demand, but of how incentives translate into locally realised profitability over time.

From a financial perspective, the results demonstrate that liability accumulation and inflation are emergent outcomes driven by behavioural timing rather than static policy settings. Redemption effectiveness depends not only on nominal reward generosity but on sustained participation, habit formation, and the synchronisation of earning and redeeming flows. Designs that appear viable under static or short-horizon evaluation may therefore destabilise over longer time horizons once adaptive responses unfold.

Taken together, these findings indicate that coalition loyalty platforms require continuous calibration instead of fixed design rules. Incentives must be sufficiently attractive to sustain participation and circulation while remaining constrained enough to prevent uncontrolled liability growth. Sustainability thus emerges as a system-level property shaped by ongoing feedback, adaptation, and coordination among stakeholders, rather than by any single design choice.

7 Conclusion and Future Directions

This work project developed and applied an agent-based model to examine the dynamics of a token-based coalition loyalty platform, using the Compro no Fundão program as a reference context. By representing consumers, merchants, and the platform as interacting agents, the model provides a structured way to explore participation patterns, behavioural adaptation, and financial flows within a shared loyalty ecosystem.

The analyses conducted in this study collectively illustrate how a coalition loyalty platform can be examined as a dynamic system shaped by interaction and feedback rather than as a collection of independent design choices. The model enables the joint consideration of token circulation, consumer engagement, merchant participation, and financial aggregation within a single simulation environment, allowing these dimensions to be analysed consistently under a common set of assumptions.

The purpose of the model is not to identify optimal parameter values or to generate predictive forecasts, but to support structured exploration of how alternative configurations and governance parameters influence system behaviour over time. The findings should therefore be interpreted as scenario-based insights that reflect the mechanisms explicitly implemented in the model, rather than as generalisable empirical claims about coalition loyalty programs.

From a methodological perspective, this project demonstrates the applicability of an agent-based model within a Design Science Research framework to an applied platform setting. The resulting simulation artefact makes it possible to experiment with behavioural rules and institutional constraints while preserving heterogeneity, bounded rationality, and feedback effects that are difficult to capture using static analytical approaches.

Several avenues for further work naturally follow from the scope of the current implementation. Future research could extend the analysis to different coalition maturity levels, incorporate richer forms of merchant decision-making, refine the assumptions used to translate simulation outputs into monetary quantities, and pursue empirical calibration using operational platform data. Within these bounds, this study provides a cautious and internally consistent foundation for analysing coalition loyalty platforms as complex adaptive systems, while remaining closely aligned with the mechanisms and assumptions explicitly encoded in the model.

References

Abadi, Joseph, and Markus Brunnermeier. 2024. "Token-Based Platform Governance." Journal of Financial Economics 162 (December): 103951. https://doi.org/10.1016/j.jfineco.2024.103951.

Abdelkafi, Nizar, and Karl Täuscher. 2016. 'Business Models for Sustainability From a System Dynamics Perspective'. Organization & Environment 29 (1): 74–96. https://doi.org/10.1177/1086026615592930.

Abooleet, Saeed, and Seth Kinnett. 2023. "A Systematic Review of Blockchain-Based Loyalty Programs." Paper presented at Hawaii International Conference on System Sciences. https://doi.org/10.24251/HICSS.2023.650.

Anudeep, Ch., Kathari Santosh, M Sandeep Kumar, Jahnavi M, Sabareesh R, and B Kiran Bala. 2024. "Constructing a Blockchain-Based Loyalty Program Model for Secure and Transparent Customer Rewards in Retail Marketing." 2024 IEEE 13th International Conference on Communication Systems and Network Technologies (CSNT), April 6, 241–47. https://doi.org/10.1109/CSNT60213.2024.10546113.

Arnejo, Zenith, Benoit Gaudou, Mehdi Saqalli, and Nathaniel Bantayan. 2025. "Communicating Agent-Based Models to Stakeholders: A Scoping Review." Journal of Artificial Societies and Social Simulation 28 (2): 2. https://doi.org/10.18564/jasss.5623.

Arthur, W. Brian. 2021. 'Foundations of Complexity Economics'. Nature Reviews Physics 3 (2): 136–45. https://doi.org/10.1038/s42254-020-00273-3.

Axtell, Robert L., and J. Doyne Farmer. 2025. 'Agent-Based Modeling in Economics and Finance: Past, Present, and Future'. Journal of Economic Literature 63 (1): 197–287. https://doi.org/10.1257/jel.20221319.

Bagchi, Rajesh, and Xingbo Li. 2011. "Illusionary Progress in Loyalty Programs: Magnitudes, Reward Distances, and Step-Size Ambiguity." Journal of Consumer Research 37 (5): 888–901. https://doi.org/10.1086/656392.

Belli, Alex, Anne-Maree O'Rourke, François A. Carrillat, Ljubomir Pupovac, Valentyna Melnyk, and Ekaterina Napolova. 2022. "40 Years of Loyalty Programs: How Effective Are They? Generalizations from a Meta-Analysis." Journal of the Academy of Marketing Science 50 (1): 147–73. https://doi.org/10.1007/s11747-021-00804-z.

Berman, Barry. 2006. "Developing an Effective Customer Loyalty Program." California Management Review 49 (1): 123–48. https://doi.org/10.2307/41166374.

Bijmolt, Tammo H. A., and Peter C. Verhoef. 2017. "Loyalty Programs: Current Insights, Research Challenges, and Emerging Trends." In Handbook of Marketing Decision Models, edited by Berend Wierenga and Ralf van der Lans. Springer International Publishing. https://doi.org/10.1007/978-3-319-56941-3_5.

Bijmolt, Tammo H. A., Matilda Dorotic, and Peter C. Verhoef. 2011. "Loyalty Programs: Generalizations on Their Adoption, Effectiveness and Design." Foundations and Trends® in Marketing 5 (4): 197–258. https://doi.org/10.1561/1700000026.

Bonabeau, Eric. 2002. "Agent-Based Modeling: Methods and Techniques for Simulating Human Systems." Proceedings of the National Academy of Sciences 99 (suppl_3): 7280–87. https://doi.org/10.1073/pnas.082080899.

Boukis, Achilleas. 2024. "The Effect of Tokenized Rewards on Customer Loyalty Programs." Annals of Tourism Research 109 (November): 103851. https://doi.org/10.1016/j.annals.2024.103851.

Breugelmans, Els, Tammo H. A. Bijmolt, Jie Zhang, et al. 2015. "Advancing Research on Loyalty Programs: A Future Research Agenda." Marketing Letters 26 (2): 127–39. https://doi.org/10.1007/s11002-014-9311-4.

Brocke, Jan vom, Alan Hevner, and Alexander Maedche. 2020. "Introduction to Design Science Research." In Design Science Research. Cases, edited by Jan vom Brocke, Alan Hevner, and Alexander Maedche. Springer International Publishing. https://doi.org/10.1007/978-3-030-46781-4_1.

Bruneau, Virginie, Valérie Swaen, and Pietro Zidda. 2018. "Are Loyalty Program Members Really Engaged? Measuring Customer Engagement with Loyalty Programs." Journal of Business Research 91 (October): 144–58. https://doi.org/10.1016/j.jbusres.2018.06.002.

Cao, Yuheng, Aaron L Nsakanda, and Moustapha Diaby. 2015. "Planning the Supply of Rewards with Cooperative Promotion Considerations in Coalition Loyalty Programmes Management." Journal of the Operational Research Society 66 (7): 1140–54. https://doi.org/10.1057/jors.2014.81.

Capizzi, Michael T., and Rick Ferguson. 2005. "Loyalty Trends for the Twenty-first Century." Journal of Consumer Marketing 22 (2): 72–80. https://doi.org/10.1108/07363760510589235.

Chainlink. 2024. "Asset Tokenization: Digital Assets Explained." Chainlink. https://chain.link/education/asset-tokenization.

Chen, Yanyan, Timo Mandler, and Lars Meyer-Waarden. 2021. "Three Decades of Research on Loyalty Programs: A Literature Review and Future Research Agenda." Journal of Business Research 124 (January): 179–97. https://doi.org/10.1016/j.jbusres.2020.11.057.

Chun, So Yeon, Dan A. Iancu, and Nikolaos Trichakis. 2020. 'Loyalty Program Liabilities and Point Values'. Manufacturing & Service Operations Management 22 (2): 257–72. https://doi.org/10.1287/msom.2018.0748.

Cong, Lin William, Ye Li, and Neng Wang. 2021. 'Tokenomics: Dynamic Adoption and Valuation'. The Review of Financial Studies 34 (3): 1105–55. https://doi.org/10.1093/rfs/hhaa089.

Cong, Lin William, Ye Li, and Neng Wang. 2022. 'Token-Based Platform Finance'. Journal of Financial Economics 144 (3): 972–9De Noni, Ivan, Luigi Orsi, and Luca Zanderighi. 2014. "Coalition Loyalty-Programme Adoption and Urban Commercial-Network Effectiveness Evaluation." International Journal of Retail & Distribution Management 42 (9): 818–38. https://doi.org/10.1108/IJRDM-05-2012-0043.

Deloitte. 2016. "Making Blockchain Real for Customer Loyalty Rewards Programs." Deloitte Development LLC. https://www.deloitte.com/us/en/Industries/financial-services/articles/making-blockchain-real-customer-loyalty-rewards-programs.html.

Dorotic, Matilda, Dennis Fok, Peter C. Verhoef, and Tammo H. A. Bijmolt. 2011. "Do Vendors Benefit from Promotions in a Multi-Vendor Loyalty Program?" Marketing Letters 22 (4): 341–56. https://doi.org/10.1007/s11002-010-9128-8.

Dorotic, Matilda, Dennis Fok, Peter C. Verhoef, and Tammo H. A. Bijmolt. 2021. "Synergistic and Cannibalization Effects in a Partnership Loyalty Program." Journal of the Academy of Marketing Science 49 (5): 1021–42. https://doi.org/10.1007/s11747-020-00759-7.

Dorotic, Matilda, Peter C. Verhoef, Dennis Fok, and Tammo H. A. Bijmolt. 2014. "Reward Redemption Effects in a Loyalty Program When Customers Choose How Much and When to

Redeem." International Journal of Research in Marketing 31 (4): 339–55. https://doi.org/10.1016/j.ijresmar.2014.06.001.

Dowling, Grahame R., and Mark Uncles. 1997. "Do Customer Loyalty Programs Really Work?" Sloan Management Review 38 (4): 71–82. https://www.researchgate.net/publication/290890144_Do_Customer_Loyalty_Programs_Really_Work.

Drechsler, Andreas, and A. R. Hevner. 2018. Utilizing, Producing, and Contributing Design Knowledge in DSR Projects. January 1. https://doi.org/10.26686/wgtn.13088369.v1.

Evans, David S., and Richard Schmalensee. 2016. 'The New Economics of Multi-Sided Platforms: A Guide to the Vocabulary'. SSRN Electronic Journal, ahead of print, 2016. https://doi.org/10.2139/ssrn.2793021.

EY. 2025. Loyalty Market Study 2025. Market Study / Industry Report. EY Global. https://www.ey.com/content/dam/ey-unified-site/ey-com/en-us/campaigns/cmo/documents/ey-loyalty-market-study-2025.pdf.

Freni, Pierluigi, Enrico Ferro, and Roberto Moncada. 2022. 'Tokenomics and Blockchain Tokens: A Design-Oriented Morphological Framework'. Blockchain: Research and Applications 3 (1): 100069. https://doi.org/10.1016/j.bcra.2022.100069.

Garcia, Rosanna, and Wander Jager. 2011. "From the Special Issue Editors: Agent-Based Modeling of Innovation Diffusion." Journal of Product Innovation Management 28 (2): 148–51. https://doi.org/10.1111/j.1540-5885.2011.00788.x.

Ghanem, Nada, Stephan Leitner, and Dietmar Jannach. 2022. "Balancing Consumer and Business Value of Recommender Systems: A Simulation-Based Analysis." Electronic

Commerce Research and Applications 55 (September): 101195. https://doi.org/10.1016/j.elerap.2022.101195.

Grimm, Volker, Jacqueline Augusiak, Andreas Focks, et al. 2014. "Towards Better Modelling and Decision Support: Documenting Model Development, Testing, and Analysis Using TRACE." Ecological Modelling 280 (May): 129–39. https://doi.org/10.1016/j.ecolmodel.2014.01.018.

Grimm, Volker, Uta Berger, Donald L. DeAngelis, J. Gary Polhill, Jarl Giske, and Steven F. Railsback. 2010. "The ODD Protocol: A Review and First Update." Ecological Modelling 221 (23): 2760–68. https://doi.org/10.1016/j.ecolmodel.2010.08.019.

Hagiu, Andrei, and Julian Wright. 2015. 'Multi-Sided Platforms'. International Journal of Industrial Organization 43 (November): 162–74. https://doi.org/10.1016/j.ijindorg.2015.03.003.

Hevner, Alan R., Salvatore T. March, Jinsoo Park, and Sudha Ram. 2004. "Design Science in Information Systems Research1." MIS Quarterly 28 (1): 75–106. https://doi.org/10.2307/25148625.

Koide, Ryu, Haruhisa Yamamoto, Keisuke Nansai, and Shinsuke Murakami. 2023. "Agent-Based Model for Assessment of Multiple Circular Economy Strategies: Quantifying Product-Service Diffusion, Circularity, and Sustainability." Resources, Conservation and Recycling 199 (December): 107216. https://doi.org/10.1016/j.resconrec.2023.107216.

Li, Jiasun, and William Mann. 2025. 'Digital Tokens and Platform Building'. The Review of Financial Studies 38 (7): 1921–54. https://doi.org/10.1093/rfs/hhaf015.

Lin, Chen, and Douglas Bowman. 2022. "The Impact of Introducing a Customer Loyalty Program on Category Sales and Profitability." Journal of Retailing and Consumer Services 64 (January): 102769. https://doi.org/10.1016/j.jretconser.2021.102769.

Liu, Yun, Bhakti Stephan Onggo, and Jerry Busby. 2024. "Agent-Based Modelling of User Engagement in New Product Development." Technovation 135 (July): 103062. https://doi.org/10.1016/j.technovation.2024.103062.

Luo, Dan. 2024. 'Corporate Finance Through Loyalty Programs'. Preprint, SSRN. https://doi.org/10.2139/ssrn.4971046.

Macal, C M. 2016. "Everything You Need to Know about Agent-Based Modelling and Simulation." Journal of Simulation 10 (2): 144–56. https://doi.org/10.1057/jos.2016.7.

Macal, C., and Michael North. 2009. Agent-Based Modeling and Simulation. In Proceedings of the 2009 Winter Simulation Conference (WSC). https://doi.org/10.1109/WSC.2009.5429318.

Mando-Connect and YouGov. 2024. Understanding Loyalty in Europe 2.0. Mando Agency UK. https://mando.co.uk/wp-content/uploads/2025/03/Mando-Understanding-Loyalty-in-Europe-2.0.pdf.

Meyer-Waarden, Lars, Johan Bruwer, and Jean-Philippe Galan. 2023. "Loyalty Programs, Loyalty Engagement and Customer Engagement with the Company Brand: Consumer-Centric Behavioral Psychology Insights from Three Industries." Journal of Retailing and Consumer Services 71 (March): 103212. https://doi.org/10.1016/j.jretconser.2022.103212.

Miller, Kent D. 2015. "Agent-Based Modeling and Organization Studies: A Critical Realist Perspective." Organization Studies 36 (2): 175–96. https://doi.org/10.1177/0170840614556921.

Nguyen, Linh Thi My, and Phong Thanh Nguyen. 2025. "The Interplay between Governance Mechanisms of Blockchain Platforms." Economics of Innovation and New Technology 34 (3): 367–91. https://doi.org/10.1080/10438599.2024.2346723.

Nishio, Kazuki, and Takahiro Hoshino. 2024. "Quantifying the Short- and Long-Term Effects of Promotional Incentives in a Loyalty Program: Evidence from Birthday Rewards in a Large Retail Company." Journal of Retailing and Consumer Services 81 (November): 103957. https://doi.org/10.1016/j.jretconser.2024.103957.

Nunes, Joseph C., and Xavier Drèze. 2006. "The Endowed Progress Effect: How Artificial Advancement Increases Effort." Journal of Consumer Research 32 (4): 504–12. https://doi.org/10.1086/500480.

Obiegbu, Chinedu James, and Gretchen Larsen. 2024. "Algorithmic Personalization and Brand Loyalty: An Experiential Perspective." Marketing Theory 25 (2): 199–219. https://doi.org/10.1177/14705931241230041.

Onggo, Bhakti Stephan, and Joël Foramitti. 2021. "Agent-Based Modeling and Simulation For Business and Management: A Review and Tutorial." 2021 Winter Simulation Conference (WSC), December 12, 1–15. https://doi.org/10.1109/WSC52266.2021.9715352.

Peffers, Ken, Tuure Tuunanen, Marcus A. Rothenberger, and Samir Chatterjee. 2007. "A Design Science Research Methodology for Information Systems Research." Journal of Management Information Systems 24 (3): 45–77. https://doi.org/10.2753/MIS0742-1222240302.

Rand, William, and Christian Stummer. 2021. "Agent-based Modeling of New Product Market Diffusion: An Overview of Strengths and Criticisms." Annals of Operations Research 305 (1–2): 425–47. https://doi.org/10.1007/s10479-021-03944-1.

Rand, William, and Roland T. Rust. 2011. “Agent-Based Modeling in Marketing: Guidelines for Rigor.” International Journal of Research in Marketing 28 (3): 181–93. https://doi.org/10.1016/j.ijresmar.2011.04.002.

Ribeiro-Rodrigues, Evelin, and Ana Paula Bortoleto. 2024. “A Systematic Review of Agent-Based Modeling and Simulation Applications for Analyzing pro-Environmental Behaviors.” Sustainable Production and Consumption 47 (June): 343–62. https://doi.org/10.1016/j.spc.2024.04.017.

Rizzati, Massimiliano, and Matteo Landoni. 2024. “A Systematic Review of Agent-Based Modelling in the Circular Economy: Insights towards a General Model.” Structural Change and Economic Dynamics 69 (June): 617–31. https://doi.org/10.1016/j.strueco.2024.03.013.

Rochet, Jean-Charles, and Jean Tirole. 2003. “Platform Competition in Two-Sided Markets.” Journal of the European Economic Association 1 (4): 990–1029. https://doi.org/10.1162/154247603322493212.

Secchi, D., V. Grimm, D.B. Herath, and F. Homberg. 2024. “Modeling and Theorizing with Agent-Based Sustainable Development.” Environmental Modelling & Software 171 (January): 105891. https://doi.org/10.1016/j.envsoft.2023.105891.

Secchi, Davide. 2015. “A Case for Agent-Based Models in Organizational Behavior and Team Research.” Team Performance Management 21 (1/2): 37–50. https://doi.org/10.1108/TPM-12-2014-0063.

Shah, Swayam, and Slinger Jansen. 2021. ‘A Framework for Designing Self-Sustaining Ecosystems with Blockchain’. In Software Business, edited by Eriks Klotins and Krzysztof Wnuk, 407:184–92. Lecture Notes in Business Information Processing. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-67292-8_14.

Sharp, Byron, and Anne Sharp. 1997. "Loyalty Programs and Their Impact on Repeat-Purchase Loyalty Patterns." International Journal of Research in Marketing 14 (5): 473–86. https://doi.org/10.1016/S0167-8116(97)00022-0.

Sonnenberg, Christian, and Jan vom Brocke. 2012. "Evaluations in the Science of the Artificial – Reconsidering the Build-Evaluate Pattern in Design Science Research." In Design Science Research in Information Systems. Advances in Theory and Practice, edited by Ken Peffers, Marcus Rothenberger, and Bill Kuechler. Springer. https://doi.org/10.1007/978-3-642-29863-9_28.

Steinbacher, Mitja, Matthias Raddant, Fariba Karimi, et al. 2021. "Advances in the Agent-Based Modeling of Economic and Social Behavior." SN Business & Economics 1 (7): 99. https://doi.org/10.1007/s43546-021-00103-3.

Stourm, Valeria, and Eric T. Bradlow. 2023. "Cross-Reward Effects in a Coalition Loyalty Program: The Impact of a Point Currency Devaluation." International Journal of Research in Marketing 40 (2): 276–93. https://doi.org/10.1016/j.ijresmar.2022.08.003.

Tanveer, Umair, Shamaila Ishaq, and Thinh Gia Hoang. 2025. "Tokenized Assets in a Decentralized Economy: Balancing Efficiency, Value, and Risks." International Journal of Production Economics 282 (April): 109554. https://doi.org/10.1016/j.ijpe.2025.109554.

Taylor, Gail Ayala, and Scott A. Neslin. 2005. "The Current and Future Sales Impact of a Retail Frequency Reward Program." Journal of Retailing 81 (4): 293–305. https://doi.org/10.1016/j.jretai.2004.11.004.

Tesfatsion, Leigh. 2006. Agent-Based Computational Economics: A Constructive Approach to Economic Theory. No. 12514. Staff General Research Papers Archive. Iowa State University, Department of Economics. https://ideas.repec.org/p/isu/genres/12514.html.

Troost, Christian, Robert Huber, Andrew R. Bell, et al. 2023. “How to Keep It Adequate: A Protocol for Ensuring Validity in Agent-Based Simulation.” Environmental Modelling & Software 159 (January): 105559. https://doi.org/10.1016/j.envsoft.2022.105559.

Utz, Manuel, Simon Johanning, Tamara Roth, Thomas Bruckner, and Jens Strüker. 2023. “From Ambivalence to Trust: Using Blockchain in Customer Loyalty Programs.” International Journal of Information Management 68 (February): 102496. https://doi.org/10.1016/j.ijinfomgt.2022.102496.

Van Der Merwe, Susara E., Reinette Biggs, Rika Preiser, Charmaine Cunningham, David J. Snowden, Karen O’Brien, Marcus Jenal, Marietjie Vosloo, Sonja Blignaut, and Zhen Goh. 2019. ‘Making Sense of Complexity: Using SenseMaker as a Research Tool’. Systems 7 (2): 25. https://doi.org/10.3390/systems7020025.

Waldherr, Annie, Martin Hilbert, and Sandra González-Bailón. 2021. “Worlds of Agents: Prospects of Agent-Based Modeling for Communication Research.” Communication Methods and Measures 15 (4): 243–54. https://doi.org/10.1080/19312458.2021.1986478.

Wang, Lu, Xin (Robert) Luo, and Frank Lee. 2019. “Unveiling the Interplay between Blockchain and Loyalty Program Participation: A Qualitative Approach Based on Bubichain.” International Journal of Information Management 49 (December): 397–410. https://doi.org/10.1016/j.ijinfomgt.2019.08.001.

Zott, Christoph, and Raphael Amit. 2008. ‘The Fit between Product Market Strategy and Business Model: Implications for Firm Performance’. Strategic Management Journal 29 (1): 1–26. https://doi.org/10.1002/smj.642.

Appendix

Table 5. ODD Protocol Documentation Summary

ODD Component	**This ODD component provides supporting information on:**
Purpose	The model simulates a token-based coalition loyalty platform inspired by the *Compro no Fundão* case to examine how consumer adoption, merchant participation, and platform governance co-evolve over time. It is designed to explore platform sustainability, behavioural feedback loops, strategic adaptation, and point-economy viability under heterogeneous agent behaviour.
Entities	Three interacting entities are represented: (i) Consumers, who shop, adopt, engage with, and potentially churn from the platform; (ii) Merchants, who decide whether to join, remain, or exit the platform and adapt campaign strategies; and (iii) a Platform Treasury, which manages point issuance, liabilities, and buyback operations.
State Variables	Consumers have adoption and engagement states, shopping habits, point balances, usage probabilities, and perceived/realised savings. Merchants track enrolment status, campaign archetypes, wallet balances, rolling performance metrics, and utility. The platform tracks outstanding points, financial flows, and system-wide KPIs.
Spatial and Temporal Scale	Agents are located in a continuous two-dimensional space with distance-based interactions. Time advances in daily steps. Simulations typically span 6–24 months, including a baseline period without the platform and a subsequent activation phase.
Process Overview and Scheduling	Each step includes consumer shopping decisions, transaction processing, point earning and redemption, updates to adoption and engagement states, merchant performance evaluation, campaign review and adaptation, and platform-level treasury operations. Processes are partially staggered to avoid synchronisation artefacts.
Design Concepts	The model incorporates heterogeneity, bounded rationality, adaptation, feedback, learning, social influence, inertia, and path dependence. Emergent outcomes such as adoption curves, merchant retention, and financial stability arise from decentralised interactions rather than centralised optimisation.

Initialization		Consumers and merchants are initialised with heterogeneous attributes drawn from empirically inspired distributions. Merchants are assigned adoption archetypes and campaign types. A fraction of compatible merchants may be enrolled at activation. Consumers start unregistered. Random seeds ensure reproducibility.
Input Data		The model is parameterised using stylised facts and empirical insights from loyalty-program literature and the *Compro no Fundão* pilot, rather than direct calibration to transaction-level data. Parameters govern behaviour, thresholds, and economic rules.
Submodels	**Consumer Behaviour**	Consumer submodels govern shopping frequency, merchant choice, platform awareness, registration decisions, usage probability, satisfaction, dormancy, and churn. Perceived and realised savings are tracked using exponential moving averages to drive behavioural feedback.
	Merchant Behaviour	Merchant submodels govern awareness, evaluation, entry and exit decisions, campaign archetype selection, wallet management, and periodic campaign review. Merchants adapt slowly based on utility, wallet health, and CLP penetration, capturing strategic inertia and experimentation.
	Point Economy and Treasury	The platform treasury manages point issuance, redemption liabilities, merchant top-ups, and optional point buybacks. Tiered earning rules and wallet constraints link transactional activity to financial sustainability.
Stochasticity		Randomness is used in agent initialization, shopping decisions, adoption thresholds, social exposure, and evaluation timing. Controlled random seeds allow systematic sensitivity analysis and replication.
Observation and Output		Model outputs include adoption states, engagement dynamics, merchant utility and retention, campaign evolution, transaction volumes, point flows, and platform-level sustainability indicators. Data are collected at agent and system levels for statistical analysis.

This table provides a concise summary of the model structure following the ODD protocol. The structure of this ODD documentation summary is adapted from the protocol introduced by Grimm et al. (2010), while the entries reflect a model-specific description of the key ODD components. The complete ODD specification, detailing all variables, processes, and functions, is provided as a separate appendix alongside the base model files to ensure transparency and reproducibility.

Table 6. TRACE Documentation Summary

TRACE Element	This TRACE element provides supporting information on:
Problem Formulation	**Group-level statement:** The model investigates the functioning and sustainability of a coalition loyalty platform (CLP) where consumers, merchants, and a central treasury interact through earning and redeeming points. It aims to understand adoption, engagement, merchant participation, and point-ecosystem dynamics in a spatial retail environment. The model addresses how behavioural rules, incentives, and economic flows jointly shape long-run platform performance and stakeholder outcomes **Individual-level contribution:** Focuses on merchant participation and strategy dynamics, examining how merchants decide to join, remain in, adapt campaigns, or exit a coalition loyalty platform under bounded rationality. The analysis specifically addresses how campaign review frequency and initial enrolment conditions shape participation stability, exit risk, and strategic heterogeneity. The goal is to generate insights that guide platform governance decisions rather than predictive or prescriptive outcomes.
Model Description	**Group-level statement:** Base ABM consists of three agent types; consumers, merchants, and a platform treasury, operating in a continuous 2D space. Consumers make category-based shopping decisions, form habits, adopt and engage with the CLP, and evolve based on realised and lost savings. Merchants evaluate whether to join or exit the program using evidence, utility, and inertia. The platform manages point issuance, redemption, and financial flows. The model uses scheduled daily steps, EMA-based evidence updates, and extensive state machines for adoption and engagement **Individual-level contribution:** Merchant behaviour is modelled through a discrete adoption and participation state machine, combined with periodic campaign reviews and performance-based exit rules. This structure allows merchants to adapt strategies over time while remaining subject to inertia and noisy performance signals. The design emphasises interpretability and traceability of merchant decisions, enabling clear links between model mechanisms and observed participation patterns in the experimental results.
Data Evaluation	**Group-level statement:** The model relies solely on internally generated simulation data, using stylised behavioural and economic assumptions drawn from loyalty-program literature. Parameter heterogeneity is introduced through distributions (Dirichlet for category preferences, Beta for habits, LogNormal for merchant characteristics). No external datasets are used; all insights derive from the ABM's output under varied platform and merchant configurations. **Individual-level contribution:** The experiment required a small modification to the base model and configuration to allow the initial enrolment to be adjustable before running

	the model. In terms of metrics, the only additional measures are calculated from columns already present in the data collector. (e.g., merchant_retention_rate is derived from merchants_joined and merchants_exited)
Conceptual Model Evaluation	**Group-level statement:** The conceptual foundations combine behavioural adoption theory, platform economics, and point-based incentive structures. Consumer and merchant state machines represent realistic decision pathways (awareness, evaluation, joining, churn, dissatisfaction). EMA-smoothed evidence, habit formation, spatial shopping behaviour, and treasury-level point economics collectively support a plausible representation of CLP dynamics. These mechanisms reflect well-documented patterns in loyalty ecosystems. **Individual-level contribution:** Simplifications include restricting merchants to a single active campaign and enforcing fixed campaign review intervals. These assumptions intentionally abstract from operational complexity to isolate strategic timing effects and participation dynamics. The chosen structure reflects a trade-off between realism and analytical clarity, prioritizing the identification of mechanism-driven patterns over detailed replication of real-world merchant decision processes.
Implementation Verification	**Group-level statement:** Validated core logic through stepwise debugging of agent creation, spatial placement, scheduling, point accrual/redemption, and treasury operations. Additional checks included verifying EMA updates, adoption-state transitions, merchant evidence accumulation, and DataCollector outputs. The model's structure and core processes were tested for internal consistency. **Individual-level contribution:** Implementation correctness was assessed through targeted single-run tests, parameter override checks, and inspection of logged merchant state transitions. These checks ensured that enrolment, campaign adaptation, and exit behaviours followed the intended logic. Consistency between the conceptual design and observed simulation behaviour was verified before running batch experiments used in the analysis.
Model Output Verification	**Group-level statement:** Outputs reflect expected behavioural patterns: meaningful incentives raise engagement, poor realised savings reduce usage, merchants join or exit based on local evidence and economic uplift, and point liabilities evolve according to issuance and redemption flows. Spatial clustering and habit effects also appear logically in consumer shopping choices. **Individual-level contribution:** Model outputs were evaluated against the qualitative behaviours motivating the design, including heterogeneous merchant lifecycles, staggered exits, and strategy clustering. No formal goodness-of-fit metrics are applied, as outputs are not calibrated to empirical data. Verification focuses on whether simulated behaviour remains coherent, stable, and interpretable across multiple seeds and parameter configurations.

Model Analysis	**Group-level statement:** The model supports scenario exploration across varied incentive structures, discount rules, merchant behaviours, and platform policies. Analyses include sensitivity tests, behavioural pattern inspection, and comparative experiments across simulation regimes. The ABM's design allows investigation of dynamic interactions and long-run system trajectories. **Individual-level contribution:** The analysis systematically varies campaign review frequency and initial enrolment fraction to examine sensitivity, variance, and path dependence in merchant outcomes. Multiple seeds are used to distinguish structural effects from stochastic noise. Comparative experiments are designed to isolate how review timing influences utility volatility, exit likelihood, and the emergence of conservative or aggressive campaign strategies.
Model Output Corroboration	**Group-level statement:** Observed outcomes align with theoretical expectations and stylised loyalty-program behaviour: S-shaped adoption curves, churn under low perceived value, merchant exit following insufficient uplift, treasury liability growing with generous rewards, and redemption-driven point circulation. These patterns reinforce the model's conceptual validity. **Individual-level contribution:** Independent empirical validation is not feasible in this context. Instead, corroboration relies on internal consistency checks, robustness across seeds, and alignment with stylised behaviours documented in coalition loyalty research. Additional confidence is obtained by verifying that distinct experimental conditions produce meaningfully different participation dynamics consistent with the model's theoretical mechanisms.

The structure of this TRACE summary table is adapted from the updated TRACE documentation framework outlined by Grimm et al. (2014). It organises the core elements of the TRACE approach into a consolidated overview, while the content reflects our model-specific interpretation and implementation of these eight components.

This table is organized in two layers: "Group-level statement" summarizes project-wide choices and shared foundations, while "Individual-level contribution" describes the work executed and reported in this individual document.

Table 7. One-way ANOVA results for merchant utility across campaign review intervals, by initial enrolment level

Condition	F-stat	p-value	Partial_eta_squared	Cohens_f	Significant
30%	3.9826	0.004270765	0.0990	0.3315	Yes **
60%	23.6054	4.79764E-15	0.3944	0.8070	Yes ***

Table 8. Tukey's HSD pairwise comparisons of merchant utility across campaign review intervals (30% initial enrolment)

group1	group2	meandiff	p-adj	lower	upper	reject
1	7	0.0762	0.0041	0.0175	0.1349	TRUE
1	15	0.0573	0.0593	-0.0014	0.116	FALSE
1	30	0.0517	0.1128	-0.007	0.1104	FALSE
1	60	0.0236	0.7994	-0.035	0.0823	FALSE
7	15	-0.0189	0.9008	-0.0776	0.0398	FALSE
7	30	-0.0245	0.7776	-0.0832	0.0342	FALSE
7	60	-0.0525	0.1028	-0.1112	0.0062	FALSE
15	30	-0.0056	0.9989	-0.0643	0.0531	FALSE
15	60	-0.0336	0.5102	-0.0923	0.025	FALSE
30	60	-0.028	0.6801	-0.0867	0.0307	FALSE

Table 9. Tukey's HSD pairwise comparisons of merchant utility across campaign review intervals (60% initial enrolment)

group1	group2	meandiff	p-adj	lower	upper	reject
1	7	0.0946	0.0000	0.0544	0.1348	TRUE
1	15	0.1229	0.0000	0.0827	0.1631	TRUE
1	30	0.1203	0.0000	0.0801	0.1605	TRUE
1	60	0.0819	0.0000	0.0417	0.1221	TRUE
7	15	0.0284	0.2965	-0.0118	0.0686	FALSE
7	30	0.0257	0.3977	-0.0145	0.0659	FALSE
7	60	-0.0127	0.9077	-0.0529	0.0276	FALSE
15	30	-0.0027	0.9997	-0.0429	0.0375	FALSE
15	60	-0.041	0.0431	-0.0812	-0.0008	TRUE
30	60	-0.0383	0.0694	-0.0786	0.0019	FALSE

Figure 7. Merchant utility levels and associated confidence interval widths by campaign review interval under alternative initial enrolment levels

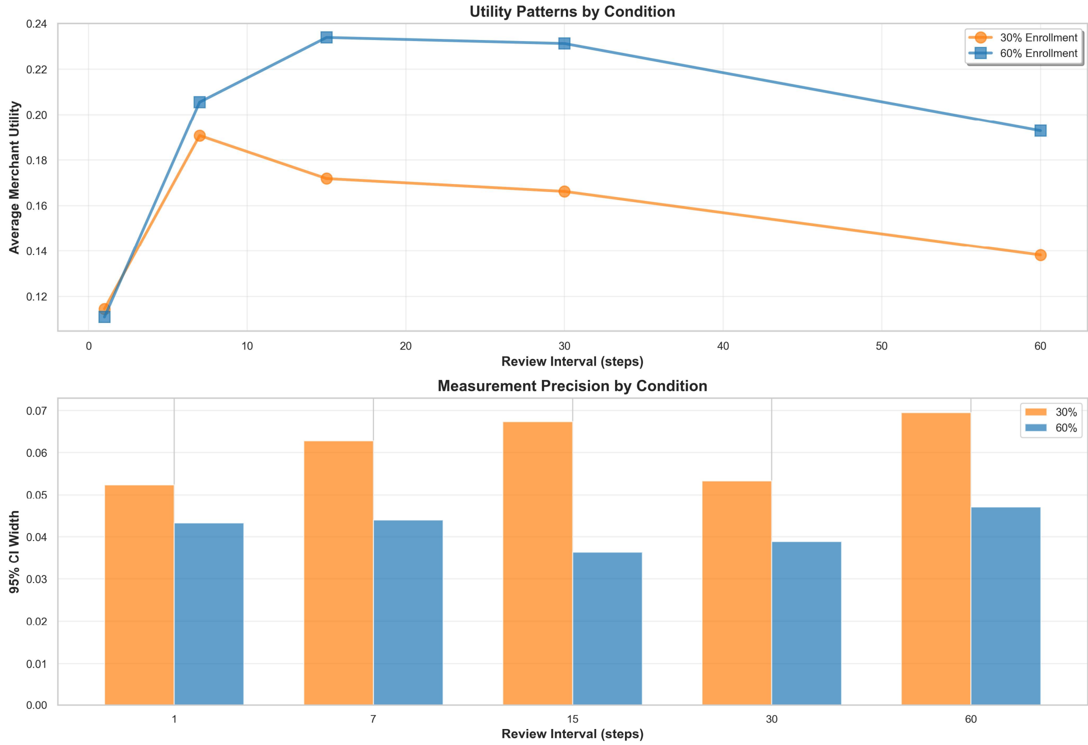

Figure 8. Merchant participation and performance outcomes by campaign review interval under 30% initial enrolment

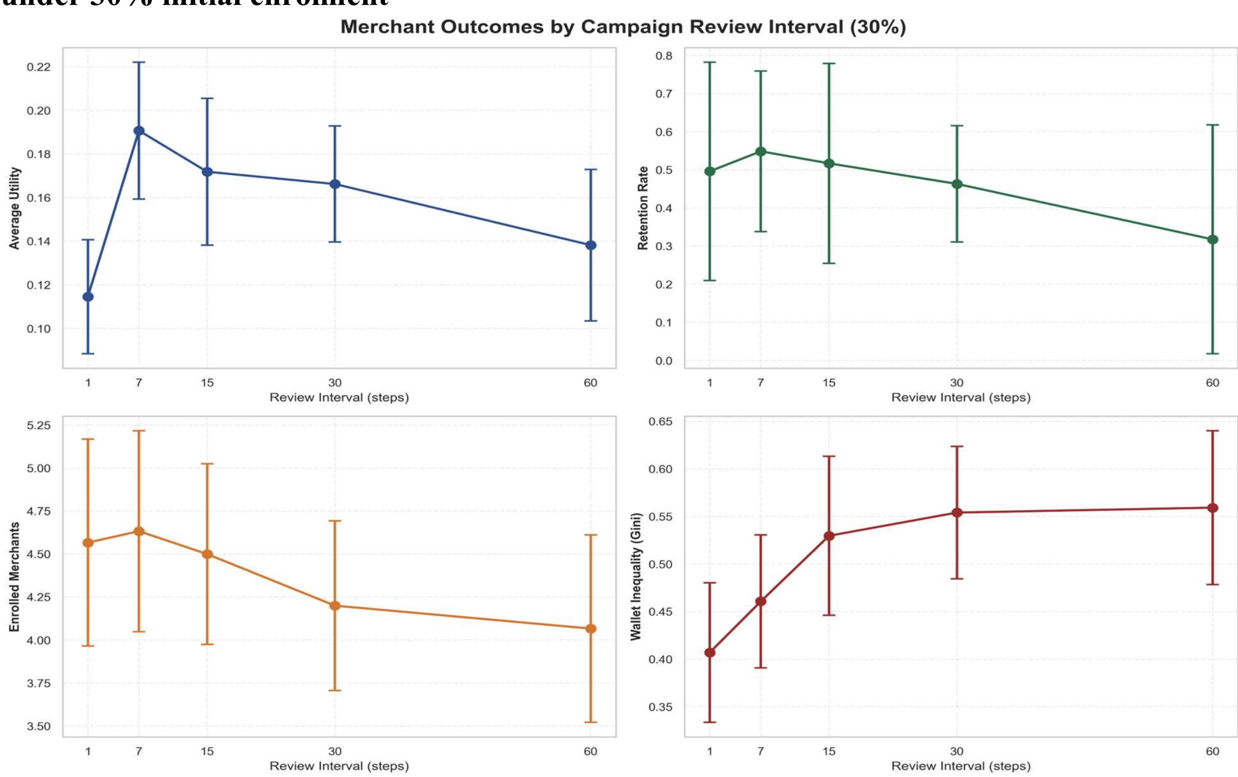

Figure 9. Merchant participation and performance outcomes by campaign review interval under 60% initial enrolment

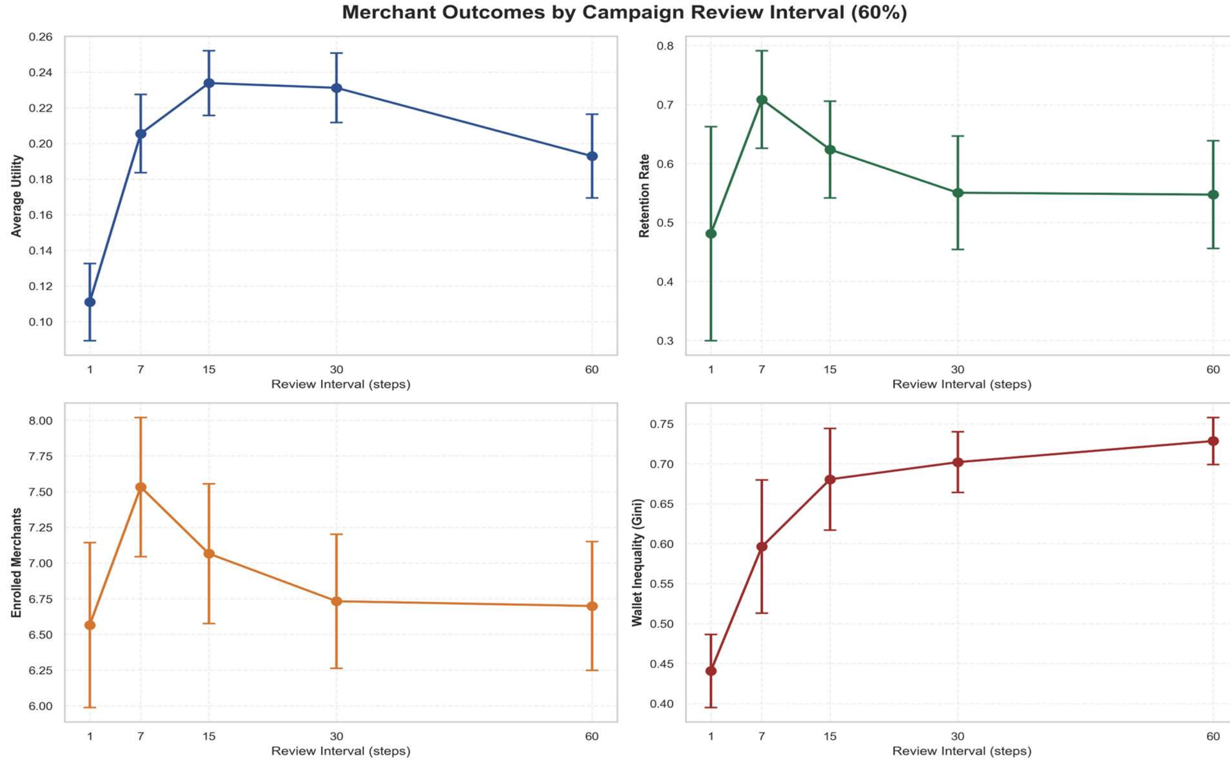

Figure 10. Consumer redemption rate trajectories by campaign review interval under alternative initial enrolment levels

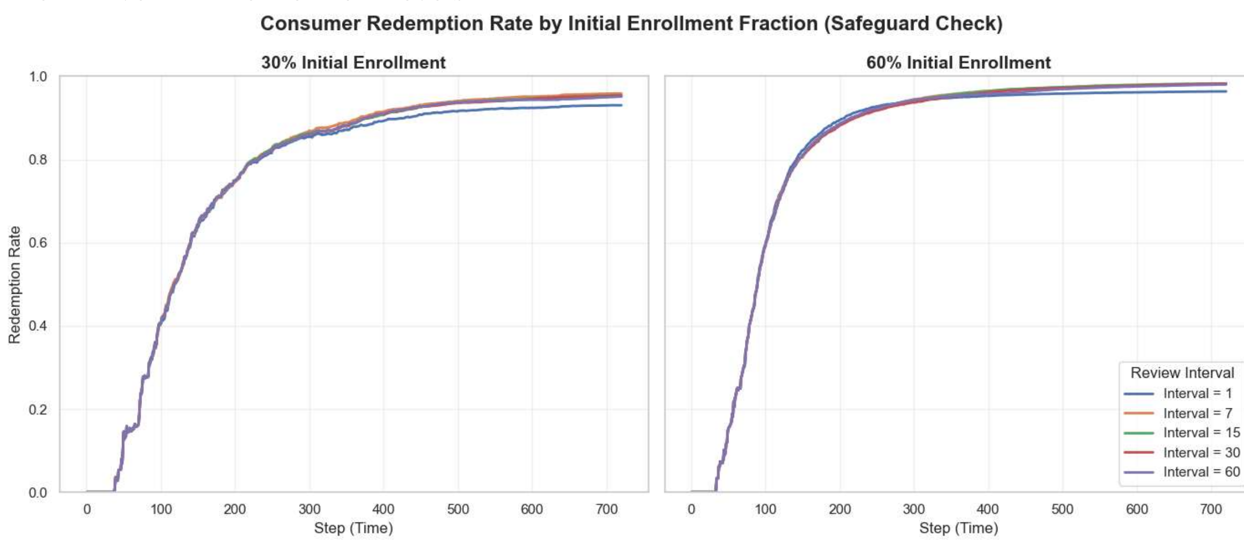

Figure 11. Consumer registration dynamics by campaign review interval under alternative initial enrolment levels

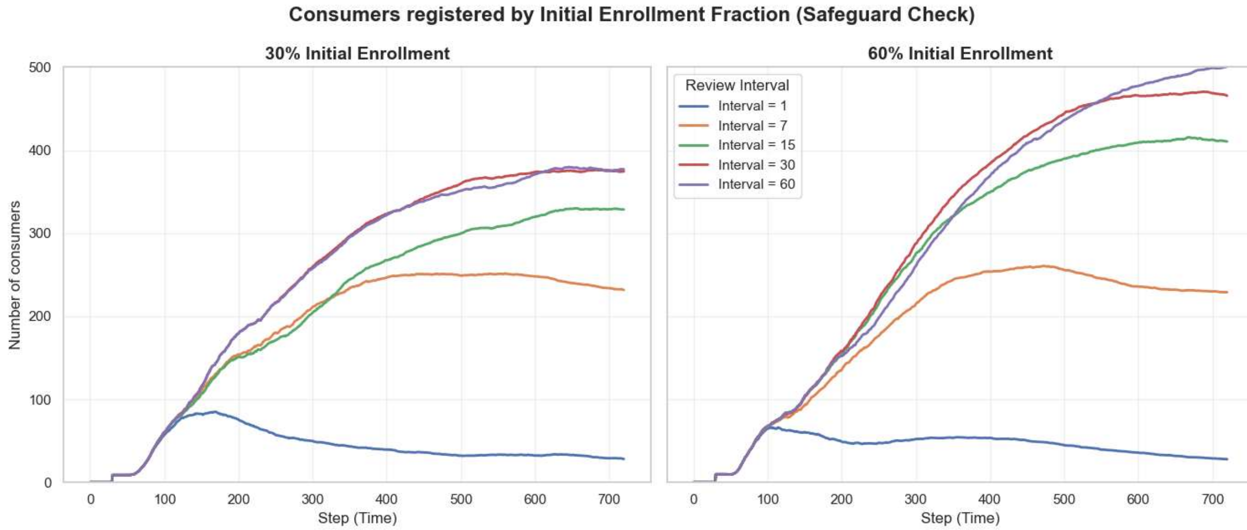

Figure 12. Consumer churn accumulation by campaign review interval under alternative initial enrolment levels