Random Harvester Technical Appendix

Methodology and diagnostic reference | Updated 2026-10-05

This appendix explains the calculations, diagnostic measures, execution conventions, and limitations used in Random Harvester portfolio reports, as well as the academic literature upon which they are based. It is a general reference and contains no customer portfolio or performance data.

System and Causality

Random Harvester is a causal, single-pass phase detector and long-only trading simulator. It processes observations from left to right. A calculation at time t may use information available up to t, but not data from t+1 or later. It does not use centered smoothing or retrospective signal placement.

The detector does not forecast a future price. It tests whether recent log-price behaviour is sufficiently trend-like to treat the current phase as predictable, and withdraws when that statistical fit breaks down. The principal fit channel is a maximum-likelihood regression of log price on time, filtered through RMSENorm (normalized root mean square error) as a scale-free measure of scatter around the fitted trend.

Execution Convention

A decision produced on bar t is executed on the next available trading bar. The report's execution ledger uses the recorded portfolio_data action rows; it does not reconstruct trades from chart prices. Transaction costs, taxes, position changes, banked profit, and investor wealth are applied at execution. Pending, skipped, or suppressed actions are not counted as trades.

Prepared Price Series

Where a provider supplies adjusted close, Random Harvester transforms OHLC (Open, High, Low, and Close prices) consistently into a prepared research series so historical splits and distributions do not create artificial jumps. The detector remains causal on that prepared series. Vendor-adjusted histories can still be revised after corporate actions.

Foreign-currency conversion is maintained separately from the price provider. Reports should therefore be read in the stated portfolio currency, while remembering that both asset-price and exchange-rate histories may affect a converted result.

Banked Profit and Investor Wealth

Banked profit is realised, protected cash removed from market risk. It is not reinvested. At time t, total investor wealth is W(t) = trading capital(t) + banked profit(t). Banked profit does not imply that the original investment has been fully recovered; the report states separately how much of the original seed has been recovered as protected cash and how much market value remains exposed.

A no-skimming baseline runs the same trading mechanism with the skimming rate set to zero. Comparing W with that baseline isolates the effect of the banking policy from the phase detector's buy and sell decisions.

Risk Channels

ATR (Average True Range, implemented as a trailing average) is a causal risk channel, not a directional signal. True range percentage and its trailing average may condition stop width, position size, and the rate at which gains are banked. ATR is never used to decide whether price direction is bullish or bearish and is not wired into buy or sell gates.

Bear mode is a separate long-only protection mechanism. It can suppress new purchases during persistent negative phases and permit a correction purchase only when the detector's stated recovery conditions are met.

Benchmarks and Comparators

The standard report compares total investor wealth with a constant-mix 60/40 equity-bond reference, a holding of the configured equity index, equal-weight buy-and-hold of the submitted assets, and the same Random Harvester trading rules without profit skimming. These are comparisons, not targets. All series use the same evaluation calendar and are normalised to a common starting investment for charts.

Diagnostic Reading Guide

Category	Measures	What it answers	Main limitation
Return	Total return, CAGR	Growth over the evaluation period and its annualised equivalent.	Short windows make annualised values unstable.
Total risk	Volatility, Sharpe	Dispersion of returns and excess return per unit of total volatility.	Sharpe can mislead when both compared returns are negative.
Downside risk	MDD, Calmar, Sortino, Ulcer, Martin	Depth, duration, and downside-adjusted return.	One severe episode can dominate MDD; inspect underwater duration too.
Benchmark relative	Active return, tracking error, information ratio	Average excess return and its consistency relative to the reference.	A high ratio from a short or quiet period is weak evidence.
Market exposure	Beta, correlation, up and down capture	How strongly the strategy participates in benchmark movements.	Low beta can make alpha look favourable despite negative active return.
Trade quality	Win rate, Wilson CI, RRR, break-even rate, edge	Whether completed trades show a repeatable statistical advantage.	Fewer than 20 completed trades is treated as descriptive only.
Cross-asset consistency	Homogeneity test	Whether observed win rates plausibly share one underlying probability.	A pooled result can conceal weak assets when performance is heterogeneous.
Capital protection	Banked profit, wealth permanence	Cash permanently removed from market risk and its share of current wealth.	Banking reduces capital available for future compounding.

Statistical Evidence Rules

- Parameter settings are fixed before a customer request and are not tuned to the submitted assets or dates. This prevents parameter-search overfitting within that customer backtest.
- Customer asset selection limits the scope of the result. It is not parameter overfitting by Random Harvester, but it prevents the report from claiming universal performance.
- Completed trades are Bernoulli observations for the binomial framework. Point estimates are reported with Wilson confidence intervals because small samples require explicit uncertainty.
- A minimum of 20 completed trades is required before the report issues a statistical edge verdict. Below that threshold, trade outcomes remain descriptive.
- Backtest evidence does not replace live forward testing. Vendor revisions, market impact, liquidity, slippage, taxes, and execution constraints can differ from the simulation.

Formula Reference

- **cagr:** Compound Annual Growth Rate; the mean growth rate of an investment if compounded annually. $CAGR = \frac{E_T^{\frac{1}{Y}}}{E_0} - 1$
- **drawdown:** Percentage decline from the running peak equity to current equity. Max drawdown is the maximum such decline over the window. It reflects the actual capital impairment an investor would have experienced before recovery. $DD_t = \frac{E_t - max\, E_{s \le t}}{max\, E_{s \le t}}$
- **time_in_market:** Fraction of days with a non-zero position (Positions > 0).
- **banked_profit:** Cumulative harvested cash paid out (net of fees/taxes). Contributes to equity/ROI but is not reinvested.
- **Tracking Error:** Annualized standard deviation of active returns (RandomHarvester minus benchmark). $TE_{ann} = \sigma(a_t)\sqrt{252}$
- **Information Ratio:** Active return per unit of active risk (tracking error). It measures the efficiency/consistency of benchmark-relative outperformance, distinguishing systematic skill from random benchmark deviation. $IR = \frac{AR_ann}{TE_ann}$
- **sharpe_ratio:** Risk-adjusted return using excess return over the risk-free rate, per unit of total volatility. $Sharpe = \frac{E(r_t - r_f,t)}{\sigma(r_t - r_f,t)}\sqrt{252}$
- **Jensen's Alpha:** The OLS intercept from regressing strategy excess returns on benchmark excess returns, annualised. Represents the return earned above what the strategy's beta exposure to the benchmark would predict. Positive alpha is evidence of skill net of systematic risk. It is distinct from active return, which does not control for beta. $\alpha_J = r_p - [r_f + \beta_p(r_m - r_f)]$
- **Calmar Ratio:** Annualised CAGR divided by maximum drawdown. Measures how much annual return was earned per unit of worst-case peak-to-trough loss. Higher is better. $Calmar = \frac{CAGR}{MDD}$
- **Sortino Ratio:** Risk-adjusted return that only penalises downside volatility. Same numerator as Sharpe (excess return over the risk-free rate) but the denominator is the annualised downside deviation — the RMS of returns that fell below the target. Better suited than Sharpe to strategies with asymmetric return profiles, where upside volatility is desirable and should not be penalised. $Sortino = \frac{E(r_t - r_f)\cdot\sqrt{252}}{\sigma_down(r_t - r_f)}, where\ \sigma_down = \sqrt{(E[\,min(r_t - r_f, 0)^2\,])}$
- **Ulcer Index:** Root-mean-square of the drawdown series (depth as a fraction in [0, 1]). Unlike maximum drawdown, which is a single-point worst-case, the Ulcer Index integrates both depth and duration of every underwater period. It is a better match to the investor experience. $UI = \sqrt{E[\,DD_t^2\,]}, with\ DD_t\ the\ drawdown\ fraction\ at\ t$
- **Martin Ratio:** Annualised CAGR divided by the Ulcer Index. Return per unit of cumulative drawdown pain. Higher is better; a higher Martin Ratio than Calmar indicates that drawdowns, while material, were shallow on average. $Martin = \frac{CAGR}{UI}$
- **Time Underwater:** Drawdown duration measured in trading days. TUW_max is the longest consecutive run of days with the equity curve below its prior peak; TUW_mean is the average length of all such runs. Complements the depth-based metrics (MDD, Ulcer) with a duration-based view that answers the question: 'How long was I underwater?'. $TUW_max = max_k\ length(streak_k)$, $TUW_mean = \frac{\sum_k\ length(streak_k)}{K},\ where\ streak_k = maximal\ run\ with\ DD_t < 0$
- **Wealth Permanence:** The share of total investor out-of-risk wealth measured against current total wealth, not against the initial capital. Because the at-risk sleeve also compounds, a high banked-to-investment ratio and a moderate wealth-permanence percentage are consistent. $WealthPermanence(T) = \frac{Banked(T)}{W(T)},\ where\ W(T) = Equity(T) + Banked(T)$
- **200-Day Moving Average:** The '200DMA' comparator applies a simple Moving Average of the last 200 trading days to the configured equity-index proxy, holding the proxy when its close price is above the DMA(200) and cash otherwise. $DMA_200(t) = \frac{1}{200}\sum_{i=0}^{199} P(t-i);\ position_t = 1\ if\ P(t-1) > DMA_200(t-1), else\ 0$
- **Active Return:** Mean of daily (strategy − benchmark) returns, multiplied by 252. This is the numerator of the Information Ratio. Unlike Jensen's alpha it does not adjust for the level of market exposure taken. $AR_{ann} = E(a_t) \times 252$
- **Sharpe Ratio — Benchmark:** The 60/40 EQUITY+BOND benchmark's own Sharpe ratio computed over the same evaluation window. Comparing 'strategy Sharpe' with 'benchmark Sharpe' measures whether the strategy offers better total risk-adjusted return, independently of active risk.
- **Market Return:** Daily return of the market proxy (configured equity index), used as the base signal for regime detection. $r_m(t) = \frac{P_t}{P_\{t-1\}} - 1$
- **Market Volatility:** Rolling standard deviation of market returns over a fixed lookback window. Used to normalize market moves into comparable units. $\sigma_m(t) = std(r_m(t-k \ldots t))$
- **Market Regime:** Discrete slow-moving trend state of the market proxy: +1 (bull), 0 (neutral), −1 (bear). Derived from a causal trend filter on the market proxy price series. $Regime(t) \in \{-1, 0, 1\}$

- **True Range %:** Causal price disturbance measured on the current day. It is the true range divided by the previous close, so it captures both intraday range and gaps without using future observations.
- **Average True Range (ATR; atr_pct):** Trailing average of the daily true-range percentage. It predicts near-term drawdown risk but carries no directional information. It conditions risk only: stop width, position size, and skim rate.
- **Var Ratio Window (k):** The detector compares the variance of the last k regression residuals against the k residuals preceding. A ratio above 1 signals that scatter around the trend is accelerating, a precursor to phase breakdown.
- **Var Ratio Exit Blend:** Controls how much the variance ratio tightens the RMSE exit threshold when elevated. The exit threshold is: rmse_enter_max + blend × (rmse_exit_max − rmse_enter_max).
- **Binomial Model and Confidence Interval (CI):** Each completed trade is modelled as a Bernoulli trial with two outcomes: success (position closed by a regular sell or take-profit) or failure (position closed by a stop-loss). The observed win rate $\hat{p} = k/n$ is the maximum-likelihood estimate of p, where its precision depends on the sample size. A 95% confidence interval quantifies this uncertainty. The intervals are computed with the Wilson score method.
- **Low Confidence:** A 'low confidence' suffix is added when the CI lower bound would place the trade in a lower category than the point estimate.
- **Reward-to-Risk Ratio (RRR):** The realised RRR is the median price gain on winning trades divided by the median price loss on losing trades, both expressed as fractions of the entry price. An RRR of 2:1 means a winning trade gains twice as much as the trade loses. A low RRR can cause the expected result of a strategy to be unprofitable, despite many wins.
- **Break-even Win Rate:** The minimum win rate at which expected gains exactly equal expected losses, given the realised RRR. Formula: $p_a = 1 \div (1 + RRR)$. Example: RRR = 0.33:1 → p_a = 75%; the strategy must win at least 75% of trades to break even. RRR = 2:1 → p_a = 33%; only one win in three is needed to cover expected losses.
- **Homogeneity Test:** The homogeneity test (Pearson χ^2) checks whether all tickers in the portfolio share the same underlying p, or whether the win rate is an average of different performers. A post-hoc check gives the win-rate dispersion across assets.
- **edge:** Expected profit per trade expressed in units of the median loss (risk unit). Formula: $E = p \times (1 + RRR) - 1$, where p is the observed win rate. $E > 0$: the strategy is expected to be profitable. $E = 0$: break-even. $E < 0$: the strategy destroys value per trade regardless of the win rate. Thresholds used in this report: below break-even ($E < 0$); fair ($0 \le E < 0.10$); good, professional standard ($0.10 \le E < 0.50$); excellent ($E \ge 0.50$).
- **stop_loss_p95_upper:** The one-sided 95% upper bound for the true Stop-Loss probability, despite observing no Stop-Losses. It tells what Stop-Loss probability would have a 5% chance of producing zero Stop-Losses.

References

- Asness, C., Moskowitz, T. & Pedersen, L. (2013). 'Value and Momentum Everywhere'. The Journal of Finance 68(3), 929-985
- Berdondini, A.(2025). 'Binomial Evolution Function (The Only Function You Need to Evaluate a Trading Strategy)'. SSRN, Working Paper 5264411
- Faber, M. T. (2007). 'A Quantitative Approach to Tactical Asset Allocation'. The Journal of Wealth Management, 9(4), 69-79.
- Fama, E. & French, K. (2010). 'Luck versus Skill in the Cross-Section of Mutual Fund Returns'. The Journal of Finance, 65(5), 1915-1947
- Grinold, R. C., & Kahn, R. N. (2000). 'Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk'. McGraw-Hill, New York
- Hong, H. & Stein, J.(1999). 'A Unified Theory of Underreaction, Momentum Trading, and Overreaction'. The Journal of Finance, 54(6), 2143–2184.
- Hamilton, J.D.(1989). 'A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle'. Econometrica, 57(2), 357–384
- Jensen (1968). 'The Performance of Mutual Funds in the Period 1945–1964'. The Journal of Finance, 23, 387–419.
- Lo, A. W. (2019). 'The Statistics of Sharpe Ratios'. Financial Analysts Journal, 58(4), 36-52.
- Magdon-Ismail, M., & Atiya, A. F. (2004). 'An Analysis of the Maximum Drawdown Risk Measure'. https://api.semanticscholar.org/CorpusID:10758362
- Politis, D.N. and Romano, J.P. (1994). 'The Stationary Bootstrap'. Journal of the American Statistical Association, 89(428), 1303-1313.
- Sharpe, W. F. (1966). 'Mutual Fund Performance'. The Journal of Business, 39(1), 119-138.
- Sharpe, W. F. (1994). 'The Sharpe Ratio'. The Journal of Portfolio Management, 23(1), 49-58.