

A Concise End-to-End Strategy Overview

Egor Shiryaev

Jupyter notebook: End-to-end GH

1 Trading Hypothesis and Motivation

Hypothesis: liquid Russian equities exhibit persistent medium-term trends. Combining *trend-following* and *volatility targeting* improves the portfolio's risk profile relative to passive buy-and-hold.

Expected effect:

- participation in directional market moves;
- reduced risk during periods of high volatility;

Economic Rationale

- **Momentum:** assets with positive recent performance often continue to trend.
- **Volatility clustering:** sizing the position via a *target vol* rule stabilizes risk.

Implementation Idea

The strategy uses a trend filter and volatility-target position scaling to achieve a more stable risk-adjusted return relative to a passive equal-weighted portfolio.

Criteria for Production Readiness

1. Out-of-sample robustness

- stable results across walk-forward periods;
- a controlled level of drawdowns.

2. Robustness to costs

- the strategy remains viable under realistic commissions and spreads;

3. Parameter robustness

- small changes in *lookback* / *vol-target* should not break the strategy;

4. Risk control

- bounded *max drawdown*;
- stable portfolio volatility;

5. Practical feasibility

- liquid instruments are used;
- a reasonable rebalancing frequency.

2 Data Acquisition and Retrieval

Data was loaded via T-Bank's *Tinkoff Invest API*: historical candles for each instrument, and dividends.

The project used both synchronous and asynchronous loading modes. The asynchronous path speeds up data collection across multiple instruments: while one request is waiting for a response from the API, the program can issue other requests.

3 Price Data Processing

After receiving data from the broker, the price series must be brought into a uniform, backtest-ready form. In this project, processing consisted of two main steps:

- handling gaps in the price series;
- adjusting prices for dividends.

Handling Gaps

Gaps were handled using the `ffill().dropna()` scheme:

- `ffill()` — fills gaps with the last available price;

`ffill()` was chosen because it uses only information available at the time of observation and does not leak data from the future. This matters for a correct backtest and for computing trading signals.

The `bfill()` approach (filling with a future value) was not used, since it can introduce *look-ahead bias*. For the same reason, the `ffill().bfill()` scheme was not applied: although it formally removes more gaps, it can fill values with future prices and make the backtest results methodologically invalid.

Dividend Price Adjustment

An additive dividend price adjustment was used to correctly compute price returns. Without this adjustment, ex-dividend dates create downward jumps in the price series that are unrelated to market movement and can distort returns, volatility, and trading signals.

Adjustment Concept

Let P_t be the price at time t , D_u the per-share dividend paid at time u , and P_t^{adj} the adjusted price. The additive adjustment is then given by

$$P_t^{adj} = P_t - \sum_{u>t} D_u.$$

In other words, the sum of all future dividends is subtracted from the historical price to remove the effect of ex-dividend dates from the price series.

Implementation Details

The implementation used per-instrument dividend data, with future payouts accumulated along the timeline. Stock splits were also accounted for: dividends paid before a split date were rescaled by the split ratio so the adjustment stays consistent with the current price scale.

A Note on Validation

If desired, the adjustment can be cross-checked against external sources (e.g., *yfinance*), but the absolute levels of adjusted prices may differ due to differences in how dividends and corporate actions are handled. It is therefore more appropriate to compare derived quantities, chiefly returns (`pct_change()`), rather than price levels.

4 Building the Benchmark: an Equal-Weighted Portfolio

The benchmark was a passive equal-weighted portfolio of the selected liquid stocks. This benchmark lets us compare the strategy's results not against a single instrument, but against a diversified buy-and-hold alternative on the same set of names.

Construction Principle

On each rebalancing date, capital is split equally across all assets. If the portfolio holds N instruments, each instrument's weight is

$$w_i = \frac{1}{N}, \quad i = 1, \dots, N.$$




Figure 1: Initial weights of the equal-weighted benchmark portfolio.

Preliminary Analysis of the Benchmark Portfolio

After building the benchmark portfolio, a basic diagnostic analysis of its properties was carried out to establish a reference point for comparison with the strategy.

Distribution of Daily Returns. The histogram of the benchmark portfolio's daily returns lets us assess the shape of the distribution, any skew, and the tail behavior.




Figure 2: Distribution of the benchmark portfolio's daily returns.

QQ-Plot of Daily Returns. The QQ-plot is used to visually check for deviations of the return distribution from normality.




Figure 3: QQ-plot of the benchmark portfolio's daily returns.

Based on this preliminary diagnostic, we can assess the shape of the return distribution and possible deviations from normality.

Correlation Structure of the Assets. The correlation matrix shows the degree of co-movement between the stocks and lets us assess the benchmark portfolio's actual diversification.




Figure 4: Correlation matrix of the benchmark portfolio's asset returns.

5 Trading Strategy: Trend Filter and Volatility Targeting

The strategy is built as an overlay on top of the benchmark portfolio (the equal-weighted stock portfolio). Rather than picking individual names, the strategy manages exposure to the benchmark: it increases or decreases the fraction of capital held in the base portfolio depending on the market regime and current volatility.

Core Idea

Two simple, interpretable rules are used:

- **trend filter:** risk is held only when the benchmark’s equity curve is above its moving average;
- **volatility targeting:** position size is scaled to maintain the strategy’s target volatility level.

The strategy thus combines filtering out unfavorable market regimes with risk control via dynamic exposure management.

Formalization

Let $r_t^{(b)}$ be the benchmark portfolio’s return, and x_t the strategy’s exposure to the benchmark.

1) Trend filter. The benchmark returns are used to build an equity curve:

$$E_t = \prod_{u \leq t} (1 + r_u^{(b)}).$$

If E_t is above its moving average, risk is “on”; otherwise exposure is reduced.

2) Volatility targeting. The realized volatility estimate $\hat{\sigma}_t$ is computed over a rolling window. The target exposure is given by

$$x_t \propto \frac{\sigma^*}{\hat{\sigma}_t},$$

where σ^* is the target annualized volatility. In practice, exposure is capped from above (a leverage cap).

3) Strategy return. For a correct backtest, the lagged exposure x_{t-1} is used (no look-ahead bias). The strategy’s return then combines the benchmark return and the cash return:

$$r_t^{(\text{strat})} = x_{t-1} r_t^{(b)} + (1 - x_{t-1}) r_t^{(\text{cash})} - c_t,$$

where c_t is the rebalancing cost, proportional to the change in exposure.

Diagnostics of Strategy Returns

As a first diagnostic of the strategy’s results, we examine the distribution of daily returns. The histogram lets us visually assess the concentration of observations near zero, the skew, and the tail behavior.



Figure 5: Histogram of the strategy's daily returns.

The high concentration of values near zero reflects periods of reduced or zero exposure to the benchmark portfolio (for example, when the trend filter is off), as well as the effect of exposure smoothing and the cash component.

6. Walk-Forward Validation with an Embargo

To assess the robustness of the results, we used walk-forward out-of-sample validation over the time series, with an embargo gap between the training and test windows.

At each step, a historical window (L_{back}) was formed, followed by an embargo gap (L_{emb}), after which performance was evaluated on the subsequent test interval (L_{fwd}). The window was then shifted forward and the procedure repeated.

The embargo is introduced to reduce information leakage between adjacent windows and to lower spurious correlation between out-of-sample metrics at neighboring validation steps.




Figure 6: Diagram of walk-forward validation with an embargo gap.

Computing Out-of-Sample Metrics

After running the walk-forward validation, key performance metrics were computed for each out-of-sample window: *Total Return*, *Annualized Return*, *Annualized Volatility*, *CVaR*, *Sharpe Ratio*, and *Drawdown*. This lets us look not only at the average result but also at the spread of metrics across different market regimes.

For the final diagnostic, the distribution of these metrics was plotted across all OOS windows, giving a more robust picture of the benchmark portfolio's behavior over time.




Figure 7: Distribution of out-of-sample metrics

7. Backtest Results: PnL, Equity, and Comparison with the Benchmark

Once the strategy's rules were specified (trend filter, volatility targeting, cash component, and costs), a backtest was run on the historical return series of the benchmark portfolio. From the backtest, the strategy's daily returns $r_t^{(\text{strat})}$ were computed, from which the equity curve and cumulative PnL were built.

The following views were used to interpret the results:

- the strategy's cumulative PnL;
- the benchmark portfolio's cumulative PnL (*buy-and-hold*), for comparison;

- cumulative rebalancing costs;
- the strategy's equity curve (*Market Value*);
- intra-month PnL reset to zero at the start of each month (*monthly-reset PnL*).

The *monthly-reset PnL* chart is used as a more readable view of the strategy's local dynamics: it lets us visually assess intra-month moves, drawdowns, and recoveries without the compounding effect of a long horizon.




Figure 8: Strategy backtest results: cumulative PnL, equity curve, and monthly-reset PnL.

8. Final Results, Metric Comparison, and Conclusions

In the final step, the key metrics of the strategy and the benchmark portfolio were computed and compared (using, among other things, the *quantstats* library): return, volatility, Sharpe ratio, drawdown, and return-distribution characteristics. This comparison lets us assess not just the absolute result, but the quality of the strategy's risk relative to the passive approach.

Conclusion on the Trading Hypothesis. This study tested the hypothesis that combining trend-following with volatility targeting can improve the risk profile of a portfolio of liquid Russian equities relative to passive buy-and-hold.

The backtest and walk-forward analysis results show that the strategy:

- exhibits a more stable return distribution;
- has a controlled level of drawdowns;
- is sensitive to transaction costs, but remains robust within reasonable ranges;
- remains effective on out-of-sample periods.

Using volatility targeting stabilizes the portfolio's risk, while trend filtering reduces participation in unfavorable market regimes.

As a result, the strategy can be considered a candidate for further testing in *paper trading* mode and subsequent production use, subject to risk limits and regular metric monitoring.