Technical Indicators on Warsaw Stock Exchange: A Computational Study

Konrad Maleńczak[1][0009−0000−5667−6323] and Alexander Prokopenya[2][0000−0001−9760−5185]

[1] Warsaw University of Life Sciences, Faculty of Applied Informatics and Mathematics, Institute of Information Technology, Department of Artificial Intelligence
konrad_malenczak@sggw.edu.pl
[2] Warsaw University of Life Sciences, Faculty of Applied Informatics and Mathematics, Institute of Information Technology, Department of Computer Information Systems
alexander_prokopenya@sggw.edu.pl

Abstract. Named technical indicators are hard to compare computationally unless data handling, thresholds, execution lag, and costs are frozen. We therefore benchmark 22 canonical outputs on daily open, high, low, close, and volume (OHLCV) data for 415 Warsaw Stock Exchange equities over 2015-06-28–2023-12-29 under one fixed protocol: shared cleaning rules, fixed signal mappings, one-day execution lag, and proportional costs (10 basis points, bps, per entry or exit). We additionally report one derived moving-average crossover rule, MA20>MA200. The study is meant as a computational audit of indicator primitives, not as a search for a new trading rule. The main diagnostic is action stability under a controlled close-price perturbation: for each stock, the close is rescaled by $\pm 1\%$, indicators are recomputed, and day-level agreement of discrete actions is measured. Slow trend-following families - moving averages, moving average convergence divergence (MACD), and directional movement index / average directional index (DMI/ADX) - retain median agreement near 100%, whereas stochastic oscillators are much more sensitive, with median stability of 68.75% for `STOCH_K_14` and 66.54% for `STOCH_D_3`. Correlation, principal component analysis (PCA), and mutual information also reveal strong within-family redundancy, including the exact identity between the Bollinger middle band and SMA(20). In this setting, near-100% stability marks structural robustness under the chosen perturbation, not predictive superiority.

Keywords: technical analysis · computational finance · reproducibility · robustness · feature redundancy

1 Introduction

Technical analysis is still widely used in trading practice and in feature engineering for predictive models [1–3]. Yet the output of a named indicator is rarely

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29909-3_21

determined by its formula alone. Data cleaning, warm-up handling, threshold conventions, execution lag, and trading costs can all change the resulting actions [4, 5]. For that reason, we start from a narrower question than the usual profitability question: under one fixed implementation, which common indicators keep their actions after a small perturbation of the input close series?

Our focus is computational rather than predictive. Each indicator is treated as a deterministic mapping from open, high, low, close, and volume (OHLCV) inputs to discrete actions, and every indicator is evaluated under the same frozen protocol on the same universe of 415 Warsaw Stock Exchange (WSE) equities. We do not tune parameters, search over rules, or compare optimised strategies. Instead, we ask three concrete questions: which indicator rules are structurally stable under small close-price changes, which ones are fragile, and which outputs are largely duplicative before any downstream model is built.

The result is a benchmark rather than a new trading rule. It covers 22 canonical outputs, adds one derived moving-average crossover rule, measures action stability under controlled $\pm 1\%$ close rescaling, and audits redundancy using correlation, principal component analysis (PCA), and mutual information. The contribution is therefore methodological at the protocol level: a reproducible computational benchmark for auditing deterministic indicator mappings under controlled perturbations. We keep the setup deliberately simple so that instability can be traced to the indicator rules and fixed execution protocol, not to strategy optimisation.

2 Related Work

Most empirical work on technical analysis evaluates trading-rule profitability [1, 6]. Another strand uses indicators as inputs to larger predictive systems [3]. A computational perspective appears in [2], where chart patterns are formalised as algorithmic objects rather than informal visual heuristics. That perspective is closer to our aim than profitability ranking alone.

What is less common is a controlled audit of indicator behaviour under one shared implementation. In practice, the same nominal rule can change materially after differences in adjustment conventions, missing-day handling, lag choices, or threshold definitions [7]. We therefore freeze the protocol and perturb the close series before any downstream optimisation. To our knowledge, a WSE-wide benchmark that combines a single implementation, a close-perturbation stability check, and an output-level redundancy audit is still under-documented.

3 Data and Fixed Protocol

The data consist of daily OHLCV series for 415 WSE equities retained after filtering for valid OHLCV fields and minimum history. Each retained stock contains at least 468 cleaned observations; the median series length is 1,444 trading days and the maximum is 1,610. The overall sample spans 2015-06-28–2023-12-29. Files are mapped to a common schema (`DATE`, `OPEN`, `HIGH`, `LOW`, `CLOSE`, `VOL`),

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29909-3_21

sorted chronologically, checked for non-positive prices, duplicated timestamps, and high–low inconsistencies, and kept without imputation. Missing trading days are not interpolated because this would distort rolling indicators.

The library contains 22 computed outputs drawn from four familiar families: trend, momentum, volatility, and volume. Trend filters are represented by simple and exponential moving averages (SMA, EMA) with windows 10/20/50/200, moving average convergence divergence (MACD), and the directional movement index / average directional index (DMI/ADX) family. Momentum is represented by the relative strength index (RSI) and stochastic $\%K(14)$ / $\%D(3)$. Volatility is represented by average true range (ATR), Bollinger bands, and on-balance volume (OBV). We also report one derived crossover rule MA20>MA200 built from SMA20 and SMA200, so it adds no new raw output. The set is canonical rather than exhaustive: broad enough to compare common defaults, but small enough to keep the benchmark interpretable.

We distinguish indicator families, individual outputs, decision rules, and evaluated items. Redundancy is computed at the output level, but long-or-cash actions are generated only after one fixed rule is assigned to the relevant family. Single SMA/EMA outputs use price-vs-average rules, whereas MA20>MA200 is treated separately as a crossover. In Table 2, `STOCH_K_14` and `STOCH_D_3` denote stochastic %K(14) and %D(3), while Lower_t and Upper_t denote the lower and upper Bollinger bands at day t. ADX14 and ATR14 are not inherently directional trading rules; their sign-of-change mapping is included only as an auxiliary proxy for within-benchmark comparability, and their action-based rankings are interpreted separately from canonical directional rules. Positions are updated with a one-day lag, transaction costs are charged at 10 basis points (bps) per entry or exit, and no parameter search is performed. The one-day lag avoids same-close execution bias, and the 10 bps per-leg cost is a simple uniform friction level rather than a market-specific estimate for every WSE name.

Cleaned inputs, indicator outputs, perturbed reruns, and final summaries are stored separately so disagreements can be traced to one stage. Tables 1 and 2 collect the main settings.

Table 1. Data and protocol summary.

Universe	415 WSE equities
Time span	2015-06-28–2023-12-29
Minimum length	468 trading days
Median series length	1,444 trading days
Fields	OHLCV
Cleaning	schema unification, sorting
Retention rule	valid OHLCV fields and $\geq$ 468 observations
Indicators	22 raw outputs + MA20>MA200
Signal protocol	fixed rules, no tuning
Execution	one-day lag, long-or-cash
Transaction cost	10 bps per entry or exit

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29909-3_21

Table 2. Canonical action mappings used throughout the benchmark. Equalities leave the current state unchanged; ADX14 and ATR14 enter only through an auxiliary sign-of-change proxy.

Evaluated output(s)	Fixed long-or-cash rule
SMA(n), EMA(n)	long if $P_t > MA_n(P)_t$; exit if $P_t < MA_n(P)_t$
MA20>MA200	long if $\text{SMA}_{20,t} > \text{SMA}_{200,t}$; exit if $\text{SMA}_{20,t} < \text{SMA}_{200,t}$
RSI(14)	canonical contrarian rule: long if $\text{RSI}_t < 30$; exit if $\text{RSI}_t > 70$
STOCH_K_14, STOCH_D_3	canonical contrarian rule: long if value < 20; exit if value > 80
MACD family (`MACD`, `MACDsig`, `MACDhist`)	long if $\text{MACD}_t > \text{Signal}_t$; exit if $\text{MACD}_t < \text{Signal}_t$
DMI family (`DMI/ADX`, `+DI14`, `-DI14`)	long if $+DI_t > -DI_t$; exit if $+DI_t < -DI_t$
OBV	long if the one-day change is positive; exit if negative
ADX14, ATR14	auxiliary proxy only: long if the one-day change is positive; exit if negative
Bollinger bands	long if $P_t < \text{Lower}_t$; exit if $P_t > \text{Upper}_t$

4 Methods

For evaluated output k, let $s_t^{(k)} \in \{-1, 0, 1\}$ denote the action at day t, where 1 enters a long position, 0 leaves the current state unchanged, and -1 exits to cash. The stability metric therefore compares discrete actions rather than realised portfolio values. The main robustness diagnostic recomputes the full pipeline under three close series,

$$P_t^{(0)} = P_t, \qquad P_t^{(+)} = (1+\varepsilon)P_t, \qquad P_t^{(-)} = (1-\varepsilon)P_t,$$

where P_t is the original close on day t, $P_t^{(0)}$ is the unperturbed series, $P_t^{(+)}$ and $P_t^{(-)}$ are the upward- and downward-rescaled closes, and $\varepsilon = 0.01$. We then compare the resulting discrete actions day by day and report the percentage of dates on which all three actions agree:

$$\text{stability}^{(k)} = 100 \cdot \frac{n_{\text{agree}}^{(k)}}{n_{\text{defined}}^{(k)}}. \tag{1}$$

In Eq. (1), $n_{\text{defined}}^{(k)}$ is the number of dates for which all three runs produce a defined action after indicator warm-up, and $n_{\text{agree}}^{(k)}$ counts the dates on which those three actions coincide. The perturbation is intentionally local: it represents small close-price discrepancies that can arise across vendors due to rounding or adjustment conventions, not a general market stress test. We keep it simple on purpose so that disagreements can be traced back to indicator definitions and signal rules. It is applied to the close series only; open, high, low, and volume are otherwise left unchanged.

Its interpretation depends on the fixed signal rules defined in Section 3. For close-derived trend rules, global multiplicative rescaling often preserves order

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29909-3_21

relations such as $P_t > \mathrm{MA}_n(P)_t$, so perfect or near-perfect stability is expected. Concretely, if $P'_t = cP_t$ with $c > 0$, then $\mathrm{MA}_n(P')_t = c\,\mathrm{MA}_n(P)_t$ and therefore $\mathrm{sign}(P'_t - \mathrm{MA}_n(P')_t) = \mathrm{sign}(P_t - \mathrm{MA}_n(P)_t)$. For such rules, 100% stability is a consequence of scale invariance under the chosen perturbation rather than evidence of predictive superiority. The test is more informative for rules that depend on the location of the close within the daily range or on threshold crossings after normalisation; stochastic oscillators are the main example in our library.

Each evaluated item also drives a minimal long-or-cash strategy executed with one-day lag on the next trading day. We report annualised Sharpe ratio with zero risk-free rate, compound annual growth rate (CAGR), drawdown, profit factor, exposure, and trade counts. Indicator-level summaries are aggregated across stocks by the median. In the composite ranking, $\mathrm{performance}_k$ denotes the median cross-sectional annualised Sharpe under the fixed lag-and-cost protocol. Stock–indicator pairs with fewer than 20 completed trades or exposure below 5% are excluded in the performance summaries in order to avoid unstable statistics generated by nearly inactive rules; CAGR per exposure and profit factor serve only as tie-breakers.

Redundancy is analysed from normalised one-day indicator changes. For each stock–output pair, we first compute one-day first differences and then standardise them to zero mean and unit variance over the valid sample. For one stock, let $X \in \mathbb{R}^{T' \times K}$ collect these valid standardised changes for the K evaluated outputs. Here $X_{\cdot i}$ and $X_{\cdot j}$ denote the ith and jth columns of X, i.e. the standardised one-day changes of outputs i and j over the T' valid dates for that stock. Linear overlap is summarised by the correlation matrix in Eq. (2)

$$C_{ij} = \mathrm{corr}(X_{\cdot i}, X_{\cdot j}), \tag{2}$$

and a compact non-redundancy proxy for output k is

$$\mathrm{independence}_k = 1 - \frac{1}{K-1} \sum_{\substack{j=1 \\ j \neq k}}^{K} |C_{kj}|. \tag{3}$$

Larger values in Eq. (3) therefore mean less average overlap with the rest of the library. Mutual information [8] and PCA are also computed as supporting structural diagnostics rather than as direct ranking drivers or substitutes for out-of-sample predictive evaluation.

For compact screening we also form one auxiliary composite score. Let q_k be an evaluated-item quantity and map it to $[0, 1]$ by min–max normalisation, where $q_{\min}$ and $q_{\max}$ are the minimum and maximum of that quantity across all evaluated items,

$$\mathrm{score}(q_k) = \begin{cases} \dfrac{q_k - q_{\min}}{q_{\max} - q_{\min}}, & q_{\max} > q_{\min}, \\ 0.5, & q_{\max} = q_{\min}. \end{cases} \tag{4}$$

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29909-3_21

Equation (4) is then applied to stability, performance, and the independence proxy from Eq. (3). The composite combines them with fixed weights,

$$\begin{aligned} \text{composite}_k = {} & 0.4\,\text{score}(\text{stability}_k) \\ & + 0.4\,\text{score}(\text{performance}_k) \\ & + 0.2\,\text{score}(\text{independence}_k). \end{aligned} \tag{5}$$

Equation (5) is used only as a compact screening convenience rather than the primary scientific result. Stability and performance are weighted equally because poor robustness or poor performance can each make an evaluated item unattractive in practice, while independence is down-weighted because it depends more strongly on representation and estimator choices.

5 Results and Discussion

5.1 Stability under close rescaling

One pattern dominates the stability results: slow trend rules remain almost unchanged under the perturbation. Median stability is effectively 100% for DMI/ADX, MACD, and most SMA/EMA rules. At the other end, stochastic oscillators drop to 68.75% for `STOCH_K_14` and 66.54% for `STOCH_D_3`. The library is also partly redundant by construction, including the exact identity between the Bollinger middle band and SMA(20). Table 3 summarises the main best and worst cases under the fixed protocol.

By contrast, stochastic oscillators are much less robust. `STOCH_K_14` reaches median stability 68.75% and `STOCH_D_3` reaches 66.54%, so roughly one third of their daily actions change after the ±1% close perturbation. The interpretation is straightforward. If a rule compares the close with a moving average computed from the same close history, both sides co-move when the entire close history is multiplicatively rescaled. If a rule depends on the close relative to the daily range or on threshold crossings after nonlinear normalisation, the same perturbation can change the action much more easily. The stability ranking therefore separates scale-invariant from non-invariant decision mechanisms rather than "good" from "bad" indicators in a predictive sense. Under the present perturbation design, the slow trend-following family is simply the most robust family.

Table 3. Key best and worst results under the fixed protocol.

Finding	Summary
Best stability group	DMI/ADX, MACD, SMA/EMA (median stability near 100%)
Most sensitive group	Stochastic oscillators
Median stability of `STOCH_K_14`	68.75%
Median stability of `STOCH_D_3`	66.54%
Exact identity	Bollinger middle band = SMA(20)
Top composite indicator	MA20>MA200 (0.890)
Next composite indicators	DMI/ADX 0.703, ADX14 0.683, MACD>Signal 0.638

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29909-3_21

5.2 Redundancy and structural interpretation

Redundancy is strong within indicator families. Some relationships are exact by construction: the Bollinger middle band is identical to SMA(20), so including both in the same library adds no new information. Moving-average families, MACD components, and DMI/ADX components also cluster tightly in the correlation view; PCA and mutual information tell the same story. As a practical screening reference point, we read absolute correlations above 0.95 as near-duplication when the same pairs also remain close in the PCA and mutual-information views. This is a heuristic threshold for compact screening rather than a formal statistical cutoff.

Otherwise a large indicator library can look more diverse than it really is and still recycle the same smoothed trend signal. For the present WSE universe, the redundancy diagnostics support a compact baseline centred on slow trend filters plus a small number of less redundant components, such as selected MACD- or ADX-related outputs.

These diagnostics should not be read as a rule for mechanically removing every highly correlated variable. They show where extra variants add almost nothing beyond signals already present in the library.

5.3 Composite screening and profitability

The composite score serves only as a screening device. Under this setup, the 20/200 moving-average crossover ranks first with composite score 0.890, followed by DMI/ADX (0.703), ADX14 (0.683), and MACD greater than Signal (0.638). This ranking should therefore be read as a benchmark-specific ordering under the present perturbation, lag, and cost assumptions, not as a universal ordering of indicator quality. That ordering partly reflects the perturbation design and the simplicity of the slow trend rules: small global changes in the close leave their decisions largely unchanged.

ADX14 stays near the top because it combines perfect stability with higher independence than several neighbouring trend filters, although its pure performance summaries are less impressive. Its rank should therefore be read as a property of the auxiliary proxy under the fixed benchmark, not as evidence that ADX14 is a standalone directional trading rule.

Single-indicator backtests do produce a sparse right tail, but the typical cross-sectional outcome remains modest once the same lag and costs are enforced everywhere. Profitability is therefore used only as a screening aid. Under this protocol, some indicators are stable but highly redundant, whereas others add variety at the cost of lower robustness.

5.4 Implications for downstream modelling

For downstream modelling, a compact trend-following core is easier to justify than a large default library. Exact duplicates can be removed, fragile oscillators make sense only when their sensitivity is part of the modelling hypothesis, and

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29909-3_21

further library growth should be motivated by genuine diversity because the dependence audit is the main bottleneck. Computationally, indicator evaluation is close to linear in series length and the number of outputs under the fixed protocol, whereas the pairwise dependence audit grows quadratically with library size and dominates further expansion of the benchmark.

6 Conclusion

The main result is simple: under one fixed protocol, slow trend-following rules are highly stable under small close rescaling, whereas stochastic oscillators are much more sensitive. This is not a statement about universal predictive superiority. In our setting, near-100% stability mostly reflects scale-tolerant decision rules; lower stability marks rules whose actions depend more strongly on the close relative to the daily range or on threshold crossings after normalisation.

Within those limits, the benchmark is most informative as a pre-screening step before downstream feature engineering, supervised modelling, or stricter out-of-sample evaluation. It highlights where a large indicator library is only superficially diverse, where a canonical output is fragile to small implementation changes, and where exact duplicates can be removed without loss. A natural next step is to extend the perturbation family beyond close-only rescaling and retest the compact baseline out of sample under liquidity-aware execution.

References

1. Park, C.-H., Irwin, S.H.: Profitability of technical analysis. J. Econ. Surveys **21**(4), 786–826 (2007)
2. Lo, A.W., Mamaysky, H., Wang, J.: Foundations of technical analysis. J. Finance **55**(4), 1705–1765 (2000)
3. Kara, Y., Boyacioglu, M.A., Baykan, O.K.: Predicting stock index direction using neural networks and support vector machines. Expert Syst. Appl. **38**(5), 5311–5319 (2011)
4. Sullivan, R., Timmermann, A., White, H.: Data-snooping and trading-rule performance. J. Finance **54**(5), 1647–1691 (1999)
5. Bailey, D.H., Borwein, J.M., Lopez de Prado, M., Zhu, Q.J.: Probability of backtest overfitting. J. Comput. Finance **20**(4), 39–69 (2017)
6. Brock, W., Lakonishok, J., LeBaron, B.: Simple technical trading rules and stochastic properties of stock returns. J. Finance **47**(5), 1731–1764 (1992)
7. Pérignon, C., Akmansoy, O., Hurlin, C., et al.: Computational reproducibility in finance. Rev. Financ. Stud. **37**(11) (2024)
8. Kraskov, A., Stögbauer, H., Grassberger, P.: Estimating mutual information. Phys. Rev. E **69**, 066138 (2004)

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29909-3_21