

Price Elasticity

Should a supermarket raise prices? Estimating category price elasticity to guide the decision.



Project context

A public retail file: 77 stores, 4 years, 169,677 store-product weeks across cold cereal, frozen pizza, bag snacks and an oral-hygiene control category.

The goal of the study: which prices can we rise, by how much, and what impact will it bring to our Gross Margin.

The file carries no unit costs, so a cost model is assumed and labelled everywhere money appears. On client data, real costs replace it and nothing else changes.

My role

Data Science × FP&A

- **Statistical modelling:** price response measured per product and store, with intervals, controlling for displays and features
- **Validation design:** a blind test against three independent analysts, written down before the engine ran
- **FP&A translation:** each move pre-cleared against the volume its own margin can afford to lose
- **Rollout design:** capped moves, staged in waves, with a matched-store test and an automatic tripwire

Raise nine cereal SKUs, in steps, starting with six

Up to \$238k / yr

of additional gross margin across the nine cereal SKUs that carry usable evidence, at assumed costs, if every move is walked all the way to its capped optimum. The first step is far smaller, and is the only thing being asked for now.

Raise 9, hold 2, refuse 4

Nine carry evidence for a rise. Two are already at or beside their optimum, so they hold. Four read implausibly and get no recommendation at all.

Start with six, at +5%

In 38 matched stores for eight weeks, with an automatic revert and a named owner. Nothing goes chain-wide before it reads clean.

What we will not do

No price is set beyond what this file has seen. No frozen pizza rise: it fails the margin line. And the four SKUs that read implausibly get no optimum, no recommendation and no money in the total.


Business Overview

Regional Supermarket Chain: 77 stores, data going from January 2009 to January 2012.


Four categories: Cold cereal carries 53.8% of the business; frozen pizza 23.1%, bag snacks 16.9%, oral hygiene 6.2%.

One deliberate price action in three years: cold cereal repriced in steps from \$2.86 to \$3.03 (12 of 15 SKUs, February to August 2011) and held since. Cereal volume ran -9.4% per trading week year on year.


The 15 cereal SKUs at the center of this audit are the file's largest category and its only recent pricing event.



Growth, then a plateau, on an unchanged mix



net revenue per trading week, by month · whole book, all four categories




share of net revenue by category · the mix barely moves across three years

Two flat years, then one repricing that held

Cold cereal was repriced in steps from February to August 2011, \$2.86 to \$3.03, and the level has held since. It is the only deliberate list action in three years.

The solid line is the shelf price, the dashed line what shoppers actually paid. The distance between them is promotion.



cold cereal, fixed basket, dollars per unit · shaded band is the repricing window

The one time this chain raised cereal prices

+4.3%

on the fixed-basket list, February to August 2011, and held since

-9.4%

cereal units per trading week, full year against full year

-2.19

the two divided: what that one event implies on its own

The model lands in the same place

Measured properly across 169,677 store-product-weeks, the cereal list response is **-2.20**. The chain's own repricing, read with nothing but a division, says **-2.19**. The engine agrees with the one real experiment on these shelves.

Why we do not stop there

One event cannot separate price from everything else moving that year. Cereal was in a long decline across the market, and a single synchronised move **cannot tell the two apart**. That is what the store-product-week panel and the week effects are for.

Nobody could say what a raise would cost

The list had not moved in years, and no measurement existed to say what moving it would do

01

No price response measured

What volume does when price moves was unknown, per product or at all, every debate ran on opinion.

02

Promotion effects credited to price

Display and leaflet weeks made demand look far more price-sensitive than it is, freezing the list.

03

Standing still is a decision too

Holding a list still has a price of its own, and nothing in the file had ever measured it.

The cost of not knowing

Measured product by product, the category was leaving \$238k a year of gross margin on the table at assumed costs, nine products priced below what their own demand would bear.

Two more sat at or beside their optimum and were held. Four returned a number too weak to price on and were refused outright: a study that refuses is one you can believe when it clears.

\$238k / yr

additional gross margin at assumed costs,
across the nine SKUs with usable
evidence

9 of 15

products cleared to rise; two held, four
refused


169,677

store-product weeks behind the
estimates

Why the simple read of your data misleads

Plot monthly price against monthly volume and a line through the cloud reads -2.02 . It looks like an answer, and it is not one.

Every point mixes a price change with a display week, a holiday, and a different store mix. The model separates them: each product in each store compared against itself, with the week absorbing whatever hit the whole panel.




each point is one month of the cereal basket · realised price against units per trading week

Where the price variation comes from

The hardest question this study gets: the list moved once, in 2011, so is the elasticity really the category's decline wearing a price label?

It is not, because the estimate never used that event. Every product is held against itself in its own store, and every week is held against the whole chain. A national decline is common to the week, so it is removed before the coefficient is fitted.

What is left is the same product, in the same store, priced away from its own norm and from what the rest of the estate did that week. Depending on the product that is 2.6% to 7.6% of price movement, and it is the only thing the number is fitted to.



pale: all the price movement in the file. solid: what survives the store-product and week effects, which is the only part the estimate is fitted to.

cold cereal · 15 SKUs · 77 shops · 156 weeks

Measured properly, then tested blind

Price response estimated per product and per store, with 80% intervals, never one category average doing the work of fifteen products.

Display and leaflet weeks are controlled for separately, so merchandising lift is not billed to price.

Estimates are fitted on early years and tested on later ones; a control category is expected to come back weakest, written down in advance.


The blind test

“Before the engine ran, three analysts estimated each category from the raw file and the results were written down. All three matched, two to the second decimal, the control category weakest, as predicted.”

What price actually does to volume here

Every category sits well below -1 : a list-price rise loses volume, and revenue falls with it. Whether margin still improves depends on the cost, and at the assumed 40% that line sits at -2.5 .


Cereal at -2.20 clears that line. Frozen pizza at -3.21 does not, so no pizza price is raised here. The control category is the least elastic of the four, exactly where a check designed to fail should land.



A “promotion effect” is mostly not the discount

A display week adds up to +73% volume with no price cut at all.

Naive models bill this exposure to price and overstate elasticity, which is why lists freeze.



extra volume in a display week, at an unchanged price · vs the same store-product's quiet weeks

the exposure half of a promotion, measured alone

Premium competes with premium, not with value

Raise DiGiorno Pepperoni 10% and Freschetta Pepperoni lifts about 8.4%. Buyers move to the other premium rising-crust brand, not to the store's own: +0.07, +0.05 and +0.03 on three own-brand products, with intervals tight enough to be sure.

A shopper buying a \$6 pizza is not one price rise away from a \$3 one. The room to take price on premium is not bounded by what the value line costs.

Two rivals could not be measured, and the file says so.

RIVAL TO DIGIORNO PEPP.	CROSS	READING
Freschetta Pepperoni	+0.84	strong substitute
Freschetta Brick Oven	+0.70	strong substitute
Freschetta 4 Cheese	+0.62	strong substitute
Own-brand Rising Crust 3 Meat	+0.07	no link
Own-brand Rising Crust Pepp.	+0.05	no link
Own-brand Rising Crust Supr.	+0.03	no link

fitted per pair on 2,700 to 11,800 store-weeks · noise floor ± 0.22 , anything inside it reads as no link

Three specialists measured blind, and the engine matched them

Before the engine ran, independent analysts estimated each category from the raw file, and the expected results were written down. Then the engine ran.

CATEGORY	BLIND ESTIMATE	ENGINE RESULT	VERDICT
Cold cereal	-2.20 / promo -2.89	-2.20 / -2.89	match to the 2nd decimal
Frozen pizza	-3.25 / -3.88	-3.21 / -3.88	match
Bag snacks	-1.82 / -1.43	-1.79 / -1.44	match
Oral hygiene (control)	-1.28 / -1.61	-1.28 / -1.61	exact, and weakest as predicted

Verified by hand. One SKU was also recomputed from the raw file with nothing but a spreadsheet: -0.55 against the engine's -0.54.

What the answer does when the years change

The file runs 37 months and the engine refuses any scope under 24, so cereal was fitted twice: once on the first 24 months, once on the last 24. Same method, different years. The list response moved 0.24, from -1.84 to -1.59.



cold cereal · list price response, merchandising held constant · 80% intervals

The optimum on paper

\$3.73 → \$4.85

moves 30%, which is why it is never the recommendation

The price actually recommended

+20.2% → +24.0%

both capped 10% past the dearest price in the window

The first +5% step

+5.4% → +7.9% gross margin

pays under either number, and keeps paying to about -2.9

The two windows share 11 months, so 0.24 is a floor on the drift and not a ceiling. The promotional half moved the same way, -1.91 to -1.55, and so does the uncontrolled -2.20 the price list is built on. The intervals still overlap: this is drift inside the noise, and the cap is what keeps the recommendation standing either way.

Eight explanations, tested and eliminated

Promotion on the products themselves

feature and display flags as controls → no change, already in

Aisle traffic from big promotions

category volume as a control → made it worse, not better

Aisle promotion intensity

share on feature, display and TPR → no effect at all

A private-label product behaving oddly

re-run with branded focal products → negatives persisted

One contaminated week

refitted 156 times, one week out each → no week flips any sign

Too few units to measure

volume compared across categories → pizza sells less, works fine

Not enough price movement

price variation compared → bag snacks has the most

Prices reacting to local demand


two-stage least squares on 75 shops → leaves only chain-wide moves

↔ The safe first step

Every move pre-cleared before anyone moves

On strong brands, +5% price costs about 2.6% of volume, while the margin gained can afford to lose 14%.

The conversation changes from “will volume fall?” to “will it fall by more than we can afford?” It is exact arithmetic, not a model.



Why the price list is cereal only

All four categories were estimated to the same standard. Three of them earned a different answer.

COLD CEREAL

Priced

53.8% of sales, and the only deliberate list move in three years. The money and the cleanest evidence are both here.

BAG SNACKS

Next in line

-1.79, interval -2.08 to -1.49, so it clears the same margin line.
16.9% of sales, not priced in this study: the obvious second wave.

FROZEN PIZZA

Refused

-3.21, with the whole interval left of the margin line. A rise there loses margin, so no pizza price moves.

ORAL HYGIENE

Held out

The control, kept out by design so the method has something to fail against. Least elastic of the four, and still not a recommendation.

Three levels of evidence, not one column of prices

Every product lands in one of three, and the level decides what may be done with it. Fifteen products each carrying a recommended price would read as more work and less evidence.

LEVEL 1 · STRONG EVIDENCE

8 products · \$200,000

WHAT QUALIFIES

A precise estimate, sitting clear of the margin line, on enough price movement to trust it.

WHAT HAPPENS

Act on it. Raise in steps, or hold where the price is already at its optimum.

LEVEL 2 · TEST ONLY

3 products · \$37,800

WHAT QUALIFIES

The direction is supported but the size is not: a margin too thin to carry the loss, or an estimate inherited from the category rather than measured on the product.

WHAT HAPPENS

An optimum is quoted, but the first move stays at 5% until the number has proved itself.

LEVEL 3 · NO RECOMMENDATION

4 products · \$0 claimed

WHAT QUALIFIES

The estimate reads between -0.41 and -0.59 , which no cereal sustains. The number is not usable, and a precise wrong number is worse than a missing one.

WHAT HAPPENS

No price is recommended. The 5% probe exists to create the missing evidence, not to collect the money.

The headline counts levels 1 and 2 only, which is \$237,830. At a 5% move the level 3 products would be worth \$64,500, and that figure is deliberately outside the total: **a number with no recommendation behind it does not belong in one.**

All fifteen SKUs, ranked by the money at stake

the top four: \$158,600, all capped

SKU	PRICE	ELASTICITY	OPTIMAL	DECISION	AT STAKE / YR
GM CHEERIOS	\$3.89	-1.44	\$5.13 (+32%, capped)	raise in steps	\$51,800
KELL BITE SIZE MINI WHEAT	\$3.07	-1.39	\$4.07 (+33%, capped)	raise in steps	\$41,600
KELL FROSTED FLAKES	\$2.62	-1.77	\$3.52 (+34%, capped)	raise in steps	\$37,100
POST FRUITY PEBBLES	\$2.30	-1.46	\$3.21 (+40%, capped)	raise in steps	\$28,100
GM HONEY NUT CHEERIOS	\$2.55	-0.54	+5% probe only	no recommendation	not counted
PL BT SZ FRSTD SHRD WHT	\$2.12	-1.53	\$2.61 (+23%, capped)	raise in steps	\$22,500
GM CHEERIOS (format 2)	\$2.68	-0.47	+5% probe only	no recommendation	not counted
POST FM SZ HNYBNCH OT ALM	\$3.14	-1.42	\$3.78 (+21%, capped)	raise in steps	\$18,900
KELL FROOT LOOPS	\$2.51	-2.35	\$3.29 (+31%)	start +5% (thin margin)	\$17,700
POST HNY BN OTS HNY RSTD	\$3.14	-0.59	+5% probe only	no recommendation	not counted
QKER LIFE ORIGINAL	\$2.37	-2.17	\$3.47 (+47%)	start +5% (thin margin)	\$11,200
PL RAISIN BRAN	\$1.75	-1.71 inherited	\$2.09 (+19%)	weakest claim, last	\$8,900
PL HONEY NUT TOASTD OATS	\$1.76	-0.41	+5% probe only	no recommendation	not counted
QKER CAP N CRUNCH	\$2.22	-3.06	\$2.33 (+5%)	hold: move worth \$30/yr	\$30
QKER CAP N CRUNCH BERRIES	\$2.30	-3.13	\$2.30 (0%)	hold: at optimum	\$0

the sentinel, frozen at \$2.30

Four SKUs read -0.41 to -0.59. No optimum, no recommendation, no money counted:
the 5% probe only creates the evidence.

\$238k is a ceiling, not a forecast

The ceiling

\$238k / yr

The nine SKUs with usable evidence, at their capped optima, in all 77 stores, at assumed costs. The four that carry no recommendation are not in it. The size of the prize, not a plan.

The first rung

6 SKUs at +5%

In 38 of 77 stores, for eight weeks. This is what is being approved, and it is a small fraction of the ceiling.

What sits between

Several rungs

Each one re-measures before the next, so reaching a +32% optimum takes multiple waves and reads. Year one captures part of the ceiling, not the ceiling.

Quote all three, never one. A ceiling read as a run-rate is how a pricing study loses a board's trust in its second year.

Nobody jumps to +32%: every rung re-measures

01 · WEEKS 1–8

Six SKUs at +5%

In the 38 test stores. Estimated losses –2.0% to –10.8% against thresholds of –11.7% to –19.2%: a safety gap of 1.6× to 6×.

02 · WEEKS 9–16

Seven more at +10%

After step 1 reads clean: the capped brands. The inherited-estimate SKU takes the smallest step, never in the same window as step 1.

03 · WEEK 17 ON

Walk the optimum

Further +10% rungs, thresholds recomputed at each new base. The optimum is the destination; the ladder is the route.

38 matched pairs, one tripwire, one sentinel

Matched pairs

Stores of similar size, adjacent in revenue: S26983 (\$452.8k) against S25229 (\$452.5k), a 0.1% gap. Test stores read against controls.

Everything else frozen

Same promotions, same displays, same shelf. A display week moves volume up to +73% on its own; one slipping in would corrupt the read.

Automatic tripwire

A SKU breaching its affordable-loss threshold two weeks running reverts to its old price. That SKU only, and a named owner signs every revert.

A sentinel SKU

Cap'n Crunch Berries stays at \$2.30 everywhere. If it moves, something other than our prices is moving the market.

What this cannot tell you

Costs are assumed

This public file invents unit costs. Optimal prices and dollar figures re-price the day real COGS arrive; elasticities do not move.

Store switching is invisible

The data sees this chain only. If every cereal rises 30%, some customers change stores; the staged caps exist partly for that.

Price image is strategy

Cereal is a known-value aisle. Whether a flagship carries a +30% optimum is a commercial decision this analysis informs, not makes.

Beyond the range

Optima above observed prices are capped at 10% past the highest price ever carried; the store test does the rest.

How to keep this honest

Re-estimate quarterly

Elasticity is not a constant; it moves with competition and category health.

Track realised against list monthly

Discount leakage is where the money went, and it is invisible in revenue reporting.

Re-run after structural change

A delisting, a new competitor, a supply interruption: any of them changes the numbers.

Key lessons learned



Checkable beats clever

A method your analyst can reproduce moves a management meeting further than a sophisticated black box.



Refusals build trust

Two competitor pairs were refused by name; the held product and the weak control made the cleared moves believable.



Exposure is not a discount

Separating what the display does from what the price cut does is the difference between a frozen list and a defensible raise.

One export, thirty minutes of your analyst's time

Sales lines

Product × store × week (or day): units, gross and net revenue.
Two or more years. CSV or Excel, any column names.

Unit costs

Per product, ideally per period. This is what turns directions into exact prices.

Promo flags

Display, feature and TPR markers per line let us split price from exposure.

Nothing else

No system access, no IT project. Files go through a private workspace with a signed DPA.

Ten working days from file to decision

DAYS 1–2

Data readiness

A written verdict on what your data can honestly support, before anything more is billed.

DAYS 3–7

Measurement

Estimation, falsification checks, and the per-SKU price list, with volume, revenue and gross margin costed at several price points rather than at the optimum alone.

DAYS 8–10

Delivery

The report on your data, the meeting deck in PowerPoint and PDF, the numbers workbook, and the store-test protocol. Then a working session where your team traces any figure back to the transactions behind it.

YOU KEEP

Everything

Every coefficient, interval and formula, the code that produced them, and the transformation logic from your raw extract to the model-ready table. Your team can rerun the whole thing, and that is the point.

What the study is worth

\$238k / yr

gross margin identified, at
assumed costs

15

products, each cleared, held or
refused with the reason

3 of 3

blind estimates matched before
the engine ran

0

moves taken on trust; every figure
traceable to transactions

Interested in what your data can support?

One export, one fixed scope, one fee agreed up front, and every figure checkable by your own team.

kevin@squ.solutions · squ.solutions

The method, in plain terms

THE ESTIMATING EQUATION

$$\log \text{ units (product, store, week)} = \beta_1 \cdot \log \text{ list price} + \beta_2 \cdot \log (\text{net} \div \text{list}) + \alpha (\text{product} \times \text{store}) + \gamma (\text{week}) + \varepsilon$$

β_1 the list elasticity: the response to a shelf-price change, the number a repricing decision uses.

β_2 the promotional elasticity: the response to discount depth, the number a promo-depth decision uses.

α one constant per store-product: each product compared against itself, made literal.

γ one constant per week: absorbs weather, holidays and traffic, so a national event never reads as price.

Elasticity, translated

-2 means: raise the price 10% and volume falls about 20%. The question is whether the extra cents per unit outweigh the lost units.

Why a regression

A line fitted through 169,677 store-product-weeks, product by product, separating price from the displays, seasons and store differences that move volume at the same time.

The 80% interval


The range we would bet on. A recommendation must hold across all of it, or it is not made.

The coefficient plot

One row per estimate: the dot is the point estimate, the bar its 80% interval.

The dashed line at -1 is where a price rise stops paying in revenue; the warm line at -2.5 is where it stops paying in margin at the assumed 40%, and it moves when real costs arrive.

Estimates wholly right of the margin line support a rise. Any interval crossing it is reported as a no-call, and the control row is the widest of all: weak evidence looks like this, and it earns a hold.



elasticity with 80% interval · clustered standard errors, per scope

The coefficient table

SCOPE	ESTIMATE	STD ERR	80% INTERVAL	P	CLUSTERS	CELLS
Cold cereal, list	-2.20	0.063	-2.28 ... -2.12	<0.001	77	169,677
Cold cereal, promotional	-2.89	0.034	-2.93 ... -2.85	<0.001	77	169,677
Cold cereal, merchandising controlled	-1.71	0.198	-1.96 ... -1.46	<0.001	15	169,677
Cold cereal, pure discount depth	-1.62	0.302	-2.00 ... -1.23	<0.001	15	169,677
Frozen pizza, list	-3.21	0.079	-3.32 ... -3.11	<0.001	12	111,176
Frozen pizza, merchandising controlled	-2.86	0.080	-2.96 ... -2.76	<0.001	12	111,176
Bag snacks, list	-1.79	0.229	-2.08 ... -1.49	<0.001	15	127,384
Bag snacks, merchandising controlled	-1.66	0.225	-1.95 ... -1.37	<0.001	15	127,384
Oral hygiene, control	-1.28	0.167	-1.49 ... -1.06	<0.001	13	116,504
Oral hygiene, control, merchandising controlled	-1.20	0.147	-1.39 ... -1.01	<0.001	13	116,504

Cereal rows are real measurements on this file: the first pair clustered by store (77 clusters), the merchandising-controlled pair by product (15 clusters, the more conservative choice with few products).

Every category, with its uncertainty and its cell count

CATEGORY	LIST ELASTICITY	80% INTERVAL	PROMOTIONAL	CELLS	PROVENANCE
Cold cereal	-2.20	-2.28 ... -2.12	-2.89	169,677	verified
Frozen pizza	-3.21	-3.32 ... -3.11	-3.88	111,176	verified
Bag snacks	-1.79	-2.08 ... -1.49	-1.44	127,384	verified
Oral hygiene (control)	-1.28	-1.49 ... -1.06	-1.61	116,504	verified

What a promotion is made of. A price cut, a display, a flyer photo, each recorded per store per week. For cereal the total promotional response is -2.9; the pure price cut alone is -1.6. The rest, some 44%, is exposure.

Which number answers which decision. The usual full event: -2.9. A discount without the flyer: -1.6. With merchandising held constant the cereal list response reads -1.71.

Cross-price, measured pair by pair

RIVAL TO DIGIORNO PEPPERONI	CROSS-PRICE	80% INTERVAL	READING
Freschetta Pepperoni	+0.84	+0.60 ... +1.07	strong substitute
Freschetta Brick Oven Italian	+0.70	+0.46 ... +0.94	strong substitute
Freschetta 4 Cheese	+0.62	+0.34 ... +0.90	strong substitute
Own-brand Rising Crust 3 Meat	+0.07	-0.01 ... +0.14	no link
Own-brand Rising Crust Pepperoni	+0.05	-0.02 ... +0.13	no link
Own-brand Rising Crust Supreme	+0.03	-0.04 ... +0.11	no link
Newman's Own Pepperoni	-0.41	-0.72 ... -0.12	moved together, not a finding

Two rivals were refused by name. DiGiorno Supreme and Three Meat move with Pepperoni at +0.98 and +0.97: one pricing decision wearing three labels. This category's noise floor is ± 0.22 , so anything inside it is reported as no link.

Three repairs, attempted and abandoned

Pool products, drop each shop-week

The standard repair, arithmetically impossible here: the rivals' average price collapses into own price, correlation exactly -1.0000 . Built anyway, it returned $+9.8$, not a number any business could produce.

Instrument with the other 75 shops

Sound in principle. But the chain prices centrally, so the other shops' price is this shop's price: estimates unchanged to three decimals.

Control for aisle promotion

The confirmed cause is shop-level display events; 100% of display weeks in this file are partial. Adding aisle intensity as a control moved nothing.

What replaced them, now standard: the placebo (cross-price fitted where the true answer is zero, pizza against mouthwash, measuring each category's noise floor in advance) and the agreement check (if A calls B a substitute, B must agree; four contradictions removed cereal cross-price after it was already written in).

One deliberate price action in three years

Prices below are a **fixed basket**: each SKU's own price, held at constant weights. An unweighted average rises whenever the mix shifts toward expensive lines with no price changing anywhere.

CATEGORY	SHARE	LIST PRICE	NET PRICE	DISCOUNT DEPTH	UNITS / TRADING WEEK
Cold cereal	53.8%	+4.3%	+5.8%	12.6% → 11.3%	-9.4%
Frozen pizza	23.1%	+2.5%	+2.7%	13.9% → 13.7%	+0.7%
Bag snacks	16.9%	+2.2%	+1.0%	5.3% → 6.4%	-8.5%
Oral hygiene	6.2%	-2.5%	-7.3%	7.6% → 12.1%	+7.7%

Cereal is the pricing event. Twelve of fifteen cereal SKUs were repriced in steps from February to August 2011, \$2.86 to \$3.03, and the level held.

Half of it is mix. The average selling price suggests +9.5%. The fixed basket says +4.3%; the rest is mix, not price.