CONSUMER DEMAND DATA

Methodology and results

The definitions, windows, and test results behind every claim in the companion materials, so a research team can reproduce each figure from the public data API. Demand is always measured before the company reported, so no result uses hindsight.

1. Scope and construction

All figures draw from the public data API. The equity universe is 471 companies with demand coverage, revenue history and a price file, about 424 of them tradable in an average month. The study window is 57 monthly return periods and 19 earnings quarters. Returns are measured as excess over an equal-weight index of the same universe, because the available index history was too short to serve as the benchmark. Returns are actual total returns, with dividends reinvested: nothing is dropped and nothing is capped. A company must trade at one dollar or more on the formation date to be held, on the price as it traded that day; once held, its return counts in full.

Survivorship. Ten of the 471 companies were delisted or acquired during the study and are held until their prices end. Nine more delisted or acquired companies in coverage lack revenue history and sit outside the reference universe. With those nine added, net Sharpe moves from 1.35 to 1.34 for the long/short, 1.31 to 1.25 for the long-only tilt, 1.10 to 1.07 and 1.30 to 1.23 for sector rotation, and 1.20 to 1.16 for the growth rule, and the pattern by year does not change. Companies that left coverage before the data was assembled cannot be added.

Core signal: for each company and metric, the trailing 28-day value versus the same 28 days a year earlier (a year-over-year change that controls for seasonality). The blended demand score standardizes purchase intent, search, and mentions each month and averages them.

2. Fundamentals

2.1 Demand to revenue

Per company, the demand score explains reported revenue growth with a typical correlation of 0.20 to 0.24. At the segment level, across 1,896 brand-and-segment series, the median correlation to segment revenue is 0.37, positive in 84 percent of series.

RELATIONSHIP	SERIES	MEDIAN R	SHARE POSITIVE
Company demand to revenue	approx 400	0.24	high
Brand/segment demand to segment revenue	1,896	0.37	0.84
Named brand matches (iPhone, HOKA, etc.)	177	0.23	0.66

Marquee brand matches run much tighter than the median: iPhone near 0.75, several consumer brands above 0.9.

2.2 Demand to revenue surprise

Ranking each quarter's reporters by demand versus the consensus-implied growth, the top group beats the bottom group on the actual revenue surprise by about 3.0 percentage points. The effect is positive in 17 of 19 quarters, with a Fama-MacBeth t-statistic of 5.96, and it holds after controlling for consensus growth, the prior quarter's surprise and the revenue trend: about half of the 3.0 points is demand's own contribution beyond those (t 3.97), and the rest comes from combining demand with consensus. Beating consensus is common in any quarter, as guidance and estimates are set conservatively. The spread has been close to zero in the last four quarters (between -0.4 and 1.0 points from 2025 Q4 to 2026 Q3), which we are watching. This is the single most robust equity-relevant result, though there is no reliable price move to capture around the print.

2.3 Operating KPIs

KPI FAMILY	SERIES	MEDIAN NOWCAST R
Payment / claim volume	3	0.67
Orders / backlog	2	0.58
Bookings / GMV / orders	28	0.56
Subscribers	approx 20	0.45

3. Equity signals

3.1 Single-name cross-sectional factor

Monthly rank correlation (information coefficient) of each demand metric versus next-month excess return, 57 months.

METRIC	IC	IC T-STAT
Purchase intent	0.034	2.96
Search	0.029	2.63
Mentions	0.020	1.94
Sentiment tone	0.007	1.06
Promotions	0.005	0.45

Purchase intent and search are significant. The blended composite has an IC of 0.030 (t 2.81) and holds within sectors (sector-neutral IC t 2.88), so it is not only a sector bet. Magnitudes near 0.03 are modest but tradable in a blend. It is not earnings momentum: with the latest revenue and EPS surprises, price momentum and one-month reversal as controls, demand keeps a Fama-MacBeth t of 3.1. The long/short does not reverse over the following twelve months. About two thirds of its return comes from the long leg, and neutral to sector, size and beta it runs at Sharpe 0.98.

3.2 Sector rotation

Ranking the 12 consumer sectors each month by aggregate demand predicts next-month sector excess returns with an IC of 0.135 and a t-statistic of 2.98, more than four times the single-name factor. Traded as a top-four / bottom-four sector long/short, it earns a net annualized 17.4% at Sharpe 1.10 (market alpha t 1.78, beta 0.35) and is positive in every calendar year (see 3.3). A broader top-six / bottom-six version is steadier: 15.5% at Sharpe 1.30. A 90-day signal window is weaker here (Sharpe 0.91). The table shows the top-four book by sub-period, including the weak trailing 12 months.

SECTOR ROTATION, TOP-4 / BOTTOM-4	NET SHARPE
Full sample (57 mo)	1.10

First half	1.77
Second half	0.85
Excluding the last 12 months	1.67
Last 12 months	0.09

3.3 Portfolio results

Net of 10 basis points charged on every dollar traded, over February 2022 to October 2026. Full construction detail and an alpha/beta/factor treatment are in the quant whitepaper.

BOOK	ANNUAL	SHARPE	MAX DD	POSITIVE YEARS
Demand long/short, market-neutral (quintile)	22.0%	1.35	-14%	4 of 5
Sector rotation (top-4 / bottom-4)	17.4%	1.10	-17%	5 of 5
Sector rotation (top-6 / bottom-6)	15.5%	1.30	-13%	5 of 5
Long-only top quintile (excess vs universe)	14.9%	1.31	-10%	5 of 5
Trend rotation (105 months from 2018, section 3.6)	14.9%	0.85	-15%	7 of 9

Against the S&P 500 (total return) the long/short has a beta of 0.18 and an alpha t-statistic of 2.55. On the Fama-French five factors plus momentum the alpha is 18.2% a year (t 2.81). A three-factor model built within this universe (S&P 500, a size proxy, price momentum) adds detail: it loads on momentum (beta 0.34, t 7.01) and the alpha is 11.4% a year (t 1.95), with about two fifths of the return shared with price momentum. With momentum removed from the signal before ranking, the long/short earns a net Sharpe of 1.01. That momentum factor earned 27% a year inside this universe, against 6% for the standard factor, which is why it absorbs more. Sector rotation loads less on momentum but carries a market beta of 0.25 to 0.35. Its market alpha has a t-statistic of 1.78 (top-four) and 2.25 (top-six); its standard-factor alpha, 1.77 and 2.13. With momentum removed its net Sharpe is 1.11.

3.4 Sub-period stability and the recent softening

The sector-rotation book was positive in every calendar year and the market-neutral long/short in four of five (2023 was -4.1%). 2024 was the strongest year; without it the long/short still runs at Sharpe 0.89, the long-only tilt 0.85 and sector rotation (top/bottom six) 1.05. The headline long/short runs at Sharpe 0.98 in the first half of the sample and 1.64 in the second. The trailing 12 months are weak at Sharpe 0.34 (6.0% annualized), with the gain concentrated in the final six months. Sector rotation shows the same pattern (trailing 12 months Sharpe 0.09) and the long-only tilt was flat. Over the same window the single-name rank correlation with price weakened while the link from demand to sector revenue strengthened.

SIGNAL, RECENT 12 MO VS PRIOR	PRIOR	RECENT 12 MO
Demand to sector revenue (IC)	0.282	0.628 (t 12.2)
Demand to price (IC)	0.177	0.063 (t 0.4)
Cross-sectional demand dispersion	0.211	0.315

> The data tracked sector revenue **more closely than ever** in the recent window while its pull on price faded. Dispersion rose rather than fell, which argues against crowding. The fade concentrated in non-consumer sectors and strengthened in the consumer core. Our reading is a market-regime effect rather than signal decay, but twelve months cannot settle that, and the pre-earnings surprise spread also narrowed in 2026.

3.5 Corroboration

Long-only, top-quintile by each construction, 57 months. Pick hit and pick gain are measured against the equal-weight universe; the Sharpe is on the excess return after costs. Requiring all three metrics to agree lowers breadth and raises per-pick quality; the blend retains breadth and has the highest Sharpe.

CONSTRUCTION	NAMES/MO	PICK HIT	PICK GAIN	SHARPE, NET
Mentions alone	80	48.6%	0.88%	0.98
Purchase intent alone	51	49.8%	0.45%	0.58
Search alone	84	48.4%	1.16%	1.04

Two of three agree	48	50.2%	1.21%	0.92
All three agree	9	54.9%	1.79%	0.67
Blended composite	85	49.5%	1.34%	1.31

3.6 Consumer trends: rotation, not stock selection

A company's mapped trend score does not pick stocks. From 2018 it has a net Sharpe of 0.42 as a stock ranker and -0.55 as a sector ranker, and adding it to the company demand score lowers the result. One construction works. Each month, rank the 141 trends by year-over-year growth of the blended 28-day signal and hold the companies mapped to the fastest-growing fifth against the companies mapped to the slowest fifth, each company weighted by its mapping to the trend. Over 105 months from 2018 that earns 14.9% a year net at a Sharpe of 0.85 (t 2.52), with a correlation of 0.40 to the single-name long/short and 0.16 to sector rotation. The Sharpe is 0.64 for 2018 to 2021 and 1.25 since. Its best month, January 2021, came mostly from GameStop; without GameStop in any month the Sharpe is 0.80. Its market beta of -0.24 has been a cost: hedged with past data only, the Sharpe is 1.22. The other five constructions are reported in the whitepaper (Section 9).

3.7 Concern spikes: alerts, not a signal

Every day, for each company and concern, we compare the share of the last seven days' posts that carry the concern with its level over the previous 90 days, leaving earlier spikes out of that baseline. A spike is three or more standard deviations above it with at least ten concern posts. The table holds every company that spiked for 21 trading days against the rest of the universe, as a calendar-time portfolio.

CONCERN SPIKE (3+ STANDARD DEVIATIONS, 10+ POSTS), HOLD 21 TRADING DAYS	RETURN VS REST, ANNUALIZED	T-STAT
Availability / stockouts	-5.4%	-1.29
Labor concern	-2.1%	-0.48
Product defects	-4.6%	-1.46
Political backlash	-4.1%	-1.43
Outage reports	-5.5%	-1.46

Negative sentiment	+2.1%	0.79
Price complaints	-2.1%	-0.59
Safety concern	-9.2%	-2.41
Any of the nine	-2.0%	-1.09

Across about 4,200 settings (count or share, window, spike size, minimum posts, holding period) the lean is negative but the number of significant results is close to chance, and more complaints do not mean worse returns. Three narrow alerts have been followed by weakness: safety spikes (about 0.7% behind the average stock over a month, every year); dips in happiness, the share of positive among positive and negative posts (about 0.7% behind similar-size stocks); and a company's first spike of six or more standard deviations with ten or more posts in a year, after which the stock fell 4.5% on average over 21 trading days, 64% fell and 77% lagged similar-size stocks. That last result rests on 39 events and was found by searching. These are alerts and system inputs, not a trading signal.

4. Macro nowcasting

Across 61 trends mapped to 41 official series, the median best-lag correlation is 0.25, with 27 trends above 0.3 and 15 above 0.5.

TREND	OFFICIAL SERIES	BEST SOURCE R
Losing a job	Initial jobless claims	0.87
Buying new clothes	Apparel retail sales	0.80
Grabbing coffee	Food services sales	0.70

A release-day trade on the index does not clear consensus. The value here is a real-time estimate of the official number before it prints, not an index trade.

5. What did not work

- Systematic earnings long/short: the stock's move around the report is not reliably related to the signal, and the event book loses money after costs.
- EPS surprise: no signal, in contrast to the revenue surprise.
- Release-day index trade on macro nowcasts: no edge over consensus.
- Category trend to matching stock basket: all tested categories above p 0.10.
- Concern spikes in general as a trade: no reliable effect (section 3.7). Three narrow alerts are kept for monitoring and as system inputs.

---

6. Overfitting assessment

The fundamental results do not depend on portfolio construction choices and reproduce the vendor's published figures, which places them outside typical overfitting risk. The equity-portfolio Sharpes carry real risk from a 57-month sample and multiple constructions tested, and should be discounted. A Sharpe of 1.35 on 57 months has a standard error near 0.48. Portfolio t-statistics use plain standard errors, since monthly returns do not overlap; Newey-West errors give higher t-statistics for every headline portfolio, so the plain figures are the conservative ones. The six strategies amount to three distinct bets: single-name demand, sector rotation and trend rotation, alongside the revenue forecast.

Points in favor of the signal being genuine: a strong economic prior (consumer demand to consumer-company revenue to stock), significance that holds under sector-neutralization, consistency across years for the long tilt, and modest, believable magnitudes rather than implausible ones.

Recommended before committing capital

- Reproduce every figure from the API using the definitions above.
- Freeze the demand-composite definition and paper-trade forward for two to three quarters.
- Check the cost model against your own execution (whitepaper Section 5.14 uses spread, square-root impact and borrow by size).

---

7. Limitations

History is about five years of returns. All demand is observed before the print, with no look-ahead. The 1,896 segment series co-move within a company, so the effective independent count is closer to 400. Returns are gross of financing and equal-weight against the universe benchmark. Equal weight leans on smaller companies; under a conservative cost model (spread, square-root impact, borrow by size) the long-only tilt in liquid names holds a Sharpe near 1.0 at $100M, and a liquid long/short with a 2% position cap and a holding buffer about 0.7 at $100M a side and 0.6 at $250M. The strongest, most construction-independent claims are the fundamental ones.

Execution timing. The portfolio figures in Section 3 trade at the month-end close on that day's data. Rebuilt from daily data with a one-day delay and averaged over every possible rebalance day, net Sharpe is 1.21 for the single-name long/short, 1.12 for the long-only tilt, and 0.87 and 0.98 for sector rotation (top/bottom four and six). Every start day was profitable for all four. Sector rotation is the most sensitive to the day chosen because it holds only twelve buckets. We also tested 69 daily entry and exit rules: rules based on year-over-year growth or on rank performed about as well as the monthly schedule, and moving-average crossovers of raw demand did not work. Report-time data confirms the earnings windows were dated correctly; with exact before-open and after-close timing there is still no tradable price edge around the report.
