

EMon Screeners

What each screen looks for, how it works, and what it is trying to achieve

EMon (Equities Monitor)

Built September 24, 2026 · live figures pulled from the EMon database at build time

Abstract

EMon scans roughly 7,032 U.S.-listed stocks every trading day, flags the ones that are behaving unusually, and in some cases turns those flags into paper options positions that it then tracks to an exit. This document explains every screener in the app: its inputs, its exact rules and thresholds, the idea behind it, and—where the app has a track record—how that idea has performed. It is written to be read next to the pages themselves; every page links to its own section here. Figures in the results tables are regenerated from the database each time the document is rebuilt.

Contents

1	Common ground: data and the daily volatility statistics	3
1.1	Where the data comes from	3
1.2	The daily volatility statistics	3
2	Volatility outliers	3
2.1	What it looks for: two screens	3
2.2	From a flag to an options position	4
2.3	Tracking and exits (options book)	4
2.4	What it hoped to achieve	5
2.5	Results so far	5
2.6	Known limitations	5
2.7	Strategy v2: model-directed stock trades	5
3	Volatility analysis	7
3.1	What it computes	7
3.2	What it hopes to achieve	7
3.3	Known limitations	7
4	Earnings leaderboard	7

4.1	What it computes	7
4.2	What it hopes to achieve	8
4.3	Known limitations	8
5	Scatson underperformer scanner (retired as a trade signal)	8
5.1	What it looks for	8
5.2	What it hoped to achieve	8
5.3	Why it is retired as a signal	8
6	Earnings volatility edge (experimental)	9
6.1	What it computes	9
6.2	What the evidence says	9
6.3	What it hopes to achieve	9
7	Sector leaderboard	10
7.1	What it computes	10
7.2	What it hopes to achieve	10
7.3	Known limitations	10
8	King valuation	10
8.1	What it computes	10
8.2	What it hopes to achieve	10
8.3	Known limitations	11
9	ETF rankings	11
9.1	What it computes	11
9.2	What it hopes to achieve	11
9.3	Known limitations	11
10	Other tools	11
A	Glossary	12
B	How this document is built	12

1 Common ground: data and the daily volatility statistics

1.1 Where the data comes from

All market data goes through one layer (`market_data.py`) that tries three sources in order:

1. **Alpaca** (real-time consolidated stock data and OPRA options data) for prices, daily bars, option expirations and option chains. The whole universe is loaded in batches of 100 symbols.
2. **Polygon** as the backup for all of the above, and the primary source for company fundamentals (P/E, market cap, price-to-book), which Alpaca does not provide.
3. **yfinance** only as a last resort, and for the one input neither of the others offers on the current plans: the earnings calendar and EPS estimates.

1.2 The daily volatility statistics

Every 15 minutes during the session, and again after the close, EMon computes for each stock over the last three months of daily closes ($N \approx 63$ bars; the latest bar is today's, so intraday values are “so far today”):

Log returns	$r_t = \ln(C_t/C_{t-1})$
Volatility	$\sigma = \text{std}(r_1, \dots, r_{N-1}) \cdot \sqrt{N}$
Prior volatility	$\sigma_{\text{prev}} = \text{std}(r_1, \dots, r_{N-2}) \cdot \sqrt{N-1}$ (the same, excluding the latest return)
Percent move	$m = P_{\text{last}}/C_{\text{previous close}} - 1$
Average volume / price	mean daily volume and mean close over the window

Important for reading the thresholds below: these volatilities are *not* annualized and *not* daily—they cover the whole three-month window. A stock whose typical daily move is s has $\sigma_{\text{prev}} \approx s\sqrt{62} \approx 7.9s$. So “prior volatility under 50%” means a typical daily move under about 6.3%, and “a move larger than the prior volatility” means a single-day move of roughly eight typical days.

These rows live in `daily_volatility_stats`; almost every screener reads them.

2 Volatility outliers

Page: Volatility Outliers · emails: “Volatility Outliers & Positions” (intraday and after the close)

2.1 What it looks for: two screens

Both screens require a three-month average volume above 1 million shares, so the name is liquid enough to trade.

Low-volatility movers	$ m \geq 15\%$ and $\sigma_{\text{prev}} < 50\%$
	<i>A big move in a normally quiet stock</i> —usually news: an earnings surprise, guidance, a deal, a downgrade.
Volatility-adjusted outliers	$ m \geq \sigma_{\text{prev}}$
	<i>A move that is large relative to how this particular stock normally trades</i> , whatever its absolute size.

A stock can pass both; the app records which screen(s) caught each name (`flagged_symbols.screen: low_vol, vol_adjusted or both`).

Why the two screens catch different things. Because σ_{prev} spans three months, the volatility-adjusted rule asks for roughly an eight-typical-day move. Names that pass it *but not* the low-volatility rule fall into two groups: very quiet stocks making moves under 15% that are still extreme for them (e.g. an 11% drop in a stock whose prior volatility is 10%), and very volatile stocks ($\sigma_{\text{prev}} \geq 50\%$) making enormous moves. The low-volatility rule, by contrast, is mostly ordinary-volatility stocks having a large, headline-driven day. These are different situations—a shock to a stable company versus an extreme day for an already-erratic one—which is why their results are tracked separately.

2.2 From a flag to an options position

The strategy is **contrarian (mean reversion)**: it bets that part of an outsized move will unwind.

1. **Expected give-back.** A “speculative move” $\hat{m} = -m \cdot |\sigma - \sigma_{\text{prev}}|$ (how much the latest return changed the window’s volatility, applied against the move), and a target price $P_{\text{target}} = P_{\text{last}}(1 + \hat{m})$. Moves above 198% are capped (the target defaults to 20% below the price).
2. **Direction and structure.** After a *drop*: a long call, or a short put if average volume exceeds 20M shares. After a *rise*: a long put, or a short call if average volume exceeds 15M.
3. **Contract.** The listed expiration closest to 30 days out, strike closest to the target price.
4. **Filters.** Skipped if the contract has no volume or volume ≤ 1 , if volume equals open interest, or (for short positions) if implied volatility is zero. Stock-only ideas appear in the email but are not tracked.

2.3 Tracking and exits (options book)

Every open position is re-priced each day and compared with its entry. Since September 24, 2026 the exits are (`data/outlier_strategy.json`, block `options_v1_exit`):

1. **Take-profit:** exit once the option is up 50% (sign-adjusted for shorts).
2. **Time exit:** exit after 20 tracked days.
3. **No stops.** The previous rules—exit at 3% against in the stock, at -16% on the option, or on “IV crush”—cut positions that later recovered: on real option prices they averaged -26% per trade versus -19% without them.

4. **Stale exit** (unchanged): from the Outlier PnL page, positions whose option price has not moved after two tracked days are closed as “2 days without movement”.

The options book still *loses* money after these changes (Section 2.7); it keeps running as a paper record alongside strategy v2.

2.4 What it hoped to achieve

Capture the partial reversal that often follows an overreaction, with defined risk (long options) or collected premium (short options on the most liquid names).

2.5 Results so far

Every position the scan has opened since 2023-10-01, one contract each, priced at option last trades. A re-flag of a contract that is still open counts once. Screens for flags made before the screen was recorded are inferred from the saved move and prior volatility. *Flat* means the option closed at exactly its entry price—often a sign it did not trade in between.

Screen	Positions	Closed	Win	Flat	Loss	Avg return	Median	Realized \$	Avg \$ / trade
<i>All time</i>									
Low-volatility movers	2,966	2,827	21%	27%	52%	−6.4%	−3.1%	−\$115,336	−\$41
Volatility-adjusted outliers	347	332	23%	29%	48%	−4.3%	+0.0%	−\$8,207	−\$25
Passed both screens	661	641	19%	28%	53%	−7.3%	−5.7%	−\$23,287	−\$36
<i>Since January 1, 2026</i>									
Low-volatility movers	992	929	19%	16%	65%	−13.7%	−15.1%	−\$77,786	−\$84
Volatility-adjusted outliers	77	72	31%	8%	61%	−1.7%	−15.1%	−\$6,790	−\$94
Passed both screens	212	199	16%	19%	65%	−13.2%	−16.7%	−\$14,593	−\$73

How to read this: a win rate near 20%, with more than a quarter of trades flat, means most of the P&L comes from a minority of positions. Compare screens on *average return* and *average \$ per trade*, not on totals—the low-volatility screen fires far more often, so its totals are naturally larger in both directions. The page shows the same figures for the last 90 days and a cumulative-P&L chart for each screen.

2.6 Known limitations

- Option prices are last trades, not bid/ask, so illiquid contracts can show stale “flat” results and real fills would be worse than the recorded prices.
- The stock-only ideas in the email (v1) are not tracked; strategy v2’s stock trades are (Section 2.7).

2.7 Strategy v2: model-directed stock trades

Page: *Volatility Outliers* (“Strategy v2” section) · email: “Strategy v2 (paper)”
code: `apps/outlier_strategy.py` · research: `research/outlier_opt`

What the history showed. Rebuilding the flags from daily bars since 2018 (9,593 flags on stocks over \$5) and pricing 3,384 of them on real option contracts since February 2024:

- The live options book has an annualized Sharpe ratio of about -2.5 (daily \$ P&L; -1.6 equal-weighted, -0.12 per trade). Fading rallies with long puts was the largest loss (-9.9% per trade, $t = -11.6$).
- The *stock* does carry a signal, but not “always fade”: up-moves give back about 1% over ten days and extreme spikes much more, while down-moves split into bounces and continuations.
- **Buying options loses for every direction rule, tenor (7–60 days) and strike (10% in to 10% out of the money)**—about -12% to -25% per trade after costs. A 30-day at-the-money straddle on a flagged name costs about 15% of the stock price while the typical 10-day move is 7%, and half bid/ask spreads on these contracts are 10–15% of the premium. Of 9,324 option configurations, only 2 were profitable in both the fitting period and the later test period, both marginally.
- The predictable part of the move is concentrated in names *without* liquid options ($+9.6\%$ vs $+1.0\%$ per trade), so the edge is expressed in the stock.
- Industry and sector did not help: adding them to the model lowered out-of-sample accuracy.

Direction. A ridge regression scores each flag with the expected ten-day return versus SPY *in the direction of the flag’s move* (positive: the move continues; negative: it reverses). Inputs come from the flag’s own bars: move size, prior volatility, move in typical-day units, how much of the move was an overnight gap, the 5- and 20-day run-up, distance from the 52-week high, price, SPY’s day, and an ETF flag, with up/down interactions. Reversal is favoured by high prior volatility, a run-up in the same direction over the prior week, a move alongside the market, and very large up-spikes; continuation by gap-driven moves, moves near the 52-week high, and moves that extend a trend. Fit walk-forward (each year predicted by a model fit only on earlier years), 2021–2026: correlation with the outcome 0.15; flags with $|\text{score}| > 2\%$ moved $+4.8\%$ the predicted way (61% of the time), positive in every year.

Entry, exit. Trade the stock in the predicted direction when $|\text{score}| > 2\%$, price $\geq \$5$, average volume $> 1\text{M}$, from the scans after 3:30 pm ET (the model uses end-of-day bars). Two entry filters, each checked in both halves of the out-of-sample record (2021–23 and 2024–26): *quiet* flags—a move of 15% or less, or prior volatility of 25% or less—earned about zero ($+0.2\%$ and $+0.7\%$), and *shorting a drop the model expects to continue* earned about zero ($+0.5\%$ and $+0.2\%$); both are skipped. Reversal calls carry nearly all of the edge ($+5.7\%$ per trade versus $+1.1\%$ for continuation calls). With the filters the remaining trades earned $+5.9\%$ each (75% hit rate, $+6.8\%$ / $+5.2\%$ by half, worst year $+2.3\%$). A higher score threshold trades off count for quality: above 4% the average is $+13.8\%$ on about 56 trades a year. Flags on an earnings day were -3.4% in 2024–26 but $+4.2\%$ in 2021–23, so they are not filtered but are worth watching. Exit at $+10\%$ or after 20 trading days, checked on daily closes. There is **no stop**: every stop tested lowered returns. Out of sample this exit earned $+5.0\%$ per trade with a 71% hit rate ($t = 8.3$, all six years positive). Accounts that cannot short stock can follow only the long signals ($+5.8\%$ per trade, but only 351 trades in six years).

Limits. Backtests before 2023 use today’s listed universe (delisted names are missing), bid/ask costs on options come from quotes measured on one day, and the v2 book starts empty on September 24, 2026—its live record is the real test.

3 Volatility analysis

Page: Volatility Analysis · run for the whole universe after the close

3.1 What it computes

One year of daily closes is split into 12 blocks of about 21 trading days (roughly one month each). For each block, $\sigma_k = \text{std}(\text{log returns in block}) \cdot \sqrt{\text{rows}}$. From these twelve monthly volatilities it stores:

<code>average_volatility</code>	mean of the 12 monthly volatilities
<code>current_volatility</code>	the latest month’s volatility
<code>minimum_volatility,</code> <code>maximum_volatility</code>	the range over the year
<code>against_prev_volatility</code>	latest month minus the month before (is volatility rising or falling?)
<code>volatility_ratio</code>	share of months whose volatility sits within one standard deviation of the mean—a measure of how <i>stable</i> the stock’s volatility is (higher = more predictable)
<code>time_period_low/high</code>	the 1-year low and high close
<code>net_profitability</code>	latest net income / revenue (fundamental context)

3.2 What it hopes to achieve

Size roughly one-month options positions with a sense of whether today’s volatility is high or low *for this stock*. The page labels a stock “Cheap” when its current monthly volatility is below its 1-year average and “Expensive” otherwise. Other screens (the earnings scanner, Scatson) reuse these columns.

3.3 Known limitations

Every figure is *realized* (historical) volatility, not option-implied volatility, so “Cheap” means the stock has been calm lately, not that its options are cheap.

4 Earnings leaderboard

Page: Earnings Leaderboard · fed by the after-close earnings scan

4.1 What it computes

For every stock with an upcoming earnings date, the scan looks back over its last four reported quarters and measures:

`avg_pre_earnings` average run-up over the 30 calendar days *before* each report
`avg_post_earnings` average move over the 3 calendar days *after* each report

Each set of four is cleaned of outliers (values more than two median-absolute-deviations from the median) before averaging. The scan also records the listed option expirations just before and just after the next report, and copies in the volatility-analysis columns (Section 3).

The leaderboard ranks names by average pre-earnings run-up, then post-earnings move (top 100, no thresholds).

4.2 What it hopes to achieve

Surface stocks that have habitually rallied into earnings, as candidates for a pre-earnings position, together with the option expirations that bracket the report.

4.3 Known limitations

Tested across about 12,600 earnings reactions on roughly 700 liquid stocks (2021–2026), the size and sign of the run-up into earnings did *not* predict the direction of the reaction (see Section 6). The ranking describes past behaviour; it is not a directional signal.

5 Scatson underperformer scanner (retired as a trade signal)

Pages: Scatson, Scatson Earnings · email: “Scatson Underperformer Scanner”

5.1 What it looks for

Stocks that reported in the last few business days and:

- **ran up into the report** — a positive close-to-close move over the 30 calendar days before it; and
- **dropped after it** — a negative move from the report date to the scan date.

Survivors are stored in `scatson_cache` with their volatility-analysis context and shown on the page and in the email.

5.2 What it hoped to achieve

A *bounceback*: the idea was that a stock the market bid up before earnings and then sold on the print is oversold, and long-dated calls (about 270 days) struck near a volatility- and range-based target price would profit as it recovered.

5.3 Why it is retired as a signal

Across 12,600 earnings reactions (2021–2026), stocks that ran up 3%+ into earnings and then fell 5%+ ($n = 1,093$) showed no bounceback: the market-adjusted 20-day return was -0.3% ($t = -0.9$) and the 60-day return $+1.0\%$ ($t = +1.2$), both indistinguishable from zero, with negative

medians. The scanner still runs as a watch-list, and the page and email carry this note; the long-call suggestions were removed.

6 Earnings volatility edge (experimental)

Page: *Earnings Vol Edge* · scans 2:45–3:50 pm ET for names reacting on the next session

6.1 What it computes

For each liquid stock whose earnings reaction happens on the next session:

1. **Expected move** from the stock's own past reactions: a small model on the log of its average absolute move over the last 8 and last 4 reactions and its recent realized volatility, fitted on 18,646 reactions of liquid stocks (from 68,688 recorded reactions across 4,416 stocks) and calibrated so its average prediction matches the average realized move.
2. **Implied move** from the options market: the at-the-money straddle for the first expiry after the reaction, converted to an expected absolute move, with the ordinary (non-earnings) days' variance stripped out.
3. **Ratio** implied/expected: below 0.8 is labelled *cheap*, above 1.15 *rich*. Names whose first expiry is more than 8 days after the reaction are shown but not traded.
4. **Trade ideas**: defined-risk iron condors (sell premium, buy wings), valued by simulating the stock's historical distribution of earnings moves, with bid/ask spread costs. Only ideas with an expected return of at least 15% on capital at risk are shown. Ideas are paper-tracked and scored after the reaction.

6.2 What the evidence says

- **Direction is not predictable** from run-up, the previous reaction, or post-earnings drift.
- **Size is predictable** (rank correlation ≈ 0.42), but the option market already prices it: across 2,188 events priced strike-by-strike, implied moves averaged only about $1.06\times$ realized, and the implied/expected ratio did *not* predict outcomes.
- Buying volatility lost money in every bucket. Selling defined-risk premium was about break-even before costs and lost roughly 10% after realistic spreads on average.
- Only the ideas the valuation engine rated at 15%+ expected return were positive (+6.6% net, 70% win, $n = 82$, $t = 1.2$)—encouraging but not statistically proven.

6.3 What it hopes to achieve

A disciplined, non-directional way to find the occasional earnings event where options look mispriced after costs—and a live, paper-tracked test of whether that edge is real. Expect few ideas; most days have none.

7 Sector leaderboard

Page: Sector Leaderboard · also included in the volatility-outliers email

7.1 What it computes

For each sector, every stock's weight is its dollar volume today, and the sector's move is the weighted sum of its stocks' moves:

$$w_i = \frac{V_i P_i}{\sum_{j \in \text{sector}} V_j P_j}, \quad \text{sector move} = \sum_{i \in \text{sector}} w_i m_i.$$

Sectors are ranked by that move, and each sector lists the ten stocks contributing most to it.

7.2 What it hopes to achieve

Show where money is actually moving today—which sectors are leading or lagging and which names are driving them—as context for the outlier flags (is a flagged stock moving with its sector or against it?).

7.3 Known limitations

Weights are computed before rows with missing data are dropped, so a sector with missing rows is slightly under-weighted.

8 King valuation

Page: King Valuation · snapshots taken after the close

8.1 What it computes

For each stock, its trailing P/E and price-to-book relative to its sector's market-cap-weighted average:

$$\text{P/E vs sector} = \frac{P/E_i}{\sum_j MC_j P/E_j / \sum_j MC_j}, \quad \text{P/B vs sector defined the same way.}$$

Below 1.0× means cheaper than the sector; above means richer. The table is sortable by either measure.

8.2 What it hopes to achieve

A quick relative-value map: which companies trade at a discount or premium to their own sector.

8.3 Known limitations

Companies without a trailing P/E (for example, loss-makers) are skipped. A cap-weighted average of P/E is close to, but not the same as, the sector’s aggregate P/E (total market cap over total earnings).

9 ETF rankings

Page: ETF Rankings

9.1 What it computes

Every exchange-traded fund in the daily volatility statistics for the chosen date, ranked by its percent move and grouped by asset class and leverage (e.g. 2×, 3×, inverse). There is no further scoring.

9.2 What it hopes to achieve

See at a glance which asset classes, themes and leveraged products moved most—useful for hedging and for confirming whether a stock’s move is part of a broader theme.

9.3 Known limitations

Coverage of the daily ETF scan has been thinner since late 2025; the page flags thin dates as “Sparse data”. The “inverse” flag stored for ETFs appears to be reversed at insert time and should be checked before relying on it.

10 Other tools

Long straddle (per ticker). For liquid names with elevated three-month volatility, prices a straddle around the 30-day mean price or the current price and estimates the chance of reaching break-even from the stock’s volatility.

Short strangle (email). The mirror image for calm, liquid names: sells a strangle with strikes about half a volatility away from the price, betting the stock stays in range.

Hedge positions (API). Given a position, suggests hedges from correlations, ETFs, options and news sentiment.

Earnings results. Reference tables of reported and upcoming earnings with EPS estimates and surprises.

Query console. Read-only SQL access to the EMon database.

A Glossary

Implied move. The absolute move the option market prices for an event, read from the at-the-money straddle (for a normal distribution, straddle $\approx 0.8 \times \text{price} \times \text{volatility}$ over the period).

Iron condor. Sell an out-of-the-money call and put, buy further-out ones as protection: profits if the stock stays in a range, with a capped loss.

Mean reversion. The tendency of an extreme move to partly reverse.

Realized vs implied volatility. Realized: how much the stock has actually moved. Implied: how much the option market expects it to move.

Short/long bias. Long = buy the option; short = sell (write) it and collect the premium.

B How this document is built

The source is `docs/emon_screeners.tex`. Running `python3 docs/build_screeners_doc.py` pulls the live figures (5,901 flags from 2023-10-01 to 2026-09-24; 4,021 options positions) from the database, compiles the PDF with Tectonic, and publishes it to every page of the app.