

Three of four promotions cleared their hurdle. The store still lost money on them.

I derived a break-even model for vendor-funded discounts, tested four brands against it, and three passed. Then I asked a different question — not whether each brand's units moved, but whether the store made more money on the days it discounted — and the answer reversed. The unit of analysis was wrong, and it had been wrong in the direction that flatters promotions.

Ninety days of transaction data from the dispensary I run in Sherman Oaks, California — 4,449 tickets, 7,558 sales lines, April to September 2026. Brands are anonymized; they are live vendor relationships.

Wednesday \$530 Gross profit per day. No promotions running at all.	Saturday \$444 Three stacked cost-share promotions plus a daily 30% off.
---	--

The trade every vendor offers

A brand wants shelf position and volume, so it offers free goods on the purchase order — frontload — in exchange for the retailer running a deep discount on its products. The retailer's discount is real money off the top line. The vendor's free goods are real units at zero cost. The question is whether the second covers the first.

Nobody I buy from has ever presented that trade as arithmetic. It arrives as a percentage and a reciprocal favour. So I worked out what the break-even actually is.

The model

Buy N units at cost C , receive $F \cdot N$ free, sell everything at price P discounted by d . Hold gross profit equal to the undiscounted baseline of $N(P - C)$ and the whole thing collapses:

$$N(1 + F) \cdot P(1 - d) - NC = N(P - C)$$

$$(1 + F)(1 - d) = 1$$

$$F = d / (1 - d)$$

Cost drops out. The frontload needed to survive a discount does not depend on the margin of the product being discounted, which means the same answer applies to a 40% SKU and a 58% SKU — and which is why the intuition that "we can afford it on the high-margin stuff" is wrong.

Discount offered	Frontload required to break even
15%	17.6%
20%	25.0%
25%	33.3%
30%	42.9%
40%	66.7%

Read it in reverse and it gets uncomfortable: a 20% frontload — a good deal, by the standards of what actually shows up on a purchase order — supports a 16.7% discount, not the 30% that is typically asked for alongside it.

If the discount runs deeper than break-even, the promotion has to buy volume to compensate. At our 53% full-price margin, a 30% discount with a 20% frontload leaves gross profit at 70% of baseline, so the promotion needs **1.43× the units** just to draw level. That is the number to hold against any promotion, and it is the number nobody quotes.

The model assumes the free units sell through at the discounted price and that unit volume is otherwise unchanged. Both assumptions are generous to the promotion, which is deliberate — if the trade fails under generous assumptions, it fails.

The first test, and the answer I got wrong

I measured each brand's sales lines per 100 tickets on promotion days against non-promotion days, normalizing for traffic so a busy week couldn't masquerade as lift. Each brand's hurdle is its own break-even multiple, from the model above at its actual discount depth.

Brand	Promo days	Non-promo	Lift	Hurdle	Verdict
Brand A	30.86	18.15	1.70×	2.34×	misses
Brand B	39.91	21.44	1.86×	1.81×	clears
Brand C	8.60	2.23	3.85×	3.44×	clears
Brand D	8.27	3.48	2.37×	2.30×	clears

Lines per 100 tickets, 90 days. Note how thin two of the three passes are — Brand B clears by 0.05 and Brand D by 0.07. A result that survives only in the third decimal is a result asking to be checked another way.

Three of four cleared. Taken at face value, that says keep the promotions, drop Brand A, carry on. I very nearly wrote that up and sent it to the vendors as a scorecard.

The store-level test overturns it

The brand-level question is "did this brand's units move." The store-level question is "did we make more money." Those are not the same question, and only the second one pays anybody. So I put the promotional calendar against days of the week and looked at gross profit per day.

Monday		\$548
Tuesday		\$473
Wednesday		\$530
Thursday		\$543
Friday		\$518
Saturday		\$444

Gross profit per day by weekday. Wednesday, in red, is the only day carrying no promotions at all.

Day	Tickets/day	GP/ticket	GP/day	Margin	Discount rate	Promotions running
Monday	35.8	\$15.33	\$548	51.6%	19.7%	one BOGO
Tuesday	35.5	\$13.35	\$473	49.0%	16.9%	Brand A 30% off
Wednesday	36.0	\$14.71	\$530	51.0%	19.2%	none
Thursday	35.5	\$15.28	\$543	49.9%	17.0%	two
Friday	36.6	\$14.14	\$518	49.0%	21.2%	two, one at 40%
Saturday	33.2	\$14.40	\$444	47.6%	23.5%	three stacked, plus a daily 30%

The finding is in the first column, and it is the one I did not expect.

Traffic barely moves. Tickets per day sit between 33.2 and 36.6 across days carrying wildly different promotional loads — and the heaviest-discount day of the week is the *least* busy one. The customers the discounts were supposed to buy do not appear. What changes instead is gross profit per ticket, which is exactly what you would expect if the discounts are being harvested by people who were coming in anyway.

Saturday is the clearest case: three stacked cost-share promotions and a daily 30% off produce the lowest traffic, the lowest margin, the lowest gross profit and the highest discount rate in the week. Wednesday runs nothing and beats it by **\$86 a day**.

Measured against Wednesday, the three heaviest days give up about **\$155 a week** — roughly **\$8,000 a year**, or 4.5% of annual gross profit, on promotions that three of four brand-level tests said were working.

I'd rather state the weakness than have someone find it: promotions are not randomly assigned to days, and Saturday has its own traffic character that has nothing to do with discounting. This is observational, not an experiment. But the confound runs the wrong way for the promotional case — if discounts drove traffic, the most-discounted day should not be the quietest one in the week.

One brand ran the experiment for me

Brand E discounts 30% every single day. That is the "deep discount in exchange for frontload" trade in its purest form, running continuously for five months on a shelf

where nothing else does it — a natural experiment I did not have to design.

Brand E returns a **31.4% blended margin**, 25.8% on units that actually sold at a price, against a 38.7% discount rate. It is the worst-performing brand on the shelf. Meanwhile Brand F, which runs no promotion at all and carries a 4.8% discount rate, turned \$7,209 of net sales into **\$3,415 of gross profit** — more than double Brand E's contribution.

Two brands, opposite strategies, and the one giving away the most money is losing to the one giving away almost none.

Who is actually funding these

Vendors reimburse a contracted share of promotional discount, between 15% and 30% depending on the agreement. I built a claim system to collect it on a cost-of-goods basis, so I know precisely what comes back. Over sixty days:

Brand	Discount given	Credit recoverable	Recovery
Brand A	\$2,558	\$443	17.3%
Brand B	\$1,988	\$405	20.4%
Brand E	\$1,113	\$205	18.4%
Brand G	\$892	\$229	25.6%
Brand C	\$814	\$146	17.9%
Brand H	\$437	\$93	21.3%
Brand D	\$339	\$60	17.9%
All seven	\$8,141	\$1,581	19.4%

The retailer funds about **80 cents of every promotional dollar**. Across all brands, sixty days of discounts came to \$14,424 — roughly \$7,200 a month against \$15,100 a month of gross profit. Half of the store's gross profit is being handed back over the counter, four fifths of it out of the retailer's own pocket, in a program described on every sales call as vendor-funded.

The reframe: the value was never in the

discount

Chasing the cash credits led somewhere better. Free goods — units received at zero cost — have produced **\$15,163 of gross profit** since April. That is **24.8% of every gross-profit dollar the store has made**, off 13.3% of sales lines, and roughly four times what the cash credit program returns.

Verified against eight purchase orders: \$2,233 of free goods against \$12,156 paid, an **18.4% frontload**. One order was approximately 68% frontload — 97 free units on an \$835 purchase.

So put the measured frontload back into the model. An 18.4% frontload supports a discount of $18.4 \div 118.4 = 15.5\%$. The store was running 30% and 40%. The vendors were, on average, delivering about half of what their own asks required — and the gap is the \$8,000 a year the day-of-week table found.

That reframes the negotiation entirely. Cash credit is worth 19.4% of discount given and arrives two purchase orders later. Frontload is worth four times as much, arrives immediately, and carries no discount obligation at all unless one is agreed. The conversation to have with a vendor is not "what's your promo support" — it's "send more units."

And then the free goods leaked

Finding where the profit came from also exposed where it was going. Of 144 free units received across two purchase orders, 18 rang at a penny under an unrelated promotion and 8 sold at full shelf price. The rest were consumed without a record of who got them or why. That is approximately **\$830 of pure gross profit** — units with no cost against them at all — converted into roughly nothing.

Two adjacent leaks, same shape: buy-one-get-one promotions gave away \$634 of real cost of goods, and 105 sales lines with no promotion attached rang at about zero while carrying \$107 of cost. None of it was theft. All of it was a system that treats a zero-cost unit and a full-cost unit as the same object, and lets a discount rule reach a unit that was already free.

What I'd change in the product

Every finding above required exporting the data and rebuilding the question outside the system that holds it. None of it should have.

- A promotion should carry a funding source and a cost-share rate as first-class fields, so the platform can report discount net of vendor credit instead of gross. Right now

the retailer's own reporting tells them what a promotion cost and nothing about who paid for it.

- Free goods should be a modelled object, not a zero in the cost column. A receiving flow that marks zero-cost units, reporting that separates them from purchased inventory, and a rule that stops a percentage promotion from discounting a unit that is already free.
- Promotional reporting should be traffic-normalized by default. Raw unit counts make every promotion look successful, because discounting a product does increase sales of that product — that is not the question anybody is actually asking.
- The break-even belongs in the promotion builder. Enter a depth, get the frontload required and the volume multiple needed if the frontload isn't there. It's one line of arithmetic and it would change what retailers agree to on the phone.

The one that matters most is the cheapest: a single view of gross profit per day with promotional load overlaid on it. That one chart overturned four separate brand-level analyses and found \$8,000 a year. Every retailer on every platform has the data for it sitting in their own transaction table, and not one of them is shown it.

Limitations

One store, ninety days, a single market. Promotions are not randomly assigned to days or brands, so every comparison here is observational — the day-of-week result is a strong signal, not a causal estimate, and I would not present it as one. The break-even model assumes free units sell through at the discounted price. Sixty-day and since-opening windows are used where each is stated. The transaction export does not populate unit quantity, so all of this is line-based rather than unit-based, which will understate baskets where a customer took two of the same item. And the store's own house brand is recorded at zero cost on 71 lines, producing \$3,670 of gross profit that isn't real — enough to move house-brand margin from 58.7% to 52.3%, which is why it is excluded above and why I found it worth mentioning rather than quietly fixing.

Kyle Williams

Managing Partner, Pacific Wellness — Sherman Oaks, California. Retail operations, purchasing, financials, POS administration and analytics. Fifteen years in cannabis, from delivery in 2011 through a decade of wholesale brokering.

