

RetailPulse

Product Requirements Document · Pricing & Promotion Analytics · MVP

Field	Detail
Product	RetailPulse: price & promo simulator, promo scorecard, demand & stock planning
Company	FreshCart (fictional grocery chain: 40 stores, 4 regions, 120 products; synthetic data)
Author / PM	Jaykumar Raichura
Status	MVP prototype complete
Stakeholders	Category Management, Pricing, Supply Chain, Store Ops, FP&A, Merchandising leadership, Data Engineering

1. Summary

FreshCart runs close to a thousand promotions over two years, planned largely on gut feel. Promotions reliably lift volume, but nobody measures whether they make money once lost margin on the discount, sales stolen from sibling products, and pantry-loading are counted. Stores also run out of stock during deep promotions because every promo is planned with the same flat inventory rule. RetailPulse measures each product's price sensitivity from sales history and turns it into three tools: a simulator for pricing and promo decisions, a scorecard of past promotions, and a demand-and-stock check for upcoming promos.

2. Problem

- Promo profitability is invisible: of 971 promotions in FY2024–FY2025, only 33% paid back once cannibalization and pull-forward were counted; the program lost \$454K net on \$3.21M of FreshCart-funded discounts.
- Stockouts: 14% of promo weeks stocked out versus 0.1% of regular weeks; deep snack promos stocked out 88% of the time.
- Price decisions lack evidence: FY2025 price increases lifted gross margin 4.7% while units fell 5.8%. Category managers can't tell which products can absorb increases and which can't.

Problem statement: How might we help category managers set prices and promotions that grow profit, not just volume, and stock them correctly?

3. Goals and non-goals

Goals

- Show the net incremental margin of every past promotion, after cannibalization, pull-forward and vendor funding.
- Let category managers test a price or promo decision before committing, with a clear recommendation.
- Cut promo-week stockouts by forecasting promo demand instead of using a flat rule.
- Turn net promo margin from –\$454K to break-even or better within two planning cycles.

Non-goals (MVP)

- Automated price changes (humans approve every decision).
- Competitor price tracking.
- Store-level or personalized pricing.
- Long-term price-image effects (flagged as a caveat in the simulator).

4. Users

Persona	Job to be done	Pain today
Category Manager	Decide which promos to run, repeat or stop	Judged on volume; no view of true profit
Pricing Analyst	Set regular prices by product	Elasticity estimates live in one-off spreadsheets
Supply Chain Planner	Order the right stock for promo weeks	Flat 2x rule over- and under-stocks
FP&A (supporting)	Explain margin changes and promo ROI	Can't split price, volume and mix effects

5. Discovery insights

- Volume is the wrong scorecard: category managers are rewarded on unit lift, which favors deep discounts.
- Shoppers swap, not add: much of a promoted item's lift comes from its FreshCart or National-brand sibling.
- Planners know the 2x rule is wrong but have nothing better; they need a number per promo, not a model.
- Trust depends on transparency: users want to see the elasticity and its uncertainty, not a black box.

6. Solution approach

Approach	Accuracy	Effort	Decision
Same discount rules for every product	Low	Low	Rejected: ignores very different price sensitivity
SQL baseline (prior 4 weeks average)	Medium	Low	Kept for quick checks; ignores seasonality
Per-product regression: price, promo type, seasonality, sibling promos, post-promo dip	High ($r = 0.97$ vs. truth)	Medium	Selected (MVP)
Machine-learning demand model (gradient boosting)	Potentially higher	High; less explainable	Revisit at V2

7. User stories

ID	User story	Priority
US-1	As a category manager, I want to simulate a discount and see net margin after cannibalization and pull-forward.	Must
US-2	As a category manager, I want every past promo scored Repeat / Fix / Stop.	Must
US-3	As a pricing analyst, I want to see each product's elasticity with its uncertainty.	Must
US-4	As a planner, I want forecast demand and a recommended order for each upcoming promo.	Must
US-5	As a category manager, I want a plain-language recommendation for my scenario.	Should
US-6	As FP&A, I want a margin bridge splitting volume, mix and rate.	Should
US-7	As a pricing analyst, I want competitor prices in the model.	Won't (V2)

8. Functional requirements

ID	Requirement	Acceptance criteria
FR-1	Simulator: product, discount, promo type, vendor funding, inventory plan; regular-price mode	Matches Excel model to the dollar
FR-2	Waterfall: gross gain, cannibalization, pull-forward, net	Components sum to net
FR-3	Promo scorecard with filters and sorting	All 971 promos; verdict rules documented
FR-4	Demand forecast with holdout accuracy shown	Model beats naive on holdout
FR-5	Upcoming promo stock check with risk flag	Flag when forecast > planned stock
FR-6	Auto-generated recommendation	Names verdict, best depth and stock risk
FR-7	Price-image caveat on large increases	Shown when optimum is at range edge

9. Data requirements

- Star schema: fct_sales (product × region × week) with dim_product, dim_region, dim_promo, dim_date.
- Quality gate before each load: unique grain, sold ≤ inventory, price paid ≤ list, no orphan promo IDs, revenue = units × price.
- Stockout weeks excluded from model fitting, because sales understate true demand.
- Model refresh: monthly; elasticities versioned so decisions can be traced to the model used.

10. Success metrics

Metric	Baseline	Target
Net promo margin (2 cycles)	-\$454K	≥ \$0
Share of promos that pay back	33%	50%
Promo-week stockout rate	14%	≤ 5%
Promo-week forecast error (MAPE)	44% (naive)	≤ 15% (prototype: 14%)
Weekly active category managers	0	80%
Decisions using the simulator	0	Every promo in the plan

11. Experiment policy

Programs that change customer behavior are tested before rollout. The digital-coupon pilot (20 vs. 20 stores, 8 weeks) raised sales 3.8% (95% CI +2.2% to +5.3%) but gross margin dollars by only 0.8%, which is not statistically significant. Decision: do not roll out broadly; redesign coupons toward higher-margin and private-label items and re-test. Success is judged on margin, not sales.

12. Release plan

Phase	Scope	Exit criteria
MVP	Simulator, scorecard, demand & stock check	Category managers use it for one planning cycle
V1	Vendor-funding negotiation view, automatic promo plan review, bulk scenarios	Net promo margin ≥ \$0
V2	Competitor prices, ML demand model, "Ask RetailPulse" natural-language Q&A (AskData pattern)	Forecast MAPE improvement ≥ 10%

13. Risks and mitigations

Risk	Mitigation
Elasticities wrong for low-volume products	Show uncertainty; fall back to category elasticity when CI is wide
Short-term model over-recommends price increases	Caveat in simulator; cap at +5–8%; A/B test first
Category managers reject a lower-volume scorecard	Pair net margin with volume; co-design verdict rules
Stockouts bias the model	Exclude censored weeks; track lost-sales estimate

14. Open questions

- Should vendor funding targets be set per supplier using the scorecard?
- What minimum net margin should a promo need to enter the plan?