Marketing Intelligence - Executive Summary

Date: October 2025

This document uses anonymized, synthetic examples. No confidential data is disclosed. Figures are illustrative and meant for public demonstration only.

Objective

Summarize how customer and product level data were used to build segments, design-controlled email campaigns, and measure incremental revenue and repeat purchase behavior with statistical rigor.

Key Outcome KPIs

KPI	Observed	Interpretation
Email open rate	29.4%	Targeted clusters reached open rates around thirty percent.
Conversion Rate	9.1%	Roughly one in eleven recipients generated at least one order.
Revenue per dollar invested	34.2	Each unit of campaign budget returned more than thirty units of gross sales.
Repeat purchase uplift	12%	Treatment cohorts showed about twelve points higher repeat purchase rate.

Notes:

1) KPIs combine several campaigns such as win back, loyalty, free shipping, seasonal events, and cart abandonment.

Key Messages

- Segmentation combined demographic data (region, city, store radius) with behavioral data such as purchase recency, frequency, monetary value, and preferred product categories.
- Customer clusters were built with standard RFM scores and k means style clustering on category and format shares, producing segments that were both statistically distinct and easy to interpret.
- Product clustering and market basket analysis grouped more than 100k SKUs into roles and natural bundles, simplifying offer design and cross sell.
- Campaigns were run with explicit treatment and control groups, which allowed uplift and ROI to be measured rather than inferred.

- Governance routines embedded the segments in the CRM and ensured that results by segment were reviewed regularly, so learning from tests translated into the next wave of campaigns.

Analytical approach

The starting point was a collectibles retailer with a large omnichannel base and a very long tail of products. Transactional data, email logs, and store event records were consolidated at customer and SKU level.
Feature engineering followed a few standard patterns:

- RFM profiles per customer over a rolling horizon.
- Category and format preferences from purchase and event history.
- Email engagement metrics such as opens, clicks, and historical unsubscribes.
- SKU level attributes including price band, margin, demand volatility, and occurrence in baskets.

On this base, segmentation was layered in three directions:

1. **Value and lifecycle**
   Customers were grouped by RFM into high value loyal, high value at risk, new high value, and long tail segments.
2. **Interest and format**
   Category and format shares fed clustering algorithms that separated competitive players, focused collectors, and budget oriented buyers.
3. **Channel and engagement**
   Email engagement and simple propensity scores highlighted who could receive fewer contacts and who should be prioritized for testing.]

Product clustering used two lenses:

- Attribute based clusters on standardized economic features, to define roles such as traffic drivers and margin drivers.
- Basket based analysis to detect bundles and frequent combinations for cross sell and “complete your set” offers.

Crossing customer clusters with product clusters produced a simple matrix that guided which offers made sense for each audience.

Experiment design and measurement

Major campaign types were bound to specific segments: for example, high value at risk customers in competitive formats for win back offers, or new high value customers for early loyalty programs. For each campaign, eligible customers were randomly split into treatment and control groups.

Measurement applied standard tools:

- Open, click, and conversion rates computed with confidence intervals for both groups.
- Chi square tests for difference in proportions to check whether observed uplifts were statistically significant.
- Incremental revenue calculated as the difference in average revenue per customer between treatment and control, multiplied by treated customers.
- Campaign ROI derived from incremental revenue divided by incentive and distribution cost.

In selected waves, uplift style models were tested to rank customers by expected incremental effect of treatment, concentrating budget where it was most effective.

Governance and impact

Segment tags and cluster ids were stored in the marketing database and synchronized with the email platform so that campaigns could be built directly on top of them. Weekly routines monitored KPIs by segment, and monthly reviews merged or redefined clusters that no longer behaved differently in practice.

The result was a marketing engine that relied less on one off lists and more on a stable segmentation structure, with measurable gains in open rate, conversion, budget productivity, and repeat purchase, as reflected in the consolidated KPIs above.