

SOLERA CO.

Three patterns in this data are costing \$475K a year.

A Deep-Dive Analysis

24 months of data · 96,478 orders · 93,358 customers

Prepared by Real Analytical

Executive Summary

Three patterns in this data are costing the business about **\$475,000 a year** — roughly 6% of revenue. Each pattern has a specific dollar impact, a clear cause, and a concrete action.

Pattern one: paid acquisition is funding customers who never become profitable. One in three acquired customers — the lowest first-order tier — generates \$53 in lifetime revenue on average. At an industry-typical \$40 CAC, 98% of them fail to cover acquisition. This accounts for roughly \$300K of the total.

Pattern two: the brand's 3% repeat rate is the wrong number to track. The right number is the revenue weight behind it — and the unmeasured upside of moving it. Even a 1–2 point lift produces \$73K–\$250K a year.

Pattern three: about 16% of orders carry freight costs exceeding 50% of product price. Depending on freight absorption policy, this is \$55K–\$95K a year in recoverable margin.

A note on this sample. The analysis uses a representative DTC dataset of 96,478 orders over 24 months. The specific category and freight patterns would shift on a real Shopify brand. The framework — first-order tertile as the dominant LTV signal, revenue-weighted repeat as the true retention metric, mix-shift as the leading indicator — applies to any DTC dataset of this shape. The framework — first-order behavior, revenue-weighted retention, mix-shift detection, and low-intent acquisition — applies to any consumer or digital business of similar scale, not just Shopify brands.

Where to start this week. Pull the last 90 days of paid acquisition spend by campaign. Tag each campaign by the average first-order value it produced. Any campaign producing more than 40% low-tertile customers (<\$73 first order) is a candidate to pause or restructure. This audit alone — before any structural change — surfaces the 20–30% of spend most likely underwater.

The Data Picture

Solera Co. is a home goods brand selling through Shopify. The analysis covers 24 months — September 2016 through August 2018 — across 96,478 delivered orders, \$15.4M in revenue, and 93,358 unique customers. The brand sells across 70+ product categories. No single category exceeds 9.2% of revenue. The top five combine to 39%.

\$15.4M

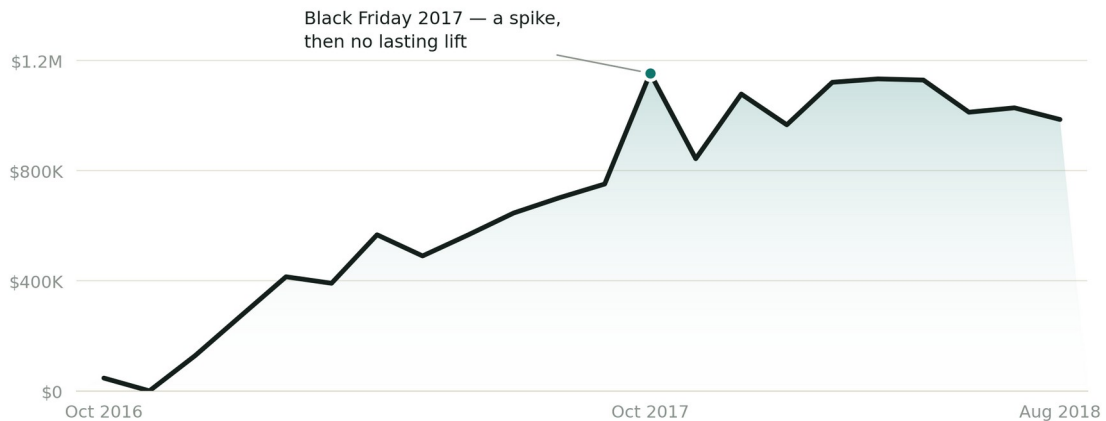
96.5K

93.4K

\$160

Revenue grew for 24 months. The repeat rate never moved.

Monthly delivered revenue



Analysis of Solera Co. data · Sept 2016 - Aug 2018 · n = 96,478 delivered orders

Revenue grew steadily for 24 months. Most DTC brands at this scale have one hero category driving 25–40% of revenue. Solera does not — its category mix is flat, more like a curated marketplace than a focused brand. That structural fact shapes every recommendation that follows.

The customer base. 93,358 unique buyers across the 24-month window. 90,557 of them — about 97% — placed exactly one order. The remaining 2,801 averaged 2.11 orders each. Repeat customers contribute 5.6% of revenue but average a smaller basket than first-time customers, \$146 versus \$161. The composition shapes the entire report. Most of the revenue is single-transaction. Most of the marketing effort that produced it is being asked to compound into retention, and it isn't.

The category mix. The top five categories combine to 39% of revenue. No category exceeds 9.2%. Most DTC brands at \$7M revenue have one category contributing 25–40% — a hero SKU that defines the brand. Solera's flat mix is structurally unusual and matters analytically. It means the findings that follow are framed in terms of acquisition patterns and customer behavior, not category strategy. There is no single SKU to optimize.

One observation worth holding onto. The largest revenue event in the dataset — Black Friday 2017 — produced a sharp spike with no sustained lift in baseline monthly revenue afterward. Acquisition events don't compound into retention on this dataset. The rest of the report quantifies why.

The Repeat-Rate Trap

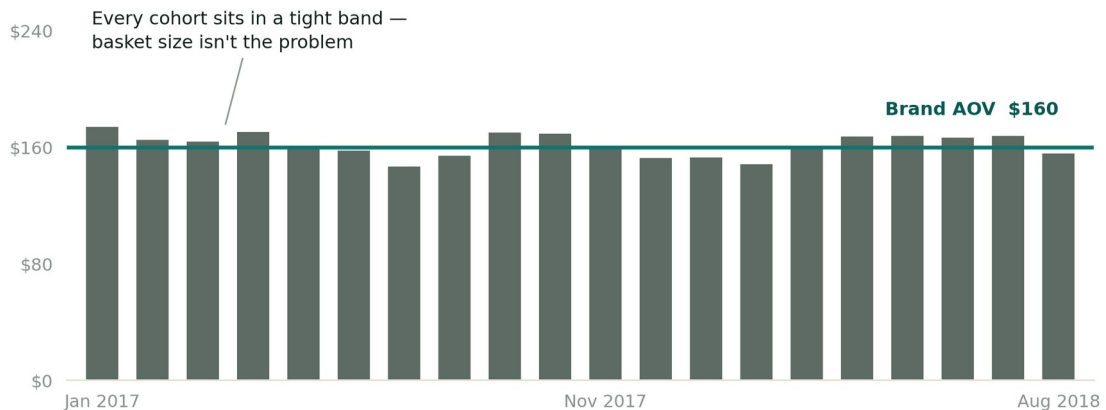
Your 3% repeat rate isn't the metric. The dollars behind it are.

The brand's repeat rate is 3%. The repeat customers generate 5.6% of revenue. More striking: their per-order basket is *smaller* than one-time customers' — \$146 versus \$161. The repeaters are valuable because they come back, not because they spend more when they do.

Of 93,358 customers, 90,557 placed exactly one order. The remaining 2,801 averaged 2.11 orders each. AOV by acquisition cohort is stable across the entire window — \$146 to \$175, well inside normal variance. The frequency, not the basket, is the lever.

First-order value is stable. Frequency, not basket, is the lever.

Average first-order value by acquisition cohort



Analysis of Solera Co. data · Sept 2016 - Aug 2018 · n = 96,478 delivered orders · cohorts < 100 customers excluded

What this is worth. Converting one-time customers into repeaters compounds quickly. A 1% conversion of the one-timer base produces **\$73K a year**. 2% — the realistic year-one target — produces **\$147K**. 5% is the aspirational ceiling at **\$368K**. Each scenario funds itself many times over. The brand isn't tracking any of them today.

Action. Track revenue-weighted repeat rate weekly. Run a 30/60/90-day reactivation campaign aimed at the 90,000+ customers who have placed exactly one order.

What Actually Predicts Lifetime Value

What customers buy first doesn't predict whether they come back. What they spend on the first order does.

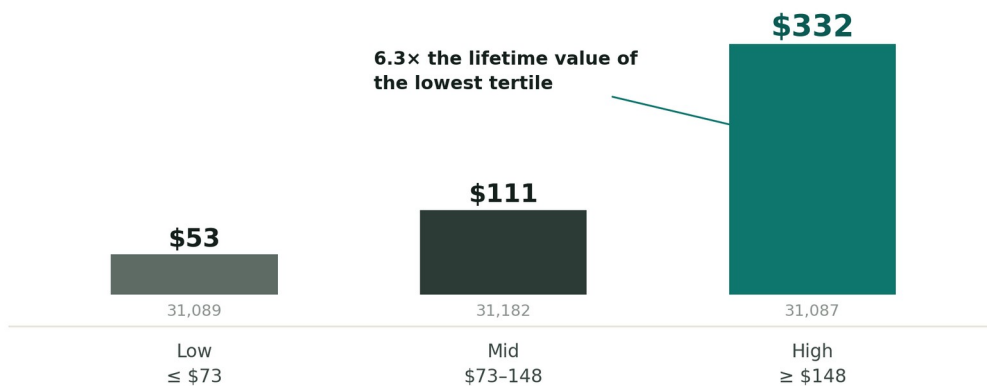
The obvious analysis would show that small electronics customers are worth \$368 and telephony customers are worth \$91. A 234% gap. The obvious recommendation would be to shift acquisition toward the high-value categories. The obvious recommendation would be wrong.

Repeat rates are virtually identical across all categories — 1.02 to 1.07 orders per customer. Categories don't differ in whether they produce loyal customers. They differ in what those customers spend on the first order. The category-LTV gap is just first-order price wearing a costume.

The data is clear about which dimension matters. Category accounts for 13% of the variation in customer value. First-order spend accounts for 26%. Combined, they explain 37%. Category is the weaker predictor by a factor of two.

First-order value predicts lifetime revenue. Category doesn't.

Average lifetime revenue per customer, by first-order tertile



Analysis of Solera Co. data · Sept 2016 – Aug 2018 · n = 96,478 delivered orders · variance explained: spend 26%, category 13%, combined 37%

What this is worth. Mid-tertile customers are worth roughly 2× low-tertile customers — \$110 versus \$53 in lifetime revenue. Shifting 15% of low-tertile acquisition into a 70/30 mid-high blend, through bundling, minimum-purchase thresholds, free-shipping tiers, and segmented landing pages, produces about **\$210K in annualized incremental revenue**. Shifting low to high is a much harder ask — those customers come from different channels and represent different intent — so high-tertile gains are a directional reference, not a target.

Action. Treat first-order value as the primary acquisition-quality signal — ahead of channel, audience, or category. Re-segment paid channels by the first-order AOV they actually produce.

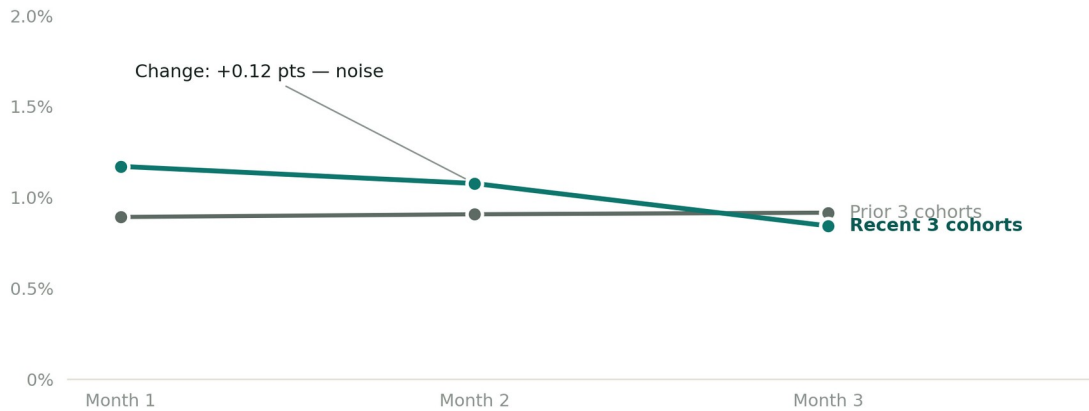
What You Should Be Watching

We expected to find a leading indicator. We didn't. That's good news — but only if you start watching.

Three standard early-warning metrics — 60-day cohort retention, first-order AOV trend, and tertile mix — all show stability or mild improvement across the last 18 months. Retention sits flat at ~1%. First-order AOV is up 4.3% comparing the most recent six cohorts to the prior six. Tertile mix has drifted slightly toward higher-value first orders. None of these movements is statistically significant. All are inside normal variance.

Repeat behavior is flat across cohorts. No trend to act on — yet.

60-day repeat rate, recent vs. prior cohorts



Analysis of Solera Co. data · Sept 2016 - Aug 2018 · n = 96,478 delivered orders

A note: this sample uses Brazilian e-commerce data, where 2017–2018 inflation ran 3–4% annually. Some of the nominal +4.3% AOV trend reflects price-level effects. On U.S. Shopify data, this confound would not apply.

The absence of a trend is not permission to ignore. It means the brand has no system to detect deterioration when it begins. There is no dashboard light that changes from green to yellow. There is no weekly review. Whatever breaks the current acquisition pattern — creative fatigue, an algorithm change, channel saturation — will not show up in the P&L for 60 to 90 days.

Three weekly metrics worth tracking:

- 1. Cohort first-order AOV (4-week rolling).** Alert if it drops more than 5% versus the prior 4 weeks.
- 2. First-order tertile mix.** Alert if low-tertile share rises more than 3 points versus the trailing 12 weeks.
- 3. New-customer count by tertile.** Tracks whether volume growth is coming at the cost of quality.

Cost of operating blind. If a 3-point tertile mix shift toward low-value customers went undetected for six months at current volumes, the revenue lost before anyone noticed would be about **\$63K**, exclusive of compounding effects on inventory, fulfillment, or budget allocation. That is the cost of not having a dashboard.

One in Three Customers Loses You Money

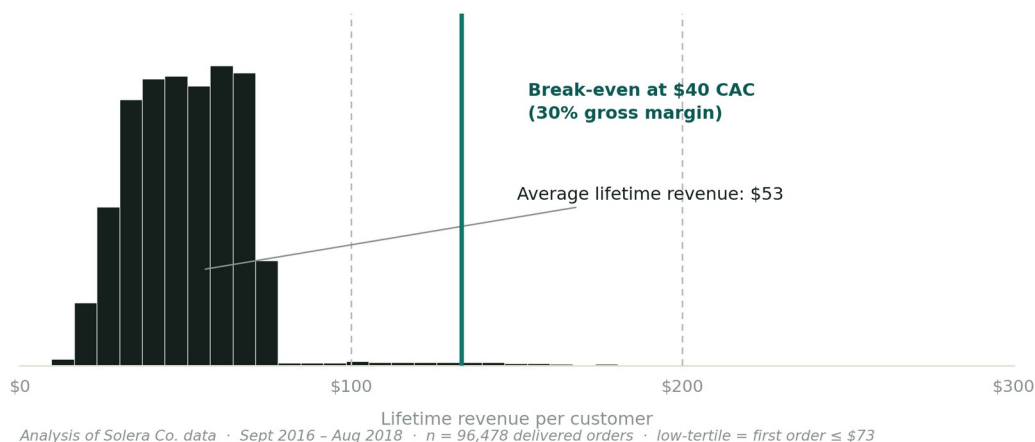
1 in 3 acquired customers loses you money. The other 2 don't make up for it.

There's a specific activity destroying margin: paid acquisition of low-intent customers. 98% of customers in the lowest first-order tier — those whose first order was \$73 or less — fail to generate enough lifetime revenue to cover their acquisition cost at any plausible CAC. The math is not marginal.

31,089 customers — exactly one-third of the customer base — fell into this tier. Their average lifetime revenue is \$53. At a 30% gross margin and an industry-typical \$40 CAC, each customer needs \$133 of revenue to break even. 99% fall below that threshold. Even at an aggressive \$30 CAC, 98% remain underwater.

At any plausible CAC, 98% of low-tertile customers never break even.

Lifetime revenue distribution — lowest first-order tertile



What this is worth. At the midpoint estimate — \$40 CAC, cut half the low-tertile acquisition — the savings are **\$188K a year**. At a \$30 CAC, the figure is **\$110K**. At \$60, it rises to **\$343K**. The figure scales linearly with actual CAC.

Why this keeps happening. Most founders measure CAC against blended AOV. Blended AOV looks healthy at \$160. But low-tertile customers generate one-third the average lifetime revenue — \$53 versus \$165. The blended number hides the problem because the high tertile (\$332/customer) mathematically masks the low tertile (\$53/customer). It is the wrong denominator.

Action. Within 30 days, audit paid acquisition by landing-page first-order AOV cohort. Pause or restructure campaigns producing more than 40% low-tertile customers — typically broad-audience promo campaigns, off-brand creative, and price-led offers. Reinvest the freed budget into channels and creative that produce mid+high tertile customers.

Prioritized Recommendations

Three actions to take, in priority order. Action 1 is the primary lever — start with its 7-day audit step. Action 1 bundles Insights 2 and 4 as a single acquisition restructure to avoid double-counting. Action 2 captures the repeat-rate opportunity. Action 3 builds the visibility infrastructure that protects the first two from drift.

#	Action	Annualized impact	Difficulty
1	Restructure paid acquisition around first-order AOV signal <i>(bundles the levers from Insights 2 + 4)</i>	\$250K – \$400K <i>midpoint \$300K</i>	Medium
2	Build repeat-rate measurement and lift system <i>(revenue-weighted tracking + reactivation campaigns)</i>	\$73K – \$250K <i>midpoint \$100K</i>	Easy–Med
3	Implement early-warning dashboard <i>(cohort AOV, tertile mix, new-customer count)</i>	\$30K – \$150K <i>midpoint \$63K</i> <i>(prevention)</i>	Easy

Sequencing. Action 1 leads despite Medium difficulty because its impact is roughly 3× the next action — and its first step, a 7-day audit, is accessible enough to start this week. Action 3 is the easiest to deploy and unlocks ongoing measurement for the others. Reasonable order: begin Action 1's audit immediately, deploy Action 3 in parallel over the first 14 days, then build Action 2's measurement as Action 3 lands.

Combined realistic-midpoint impact: \$400K–\$475K a year in current leakage recovery, plus ~\$63K a year in future drift prevented.

What to Track Going Forward

Most founders look at gross revenue, AOV, and total orders. Three additional metrics matter more:

- 1. Revenue-weighted repeat rate (weekly).** Not 'what percent of customers repeat' — but 'what percent of revenue comes from repeat customers.' This predicts margin durability six months out and surfaces small repeat lifts that the raw rate hides.

2. New-customer mix by first-order tertile (weekly). The single most important leading indicator of acquisition health. If low-tertile share rises more than 3 points versus the trailing 12 weeks, paid acquisition is degrading even if total customer count looks fine.

3. Freight-as-percent-of-price by order (monthly). Surfaces categories with structural margin issues. Any category where median freight ratio exceeds 30% warrants a pricing or shipping review.

Also worth doing. About 16% of orders carry freight costs exceeding 50% of product price. If the brand absorbs any portion of freight on these orders, that's \$55K–\$95K a year in recoverable margin. Minimum-order thresholds, freight-cost recovery, or pausing structurally bad categories all work. Test in 14 days, roll out in 30.

If you want this monitored automatically with weekly alerts and a single executive view, the next step is a Decision Dashboard. It is the closest thing to having a part-time analytics partner.

About

Real Analytical delivers decision-ready analytics for business owners who don't have a data team. You send us your data; we send back a written report identifying what's actually driving your growth and what to do about it. No retainers required, no engineering integration, no dashboards to learn unless you want one.

Our framework — first-order behavior as the dominant LTV signal, revenue-weighted retention as the true loyalty metric, mix-shift as the leading indicator, low-intent acquisition as the primary margin leak — is built on patterns common across consumer and digital businesses. We apply it to every engagement, surfacing the two or three patterns that actually move the P&L in a 10-page report you can read in 20 minutes and act on within a week.

If you want this analysis run on your own data, the Deep-Dive Report is \$399 (report-only) or \$499 (with a 20-minute walkthrough call). Five-day delivery. Up to two data sources. Reply to this email or visit realanalytical.com to start.

Get in touch

Email hello@realanalytical.com

Web realanalytical.com