Research Article

Adapt to Win: A Statistical Analysis on Strategies in Rock-Paper-Scissors

Ruohan Huang

Junior, Walnut High School, 400 N. Pierre Road, Walnut, CA 91789, USA

ABSTRACT

Game theory provides a robust framework for analyzing strategic decision-making, and Rock-Paper-Scissors (RPS), despite its simplicity, serves as an effective model due to its balanced structure and broad applicability. This study evaluates the performance of nine RPS strategies—ranging from random and fixed patterns to adaptive techniques—through pairwise matchups simulated over multiple rounds. Statistical tools, including Moving Average (MA), Cumulative Sum (CUSUM) Control Chart, and Decay-Weighted (DW) Metrics, assess each strategy's stability and adaptability in individual match pairs and overall outcomes. Results show that adaptive strategies focusing on recent history, such as multi-round observation and certain reaction-based techniques, excel in long-term performance by precisely adapting to opponents' behavioral shifts while avoiding overreactions to minor variations. These findings highlight the critical importance of adaptability in optimizing decision-making in competitive environments.

Keywords: Game Theory, Rock-Paper-Scissors (RPS), Strategic Decision-Making, Adaptability, Quantitative Analysis, Multi-Round Observation

INTRODUCTION

Imagine standing at a crossroads in a competitive game, where each decision could mean victory or defeat - this is where game theory takes center stage. By breaking down the complexities of strategic interactions, it enables individuals and teams to navigate competition with confidence and make informed choices. Figuratively, game theory is not just a trusted ally when lifting a trophy; it is also the root of your joy in the thrill of victory. Acting as an invisible supporter, it guides players through intricate challenges, uncovering the best paths to success. For instance, in poker, strategic bluffing creates uncertainty, giving players a critical edge (1). In sports, game-theoretic models support real-time decisions that enhance performance (2). Similarly, in finance, dynamic strategies inspired by these principles allow investors to adapt to market changes and maximize returns (3). These examples illustrate how this framework bridges theory and practice, making it a cornerstone of success in competitive environments.

RPS is a timeless game tied to childhood memories, celebrated for its simplicity and accessibility. Requiring only hand gestures, it allows players to participate

Corresponding author: Ruohan Huang, E-mail: tracy.ruohan.huang@gmail.com.

Received November 30, 2024**; Accepted** December 16, 2024
https://doi.org/10.70251/HYJR2348.24129148

anywhere, making it universally appealing. While often seen as a quick and fair decision-making tool, RPS also involves subtle strategies that challenge players to think quickly and adapt. Its fairness ensures equal chances of success for players of all backgrounds. Beyond its simplicity, RPS serves as a valuable platform for exploring game-theoretic strategies. Its straightforward rules are ideal for studying decision-making processes, equilibrium strategies, and dynamic adjustments (4). When extended to multi-round games, RPS reveals how strategies evolve over time, offering insights into long-term optimization (5). Additionally, it supports experimental research by enabling effective hypothesis testing and model validation (6, 7). This unique mix of simplicity and strategic depth makes RPS an excellent tool for advancing game theory in both theory and practice.

While RPS is widely recognized, research on how strategies perform and adapt in multi-round confrontations remains limited. Understanding how strategies evolve, interact, and sustain effectiveness over time is essential for optimizing decision-making in competitive environments. This study addresses this gap by examining the nuanced behavior of strategies across progressively increasing rounds, offering valuable insights into their adaptability, stability, and long-term potential.

To achieve this, nine different RPS strategies are explored, including random moves, fixed patterns, and adaptive techniques, each designed to respond to specific gameplay situations. Detailed descriptions of these strategies are provided in Table 1 in the following chapter. Simulations systematically evaluate strategy performance across early, middle, and late rounds in pairwise matchups and overall outcomes. By capturing both short- and long-term trends, the data reveal how strategies adjust and evolve over time. Metrics, primarily win rates, form the basis for statistical analysis. Using tools such as MA, CUSUM Control Chart, and DW Metrics, trends across gameplay phases are assessed, focusing on stability and adaptability. This research underscores the critical role of adaptability in achieving long-term success and highlights optimization potential for improving strategic decision-making in dynamic systems.

METHODS AND MATERIALS

Purpose of Simulation

Simulation serves as the foundation of this study, providing a controlled and scalable environment to evaluate the performance of RPS strategies. By systematically reproducing interactions between strategies, the simulation generates robust and reproducible data that captures performance trends over time.

RPS, with its simple rules and constrained state space, is particularly well-suited for simulation-based analysis. These characteristics allow for exhaustive exploration of strategic interactions while ensuring the data remains representative of all possible states and outcomes (8). Moreover, simulations allow for repetitive experiments under controlled conditions, enhancing the statistical significance and reliability of findings. These strategies can be programmed and simulated to assess their performance in different contexts (9). Simulation also allows the researcher to control variables in the experiment, such as the number of rounds and participants' strategy choices, which provides for a systematic exploration of the dynamic relationships between strategies (10).

This research leverages these advantages to examine diverse RPS strategies. Simulations enable precise control over variables such as the number of rounds and pairing rules, offering insights into how strategies evolve and perform across different phases of gameplay. The low-cost, risk-free nature of simulation further facilitates large-scale data collection without the ethical or logistical concerns of human experimentation (11).

Through this approach, simulation provides a comprehensive dataset that serves as the basis for analyzing long-term optimization potential and dynamic strategy interactions. These findings contribute to advancing both theoretical and practical applications of game theory in competitive scenarios.

Simulation Design

The simulation is structured to evaluate the performance of distinct RPS strategies across progressively increasing rounds. This design ensures a systematic analysis of both short-term behavior and long-term adaptability, providing insights into how strategies perform and evolve over time.

Strategies Overview. The nine strategies analyzed in this study represent a diverse range of decision-making approaches. Fixed strategies, such as constant and sequential, serve as baselines for examining predictable and static behaviors. Reactive strategies, including mirror, reverse mirror, reaction, and exclusion, directly respond to opponents' most recent move. Adaptive strategies, such as probability-weighted and multi-round observation, incorporate learning and make dynamic adjustments over multiple rounds to respond to opponents' evolving tendencies. Table 1 provides detailed definitions and example scenarios for each strategy, illustrating their distinct mechanisms and how they function during gameplay.

Game Setup. The simulation follows a round-robin format where each strategy competes against every other, including itself, in pairwise matchups. This systematic approach ensures comprehensive evaluation of all possible interactions. Matches are conducted over progressively increasing rounds, ranging from 10 to 1,000 in increments of 10, capturing both short-term performance and long-term adaptability. The 1,000-round limit balances the need to observe trends with computational efficiency, as extending to 10,000 rounds adds minimal additional insights while increasing costs.

Before each match, strategies are initialized to ensure fair conditions. During gameplay, strategies follow predefined rules, with results recorded for each round. Matchups such as Strategy A vs. Strategy B and Strategy B vs. Strategy A are treated equivalently to simplify evaluation. The total 45 unique pairings calculated using $n\ (n+1)\ /\ 2$ for $n = 9$, include self-pairings, which are addressed separately in the analysis.

This controlled simulation design enables the collection of comprehensive performance data, systematically examining strategy interactions to provide a strong foundation for detailed analysis.

Key Metrics. The simulation generates raw performance data for each pairing between strategies, focusing on metrics that quantify effectiveness across varying rounds.

The primary metric at this stage is the Pairing Win Rate (PWR), which represents the percentage of rounds a strategy wins in a specific pairing. It is calculated as $PWR = \left(\frac{\text{Number of Wins}}{\text{Total Rounds Played}}\right) \times 100$. This metric reflects how well a strategy performs against a specific opponent during a single game, forming the foundation for further analysis.

Another recorded metric, the Pairing Draw Rate (PDR), captures the percentage of rounds that end in a draw. While PDR is not analyzed in this study, it offers potential for future research, particularly in exploring

Table 1. Strategy Definitions and Example Scenarios

Strategy (Abbr.)	Definition	Example Scenarios
Random (Rndm)	Randomly selects a move each round, providing a baseline for unpredictability.	Opponent: Rock → Random selects Rock, Paper, or Scissors randomly.
Constant (Const)	Makes the same choice every round, representing a fixed strategy.	Constant chooses **Rock** every round, regardless of opponent's move.
Sequential (Seq)	Cycles through moves in a set order, such as rock, paper, scissors.	Opponent: Rock → Sequential plays Paper (next in sequence).
Mirror (Mirr)	Copies the opponent's last move in each round.	Opponent: Paper → Mirror plays Paper (copies opponent's move).
Exclusion (Excl)	Excludes the opponent's last move from its own choices, aiming to be less predictable.	Opponent: Scissors → Exclusion plays Rock or Paper (excludes Scissors).
Reaction (React)	Counters the opponent's previous move with the corresponding winning move.	Opponent: Paper → Reaction plays Scissors (counters Paper).
Reverse Mirror (RevMirr)	Plays a move that would lose to the opponent's last move, using a counterintuitive approach.	Opponent: Rock → Reverse Mirror plays Scissors (loses to Rock).
Probability-Weighted (ProbWt)	Selects a move based on the frequency of the opponent's previous choices to maximize effectiveness.	Opponent frequently plays Rock → ProbWt likely chooses Paper.
Multi-Round Observation (MRO)	Identifies patterns over multiple rounds and adapts moves based on observed opponent behavior.	Opponent: Scissors in previous **3** rounds → MRO plays Rock (based on pattern).

Note: Abbreviations are derived from strategy names, with the first letter capitalized for clarity, except for MRO, which is presented in uppercase as an acronym. The MRO strategy observes the last 3 rounds, while ProbWt considers all prior moves, testing the difference in responsiveness among adaptive strategies.

strategies that frequently result in ties.

These raw metrics are directly derived from the simulation results and provide the groundwork for calculating higher-level metrics, such as Win Rate Disparity (WRD) and Overall Win Rate (OWR), in subsequent stages of the analysis.

Data Summary

The simulation produces a structured dataset capturing the outcomes of pairwise interactions between strategies. Each row represents one player's perspective in a specific match pair, including identifiers such as match pair ID, player ID, and strategy abbreviation, as well as performance metrics like wins, draws, and PWR. This two-rows-per-matchup format supports detailed and scalable analysis while maintaining flexibility for potential future research involving multiplayer scenarios.

The dataset's structure is defined by two types of fields: identifiers and metrics. Identifiers help track the players and strategies involved in each simulation, while metrics quantify performance outcomes. These elements are detailed in Table 2a and Table 2b.

To illustrate the raw data structure, a sample dataset is presented in Table 3 below. Each row corresponds to one

Table 2a. Identifiers for Simulation Results

Field	Type	Description	Example
Match Pair ID	Categorical	Identifier for each pairwise matchup in the simulation, representing the combination of strategies.	12
Player ID	Categorical	Identifier for the player within a specific match pair (e.g., PL1 for Player 1).	PL1
Strategy Abbreviation	Categorical	Abbreviation for the strategy used by each player. Refer to Table 1 for definitions of strategy abbreviations.	Rndm

Table 2b. Metrics for Simulation Results

Metric	Type	Description	Range	Example
Rounds Per Game (RPG)	Integer	The pre-defined number of rounds played in each matchup.	[10, 20, ...,1000]	10
Wins (wins)	Integer	The number of rounds won by the player in the matchup.	[0, RPG]	3
Draws (draws)	Integer	The number of rounds that ended in a draw.	[0, RPG]	5
PWR	Float	Percentage of rounds won by the player in the matchup, calculated as .	[0.00, 100.00]	30
PDR	Float	Percentage of rounds that ended in a draw, calculated as .	[0.00, 100.00]	50

To illustrate the raw data structure, a sample dataset is presented in Table 3 below. Each row corresponds to one player's perspective in a match pair.

Table 3. Sample Dataset from Simulation Results

Match Pair ID	RPG	Player ID	Strategy Abbr.	Wins	Draws	PWR (%)	PDR (%)
16	10	PL1	Rndm	3	3	30	30
16	10	PL2	React	4	3	40	30
78	340	PL1	RevMirr	108	109	31.76	32.06
78	340	PL2	ProbWt	123	109	36.18	32.06
45	680	PL1	Mirr	190	231	27.94	33.97
45	680	PL2	Excl	259	231	38.09	33.97
39	1000	PL1	Seq	332	335	33.2	33.5
39	1000	PL2	MRO	333	335	33.3	33.5

player's perspective in a match pair.

The dataset forms the foundation for evaluating strategy performance, offering granular insights through PWR and broader trends via metrics like OWR and WRD. Its structure supports current analysis and allows for future extensions, such as multiplayer scenarios or additional metrics.

Data Preprocessing

The preprocessing phase transforms raw simulation data into a structured dataset by excluding self-pairing matchups, where a strategy competes against itself. Although included in simulations for completeness, self-pairings were removed to focus on meaningful pairwise comparisons.

From the cleaned dataset, WRD, or Disparity, was derived to measure performance gaps between competing strategies. Defined as $WRD = PWR_1 - PWR_2$, where PWR_1 and PWR_2 are the PWRs of player 1 and player 2, WRD evaluates both the magnitude and direction of differences. A positive WRD indicates the first strategy outperformed the second, while a negative value indicates the opposite.

The second metric, OWR, provides a summary of a strategy's performance across all pairings, excluding self-pairings. OWR is computed as the average of all PWR values for a strategy: $\mathrm{OWR} = \frac{\Sigma \mathrm{PWR}}{\mathrm{Total\ Matchups}}$. This metric offers a holistic view of a strategy's effectiveness against various opponents.

RPG values were segmented into three phases: Early (10–330 rounds), Middle (340–670 rounds), and Late (680–1,000 rounds). This framework facilitates analysis of initial performance adjustments, transitional behaviors, and eventual stabilization, providing a comprehensive view of strategy dynamics.

Fundamental Statistical Tools

Basic statistical measures, such as Mean and Standard Deviation (Std.), are foundational for quantifying the central tendencies and variability of key outcomes. These metrics provide the groundwork for assessing strategy performance and understanding overall trends, serving as a basis for further analysis.

Heatmaps, a widely used tool in data visualization, are employed to explore interaction trends between strategies. By representing both the magnitude of disparities and their variance as color gradients, heatmaps provide an intuitive way to identify relative strengths, weaknesses, and unexpected patterns in match pair dynamics (12).

Error Bar Charts serve to represent the mean and variability of results for each strategy. By showing confidence intervals or standard deviations, this tool enables a straightforward comparison of the reliability and effectiveness of strategies across matchups (13).

These methods lay the groundwork for more advanced analyses, such as performance trends over time and adaptability assessments, which are discussed in subsequent sections.

Advanced Analytical Techniques

Moving Average. The MA is a widely used statistical tool for analyzing trends in time-series data, smoothing fluctuations over a fixed observation window to highlight consistent patterns while filtering out short-term noise. At time, the MA is computed as $MA(\mathrm{C}_t) = \frac{1}{N}\Sigma_{i=t-N+1}^{t} C_i$, where N is the number of periods over which the average is calculated. This technique effectively distinguishes random noise from genuine shifts, making it particularly valuable for robust monitoring and trend analysis in dynamic contexts (14, 15).

The Moving Average Stability Index (MASI) quantifies stability by evaluating variability in smoothed metrics over a defined observation period. MASI is calculated as $MASI = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(x_i - \overline{x})^2}$, where x_i represents the smoothed value at the i-th observation, $\overline{x}$ is the mean of these values, and n is the number of observations in the window. Lower MASI values indicate greater stability with minimal fluctuations, while higher values signify variability and potential instability (16, 17).

MA was applied to analyze trends in win rates, the CUSUM of win rates, and win rate disparities, capturing patterns of consistency and variability. MASI was employed to assess phase-specific stability for win rates (early, middle, and late phases) and long-term stability for the CUSUM of win rates across all rounds. Lower MASI values in win rates indicated consistent performance under varied conditions, while higher values suggested significant variability.

Cumulative Sum. The CUSUM Control Chart, introduced by E.S. Page in 1954, is a statistical tool designed to detect shifts in the mean level of a process over time. It is widely used in quality control, finance, and industrial monitoring due to its ability to identify small, persistent shifts more effectively than traditional control charts (14, 15). While its sensitivity to gradual changes is a strength, it can also lead to false alarms in high variability processes (16).

At any given time t, the CUSUM value C_t is calculated as the cumulative sum of deviations from the target mean μ_0, expressed as $C_t = \Sigma_{i=1}^{t}(X_i - \mu_0)$, where X_i represents the observed value at time i. If C_t exceeds predefined thresholds, it signals a potential shift in the process mean.

These thresholds, referred to as the Upper Control Limit (UCL) and Lower Control Limit (LCL), are calculated as UCL = $k + H$ and LCL = $-k - H$, where k is the reference value (typically set to half the expected shift size), and H determines the chart's sensitivity. When C_t crosses these limits, it indicates a statistically significant deviation from the target mean, warranting further investigation (17).

The MA CUSUM builds on the traditional CUSUM by incorporating a moving average to smooth the cumulative sum, reducing the impact of short-term fluctuations. By combining CUSUM's sensitivity to subtle shifts with MA CUSUM's smoothing capability, these methods effectively distinguish meaningful trends from random fluctuations, providing comprehensive insights into both consistent and evolving dynamics in sequential data.

In this research, CUSUM and MA CUSUM were applied to analyze the sequential dynamics of win rates and win rate disparities of strategies. X_i denotes the observed win rate or disparity at round i, while μ_0 is the overall mean. The UCL and LCL were used to identify statistically significant shifts in strategy performance across rounds.

MA CUSUM was particularly valuable for smoothing variability in high-variance matchups, providing clearer insights into long-term trends and adaptability. This combined approach allowed for a detailed assessment of how strategies maintained or lost performance consistency over time, highlighting critical phases of adaptation or decline.

DW Metrics. DW Metrics are rooted in the principle of Exponential Decay, a mathematical concept used to model processes where recent observations are more relevant than older ones. The decay function $w_i = e^{-\lambda(t-i)}$ assigns exponentially decreasing weights w_i to observations, where λ is the decay rate constant and t is the current time. A higher λ emphasizes recent events, while a lower λ incorporates a broader historical perspective. This weighting approach is widely applied in time series analysis, stochastic modeling, and adaptive systems to prioritize timely information (18).

The DW average extends this principle to compute a smoothed value for a dynamic variable, emphasizing recent data points while gradually attenuating older ones. Itiscalculatedas $DW(X_t) = \frac{\Sigma_{i=1}^{t} w_i X_i}{\Sigma_{i=1}^{t} w_i}$, where X_i represents the observed value at time i, and w_i are decay-adjusted weights. This averaging method enhances sensitivity to recent trends, making it particularly suited for systems undergoing rapid changes.

Metrics derived from DW average enable deeper analyses. The DW Slope (S) leverages linear regression principles to measure the rate of change in a decay-weighted context.

It is expressed as $S = \frac{\Sigma_{i=1}^{t} w_i (t_i - \bar{t})(X_i - \bar{X})}{\Sigma_{i=1}^{t} w_i (t_i - \bar{t})^2}$, where $\bar{t}$ and $\bar{X}$ denote the weighted means of time indices and observations, respectively. This slope quantifies trends in dynamic processes, indicating acceleration, deceleration, or stabilization.

The Adaptability Index (AI) measures variability in decay-weighted values over time, capturing the extent of a system's dynamic adjustments. It is computed as $AI = \sigma(DW(CUSUM(X_t)))$, where σ is the Std. and $CUSUM(X_t)$ represents the cumulative sum of observations at time t, weighted using the decay function. Higher AI values suggest frequent and substantial adjustments, while lower values indicate steadier, less reactive behaviors.

In this research, decay-weighted metrics are employed to analyze two critical aspects of strategic behavior in RPS, with X_i representing win rates. The S quantifies the rate of improvement or decline in performance over time, emphasizing recent observations through the exponential decay weighting function. On the other hand, the AI measures the variability in a strategy's dynamic adjustments to opponents' evolving tactics, offering insights into its capacity to adapt over time. The decay rate λ is carefully calibrated to strike a balance between responsiveness and stability, ensuring that immediate reactions are captured without sacrificing attention to longer-term trends.

RESULTS

Performance Summary

This section presents aggregated statistics of strategy performance using the core metrics PWR, WRD, and OWR.

The analysis begins with match pair metrics, detailing performance dynamics between strategies. Table 4a presents PWR and WRD statistics, with WRD mean showing the directional performance gap (positive for the

Table 4a. Detailed Metrics for Strategy Match Pairs

Match Pair	PWR Mean	PWR Std.	WRD Mean	WRD Std.
(Const, MRO)	(0.11, 99.74)	(0.34, 0.5)	-99.64	0.75
(Const, RevMirr)	(99.77, 0.18)	(0.61, 0.62)	99.59	1.22
(Const, React)	(0.1, 99.68)	(0.21, 0.66)	-99.58	0.73
(Seq, Mirr)	(99.63, 0.11)	(1.17, 0.38)	99.52	1.32
(Seq, RevMirr)	(0.15, 99.57)	(0.42, 1.17)	-99.42	1.34
(Seq, Excl)	(0.28, 49.8)	(1.06, 4.64)	-49.52	5.45
(RevMirr, MRO)	(18.2, 63.6)	(23.98, 47.95)	-45.4	**71.93**
(React, MRO)	(66.19, 33.12)	(2.69, 1.4)	33.08	1.41
(Mirr, MRO)	(0.41, 33.29)	(1.36, 0.93)	-32.88	2.09
(React, RevMirr)	(21, 39.5)	(40.94, 20.47)	-18.5	61.4
(Excl, MRO)	(40.31, 45.51)	(2.34, 3.08)	-5.19	4.84
(Mirr, React)	(31, 34.5)	(46.48, 23.24)	-3.5	69.72
(ProbWt, MRO)	(31.36, 34.28)	(4.23, 3.83)	-2.92	7.52
(RevMirr, ProbWt)	(32.6, 34.55)	(3.73, 4.12)	-1.95	6.98
(Seq, ProbWt)	(34.43, 32.8)	(3.67, 3.08)	1.63	5.83
(React, ProbWt)	(34.06, 32.99)	(3.08, 3.01)	1.07	5.14
(Rndm, React)	(33.75, 32.7)	(3.09, 3.53)	1.05	5.76
(Mirr, RevMirr)	(33, 34)	(23.8, 47.61)	-1	71.41
(Mirr, Excl)	(33.57, 32.57)	(4.27, 4.19)	0.99	8.2
(Excl, RevMirr)	(32.72, 33.48)	(5.87, 2.02)	-0.76	6.34
(Excl, ProbWt)	(33.8, 33.22)	(2.88, 3.48)	0.58	5.9
(Rndm, ProbWt)	(32.85, 33.11)	(2.9, 2.94)	-0.26	4.97
(Rndm, RevMirr)	(33.54, 33.71)	(3.12, 5.54)	-0.18	7.87
(Rndm, Excl)	(33.1, 33.26)	(3.04, 4.01)	-0.16	6.22
(Const, Mirr)	(0.22, 0.1)	(1.07, 0.29)	0.12	1.13
(Const, Excl)	(49.92, 49.84)	(4.09, 4.44)	0.09	8.47
(Rndm, Const)	(33.34, 33.42)	(3.11, 3.3)	-0.08	5.81
(Mirr, ProbWt)	(14.54, 14.62)	(9.98, 10.1)	-0.08	0.54
(Const, ProbWt)	(0.19, 0.26)	(0.62, 1.07)	-0.07	1.28
(Const, Seq)	(33.31, 33.37)	(0.4, 0.81)	-0.06	1.21
(Rndm, Seq)	(33.42, 33.36)	(3, 3.31)	0.05	5.1
(Seq, React)	(0.2, 0.24)	(1.02, 0.68)	-0.04	1.27
(Excl, React)	(33.35, 33.39)	(1.76, 6.1)	-0.04	7
(Seq, MRO)	(32.9, 32.86)	(1.44, 1.01)	0.03	1.29
(Rndm, Mirr)	(33.01, 33.04)	(3.46, 3.67)	-0.02	6.26
(Rndm, MRO)	(33.47, 33.46)	(2.81, 2.95)	0.02	5.06

first strategy, negative for the second). Figure 1a visualizes absolute WRD mean and variability, emphasizing the magnitude of differences. WRD variability reflects natural fluctuations, derived from non-absolute values for consistency.

Table 4a and Figure 1a reveal interesting patterns in strategy matchups. Some pairs, like Seq vs. Mirr, show one strategy clearly outperforming the other, while others, such as Excl vs. ProbWt, are more evenly matched, suggesting closer competition. High variability in pairs like RevMirr vs. MRO points to inconsistent performance, raising questions about what causes these fluctuations. These findings highlight the diverse dynamics between strategies and suggest areas for deeper analysis in future sections.

Aggregated strategy metrics offer a broader perspective on performance trends across all match pairs. Table 4b and Figure 1b summarize overall strategy effectiveness, smoothing out pair-specific variations to reveal comprehensive strengths and weaknesses.

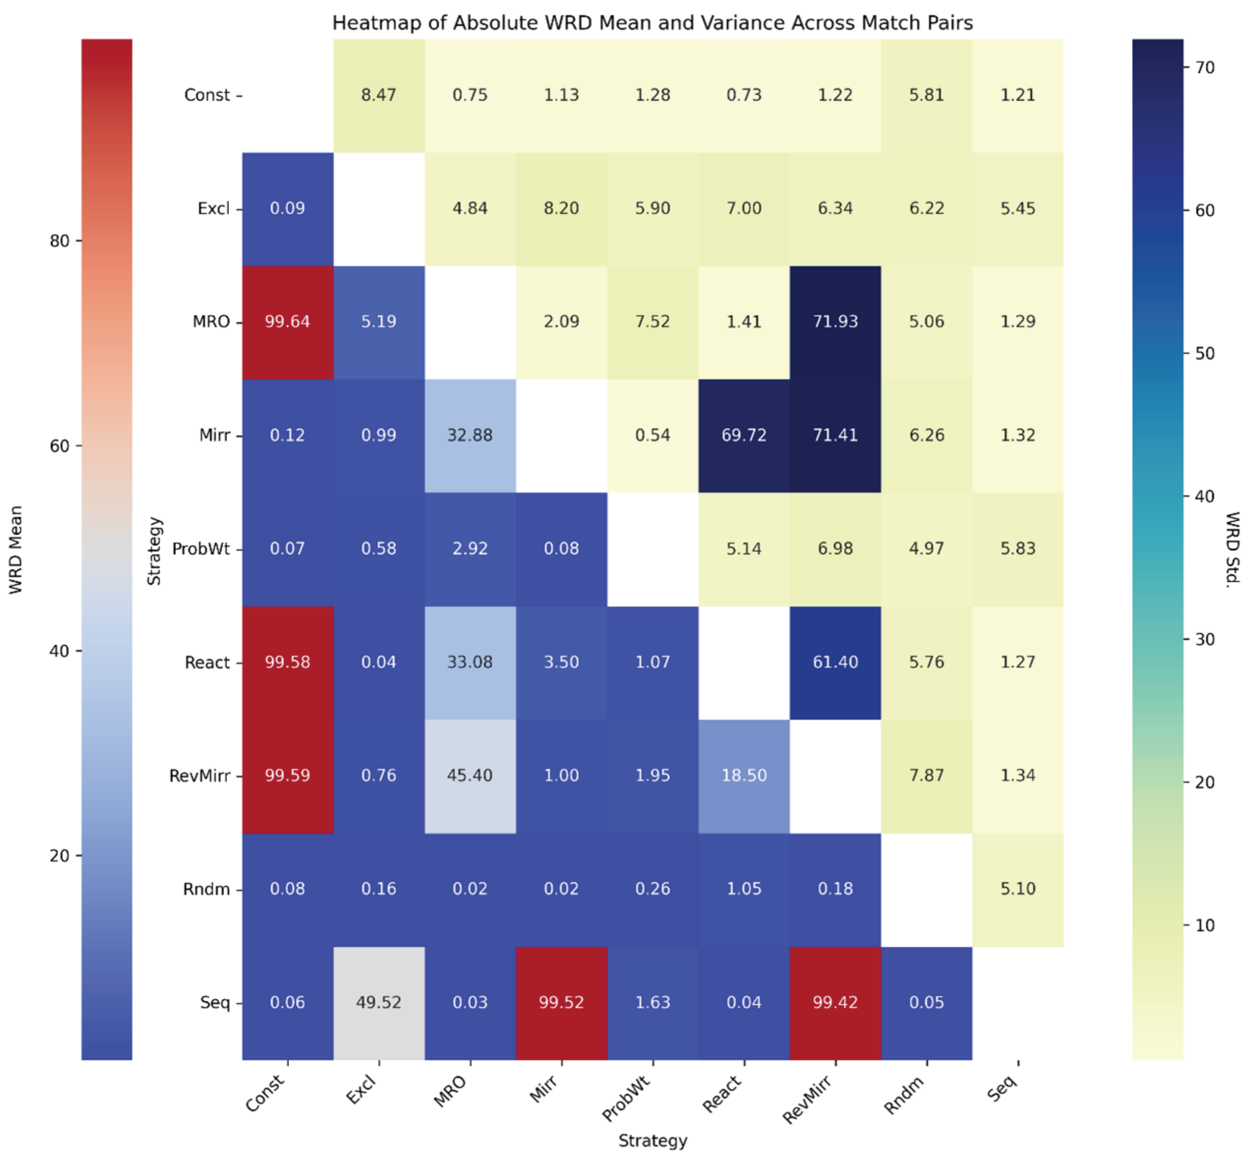

Figure 1a. Heatmap of Absolute WRD Mean and Variability Across Match Pairs.

Table 4b and Figure 1b reveal clear differences in overall strategy performance. MRO leads with the highest mean OWR (46.98%) but shows notable variability (6.09), indicating inconsistent dominance. React follows with a high mean OWR (40.22%) but greater variability (6.35), while Excl balances strength (38.21%) with stability (0.97). Strategies like Mirr and ProbWt show lower OWRs, reflecting weaker effectiveness. Higher variability in strategies such as RevMirr (7.62) and Mirr (7.03) highlights challenges in maintaining consistent performance, raising questions about trade-offs between adaptability and stability for further analysis.

The next section focuses on stability and adaptability evaluation by excluding fixed strategies like Const and Seq, which offer limited insights into variability. The Random (Rndm) strategy is retained as a baseline reference, enabling a clearer comparison with dynamic strategies in staged evolution analysis.

Temporal Analysis

This section examines the evolution of match pairs and individual strategies across Early, Middle, and Late phases. Observations are based on visualized trends, including MA, CUSUM, and MA CUSUM plots for PWR, WRD

Table 4b. Overall Strategy Effectiveness Summary

Strategy	OWR Mean	OWR Max	OWR Min	OWR Std.	OWR Median
MRO	**46.98**	53.57	37.16	6.09	51.11
React	40.22	55.94	31.09	6.35	39.38
Excl	38.21	41.11	33.93	0.97	38.34
RevMirr	36.4	50.94	23.98	7.62	34.58
Rndm	33.31	36.77	28.75	1.05	33.36
Seq	29.29	32.5	27.65	0.64	29.3
Const	27.13	29.75	25.31	0.59	27.12
ProbWt	26.61	31.03	21.94	1.63	26.73
Mirr	18.22	33.75	7.69	7.03	17

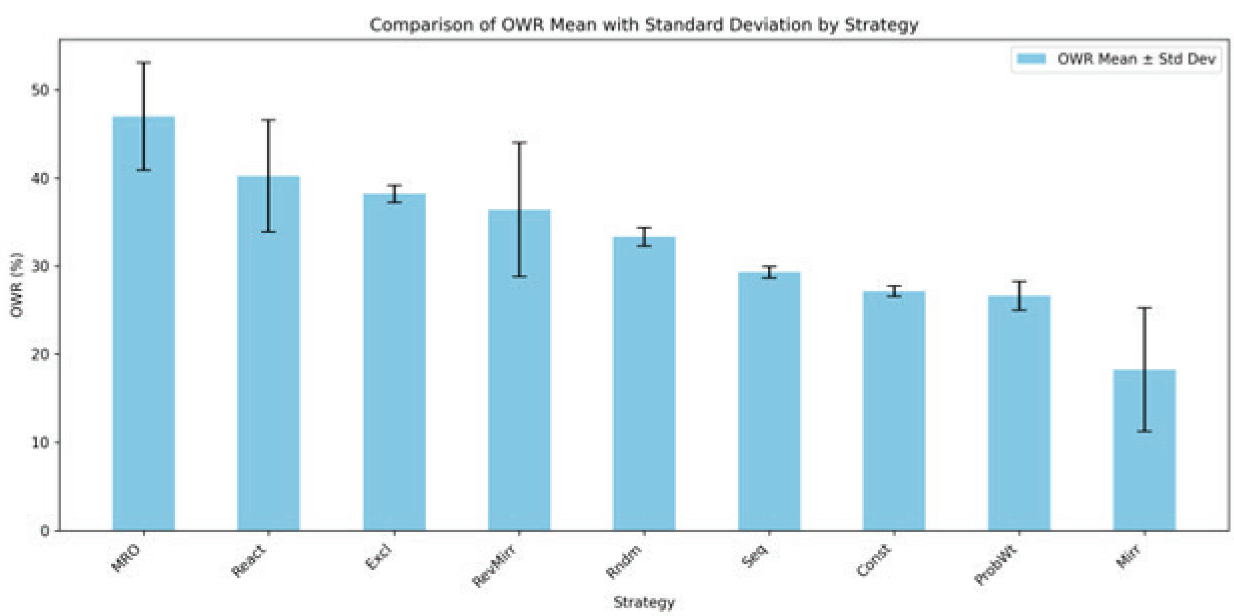

Figure 1b. Comparison of OWR Mean with Std. by Strategy.

and OWR. Summarized insights are presented in tables, while selected examples provide a deeper, illustrative analysis.

The evolution of match pairs is summarized in Tables 5a, 5b, and 5c, grouped by disparity categories: High (H), Balanced (B), and Low (L). These categories were derived from the WRD data presented in Table 4a. These tables, based on visualized data analysis, provide a comprehensive overview of performance trends and stability dynamics across the three phases.

To link the aggregated insights to their visual foundations, three representative match pairs - (React, MRO), (RevMirr, ProbWt), and (Rndm, Excl) - are selected for detailed analysis. Figures 2a, 2b, and 2c present their respective plots, showcasing specific trends and behaviors

In React vs. MRO, React quickly takes the lead in the early phase, with its Moving Average PWR rising above 60%, while MRO stays around 30%. The CUSUM trends show that both React and MRO start below the LCL, showing weak performance early on. As the rounds progress, both steadily improve, crossing the LCL around the middle phase and moving closer to their baselines. By the late phase, React holds a strong lead, shown by stable WRD trends in its favor. React shows greater short-term variability (MASI = 3.30) compared to MRO (MASI = 1.60). Over the long term, React's fluctuations are higher (MASI = 7.89) than MRO's (MASI = 4.00).

In RevMirr vs. ProbWt, ProbWt dominates early, with its Moving Average PWR peaking near 55% while RevMirr trails at 30%. Disparity CUSUM dips below

Table 5a. Phase Evolution of Match Pairs with High Disparity

Match Pair	Early Phase	Middle Phase	Late Phase
(RevMirr, MRO)	MRO leads with fluctuations (7.23, 14.45)	MRO rebounds as RevMirr dips (4.85, 9.71)	Converge with MRO edge, fluctuating (3.26, 6.52)
(React, MRO) *	Both rise sharply; React leads (5.16, 2.49)	React stabilizes, MRO rebounds (0.07, 0.02)	Converge with React leading solidly (0.01, 0.01)
(Mirr, MRO)	MRO peaks early (1.53, 1.22)	MRO hold leads; Mirr rebounds (0.03, 0.04)	Converge with MRO steady lead (0.02, 0.01)
(React, RevMirr)	RevMirr fluctuates and leads (7.42, 3.71)	RevMirr edges React at peak (7.02, 3.51)	Converge with RevMirr lead (5.45, 2.73)

Table 5b. Phase Evolution of Match Pairs with Balanced Disparity

Match Pair	Early Phase	Middle Phase	Late Phase
(Excl, MRO)	MRO leads with fluctuations (0.89, 1.32)	MRO holds steady; recovers (0.24, 0.48)	Converge with MRO fluctuating lead (0.24, 0.37)
(Mirr, React)	React leads with fluctuations (8.1, 4.05)	React holds lead; Mirr rebounds late (7.73, 3.87)	Converge with early Mirr edge, fluctuating disparity (6.86, 3.43)
(ProbWt, MRO)	MRO leads with high fluctuations (3.94, 3.46)	MRO lead dips as ProbWt recovers (0.48, 0.28)	Converge with slight ProbWt advantage (0.16, 0.16)
(RevMirr, ProbWt)*	ProbWt leads with high fluctuations (3.85, 3.71)	ProbWt dips as RevMirr rebounds (0.49, 0.31)	Converge with slight RevMirr advantage (0.15, 0.29)
(React, ProbWt)	Fluctuations as React closes gap (1.88, 1.54)	React lead narrows; ProbWt rebounds (0.3, 0.58)	Converge with slight React advantage (0.26, 0.19)
(Rndm, React)	React leads with fluctuations (1.99, 2.15)	Rndm gains; disparity steadies (0.48, 0.36)	Converge with slight Rndm advantage (0.24, 0.17)
(Mirr, RevMirr)	Competitive with fluctuations (6.62, 13.24)	Oscillating dynamics (2.35, 4.69)	Converge with RevMirr fluctuating lead (1.92, 3.84)

Table 5c. Phase Evolution of Match Pairs with Low Disparity

Match Pair	Early Phase	Middle Phase	Late Phase
(Mirr, Excl)	Mirr leads with fluctuations (1.17, 1.32)	Mirr fluctuates as Excl rebounds (0.39, 0.3)	Converge with Excl fluctuating lead (0.4, 0.33)
(Excl, RevMirr)	Excl gains early advantage (6.7, 1.4)	Excl's lead narrows toward balance (0.52, 0.14)	Converge with minor RevMirr advantage (0.24, 0.17)
(Excl, ProbWt)	Excl lead shrinks (3.86, 5.29)	Excl declines as ProbWt recovers (0.27, 0.42)	Converge with minor ProbWt advantage (0.32, 0.23)
(Rndm, ProbWt)	Competitive; Rndm gains early edge (1.73, 1.33)	Alternating dominance (0.2, 0.67)	Converge with minor Rndm advantage (0.19, 0.35)
(Rndm, RevMirr)	RevMirr leads with fluctuations (3.7, 9.6)	RevMirr declines as Rndm gains (0.3, 0.34)	Converge with minor Rndm advantage (0.34, 0.15)
(Rndm, Excl) *	Excl leads with fluctuations (1.64, 3.53)	Fluctuation narrows; minor Excl edge (0.54, 0.48)	Converge with fluctuating disparity (0.35, 0.24)
(Mirr, ProbWt)	ProbWt narrows lead (2.35, 2.73)	Balance with minor ProbWt advantage (0.91, 0.91)	Converge with diminishing disparity (1.53, 1.54)
(Excl, React)	React leads with fluctuations (0.82, 7.63)	Excl nears balance (0.18, 0.83)	Converge as disparity narrows (0.11, 0.4)
(Rndm, Mirr)	Mirr leads with fluctuations (2.83, 2.93)	Rndm gains; Mirr declines (0.28, 0.28)	Converge with minor Rndm advantage (0.31, 0.24)
(Rndm, MRO)	MRO leads with fluctuations (1.27, 1.23)	Rndm gains; steady slim disparity (0.25, 0.29)	Converge with minor Rndm advantage (0.18, 0.31)

Note: * (Selected for detail analysis), Numbers in parentheses indicate MASI values, where lower values signify higher stability and reduced variability in the phase.

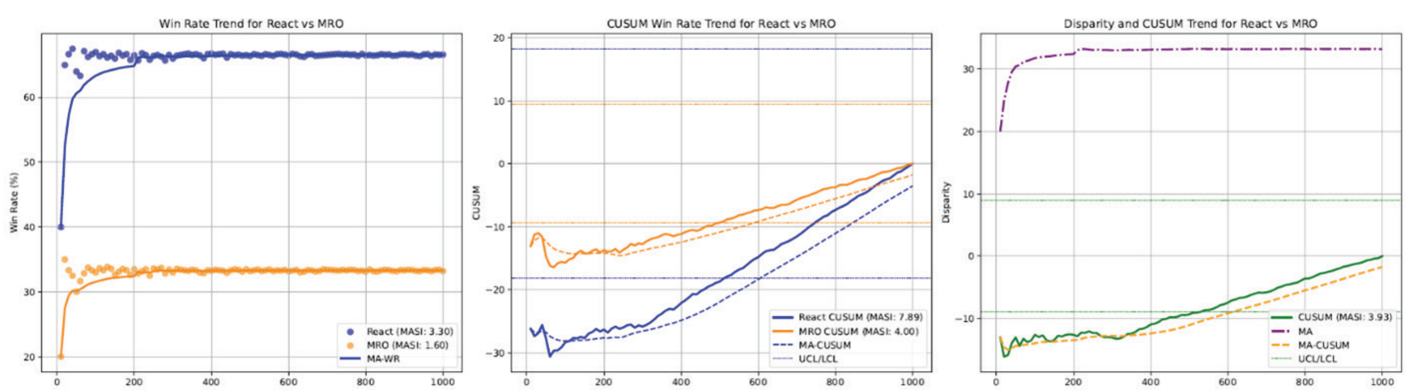

Figure 2a. Performance Trend for React vs MRO.

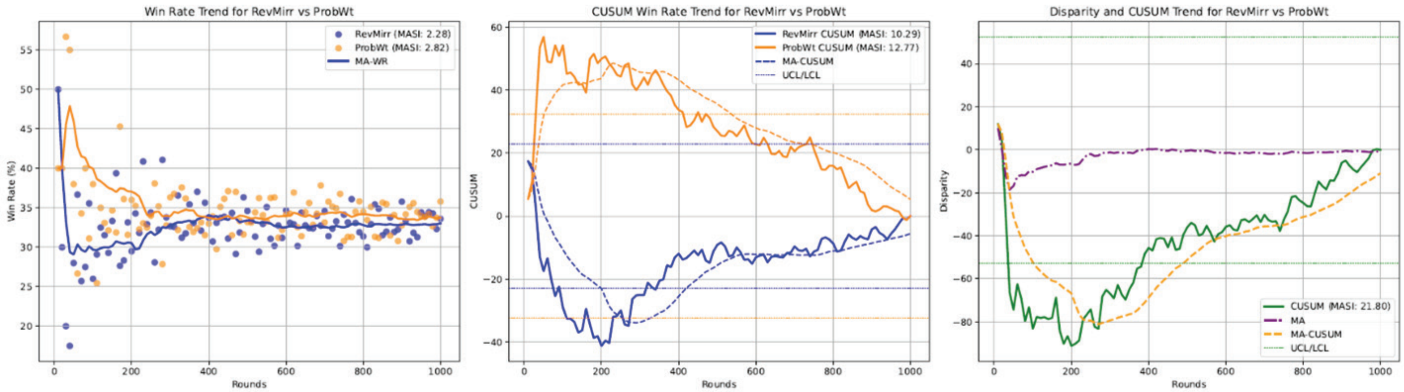

Figure 2b. Performance Trend for RevMirr vs ProbWt.

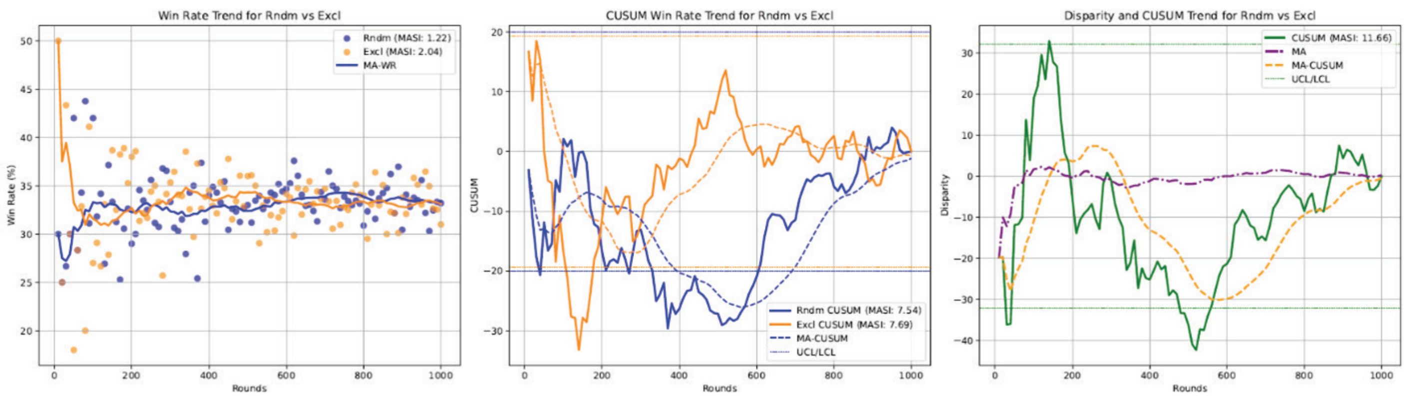

Figure 2c. Performance Trend for Rndm vs Excl.

the LCL late in this phase but recovers steadily. During the middle phase, ProbWt declines as RevMirr rebounds, evident in crossing CUSUM trends and a narrowing WRD. By the late phase, RevMirr maintains a consistent lead as disparity continues to decrease, showcasing its stability. ProbWt fluctuates more in the short term (MASI = 2.82) compared to RevMirr (MASI = 2.28) and in the long term (MASI = 12.77 vs. 10.29). WRD exhibits significant fluctuation (MASI = 21.80) across all rounds.

In Rndm vs. Excl, Excl starts strong with its Moving Average PWR above 40%, while Rndm lags at around 30%. In the early phase, Excl's lead drops quickly, with its CUSUM falling from the UCL to far below the LCL before recovering. In the middle phase, Excl holds a slight lead as its CUSUM peaks and then drops, while Rndm catches up, narrowing the gap. In the late phase, they take turns leading, eventually converging as disparity levels out. The match shows back-and-forth dynamics before settling. Rndm shows lower short-term variability (MASI = 1.22) than Excl (MASI = 2.04), with similar long-term variability (MASI = 7.54 vs. 7.69). WRD shows moderate fluctuation (MASI = 11.66)

The next part of this section shifts focuses from match pairs to the evolution of individual strategy performance over time. Observations are summarized in Table 6, which provides an overview of trends in stability and adaptability for each strategy.

To help explain these trends, two strategies - MRO and Excl - are chosen as examples for closer analysis. MRO shows big swings early on but eventually settles down, giving us an idea of how it handles long-term changes. Excl, on the other hand, quickly adjusts early and adapts well to other strategies. Figures 3a and 3b show how these two strategies perform over time.

The MRO strategy exhibits significant long-term variability (MASI = 16.81 for CUSUM) and moderate short-term fluctuation (MASI = 1.60 for WR). Its win rate fluctuates between 37.5% and 52.5% during the early

Table 6. Phase Evolution of Strategy

Strategy	Early Phase	Middle Phase	Late Phase
Rndm	Declines initially, then rises to a peak (0.62)	Gradually declines, at times below LCL (0.15)	Converges with fluctuating recovery (0.07)
Mirr	Peaks then down with fluctuations (1.71)	Declines further, crossing LCL multiple times (1)	Converges with fluctuating recovery (0.85)
Excl	Peaks early, then down below LCL (0.47)	Gradually recovers with fluctuations (**0.07**)	Converges with steady recovery (0.08)
React	Rises with fluctuations (1.17)	Further rises, up UCL before declining (1.13)	Converges with declining fluctuations (0.86)
RevMirr	Rises with fluctuations (1.49)	Peaks near UCL, then gradually declines (1.39)	Converges with fluctuating recovery (0.63)
ProbWt	Initial rise, fluctuating decline (0.58)	Gradual decline, crossing below LCL (0.21)	Converges with rising trend (0.2)
MRO	Initial rise, steady decline (**1.73**)	Below LCL before gradual recovery (1.25)	Converges with steady recovery (0.81)

Note: Numbers in parentheses indicate MASI values, where lower values signify higher stability and reduced variability in the phase.

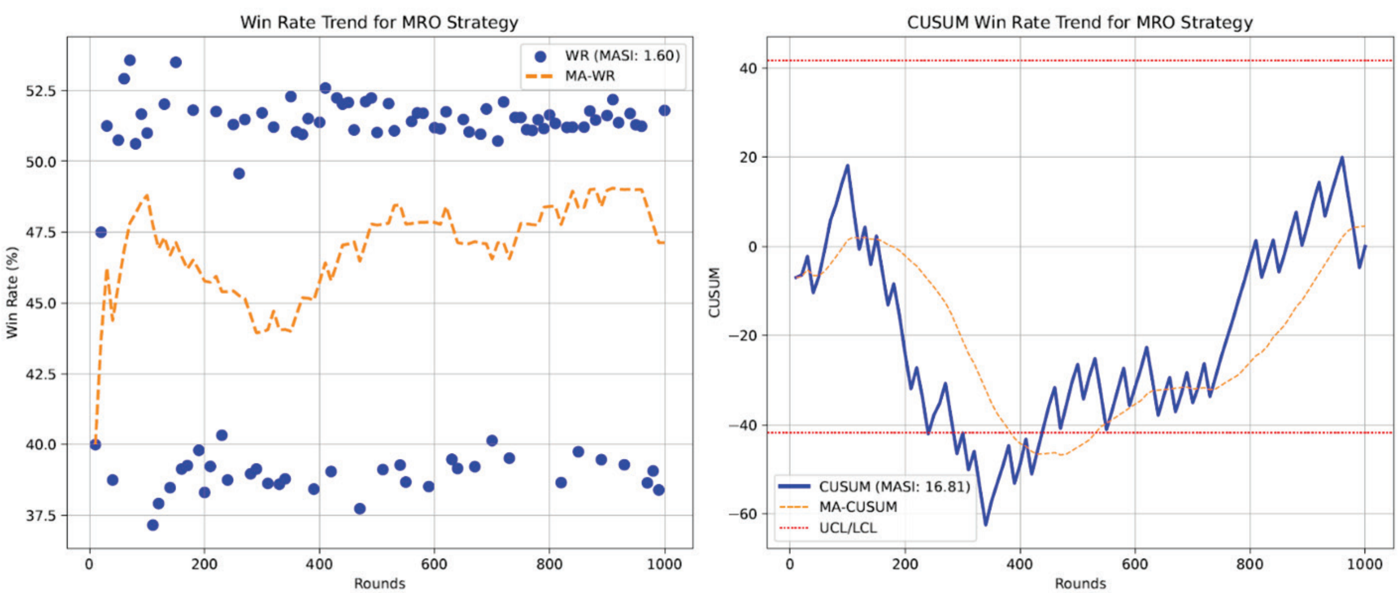

Figure 3a. Performance Trend for MRO Strategy.

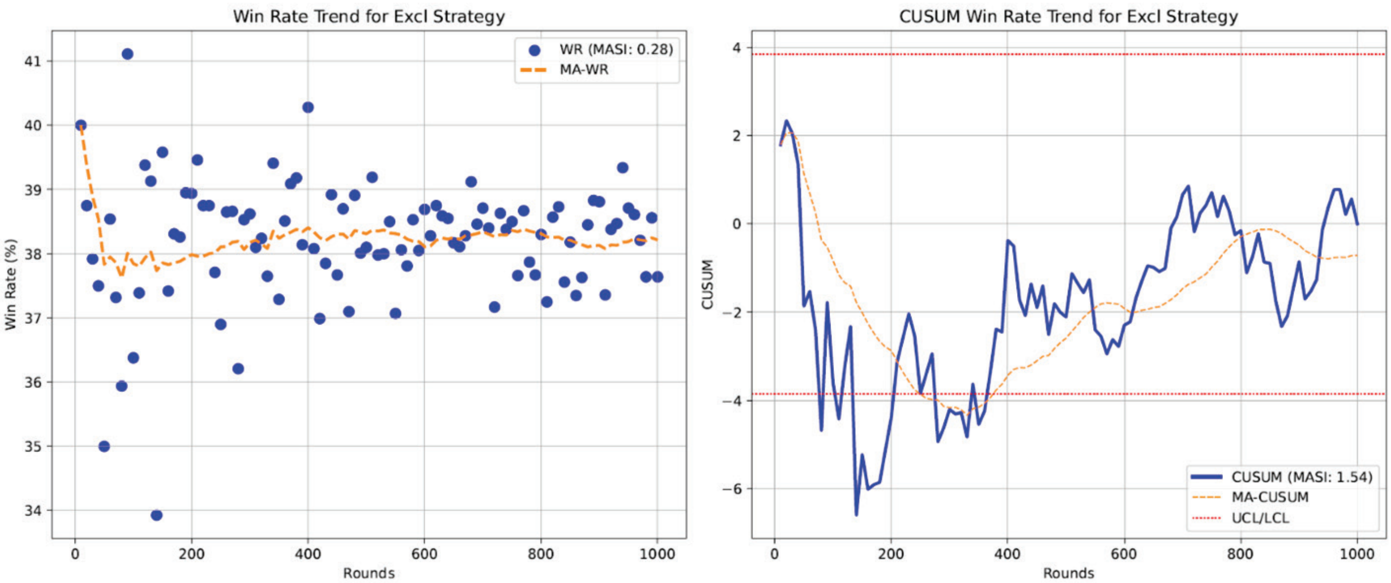

Figure 3b. Performance Trend for Excl Strategy.

phase, with the MA WR stabilizing near the mean of 46.98%. The CUSUM plot shows an initial rise followed by a sharp decline after the 100th round, reaching a minimum of -62.51, well below the LCL, during the middle phase. By the late phase, MRO steadily recovers, showing smaller fluctuations and converging toward balance. This demonstrates MRO's adaptability despite early volatility.

Excl demonstrates steadier performance, with low short-term variability (MASI = 0.28 for WR) and modest long-term variability (MASI = 1.54 for CUSUM). Its MA WR remains close to the overall mean of 38.21%. Early fluctuations subside quickly, stabilizing Excl through the middle and late phases. The CUSUM plot shows an early decline, bottoming at -6.6, well below the LCL, followed by a slow recovery. By the late phase, Excl converges near its baseline, reflecting consistent recovery and balance. Its low variability highlights a stable but less dynamic performance.

Adaptability Insights

This section examines how match pairs and individual strategies adapt over time. The *AI* captures the magnitude of performance shifts over time, while the *S* focuses on short-term responsiveness. By combining these metrics with visual trends like DW CUSUM, deeper insights are gained into how strategies evolve and adapt over time.

The adaptability dynamics of match pairs are summarized in Tables 7a, 7b, and 7c, aligned with the categories in Tables 5a, 5b, and 5c. Match pairs in each table are ranked by their combined values, which show how reactive the strategies are over the long term. Steeper slopes reflect short-term adaptability, showing how quickly strategies respond to changes. Together, these metrics give a clear picture of how match pairs adjust and evolve over time.

To ensure consistent comparison from various perspectives, the same three match pairs - (React, MRO), (RevMirr, ProbWt), and (Rndm, Excl) - are revisited for further analysis. Figures 4a, 4b, and 4c provide a detailed visualization of their adaptability dynamics, further supporting the analysis.

In the React vs. MRO match (Figure 4a), React demonstrates a higher decay-weighted win rate slope of 0.0490, reflecting its steady and significant improvement in performance over the rounds. Its AI of 0.36 confirms React's ability to adjust consistently while responding to changes. In contrast, MRO's slope of 0.0245 and AI of 0.18 indicate a more subdued adaptability, stabilizing earlier with limited responsiveness to performance shifts. The CUSUM trends show React's sharp progression, while MRO levels off, reflecting their distinct adaptability styles.

The RevMirr vs. ProbWt match (Figure 4b) highlights contrasting dynamics. RevMirr achieves a slope of 0.0246, signifying a steady but moderate improvement in adaptability. Its AI of 0.41 shows consistent adjustments without sharp shifts. ProbWt, with a slope of 0.0224 and a higher AI of 0.64, adapts more dynamically over time, reflecting greater sensitivity to performance shifts. The CUSUM plot reveals that ProbWt takes a more responsive path, while RevMirr maintains a balanced and stable adjustment over the rounds.

In the Rndm vs. Excl match (Figure 4c), both strategies show comparable slopes of 0.0249 for Rndm and 0.0242 for Excl, indicating similar rates of short-term adaptability. However, their AIs diverge, with Rndm achieving a higher AI of 0.41 compared to Excl's 0.30, suggesting Rndm adapts more actively over time. The CUSUM trends illustrate Rndm's consistent improvements, while Excl exhibits fluctuations, reflecting its less stable adaptability pattern.

Table 7a. Adaptability Dynamics of Match Pairs with High Disparity

Match Pair	AI	Slope	Observation
(RevMirr, MRO)	(2.82, 5.63)	(0.0053, 0.0612)	MRO adapts late; RevMirr adjusts steadily
(React, RevMirr)	(4.2, 2.1)	(0.0099, 0.0309)	React adapts dynamically; RevMirr stabilizes consistently
(Mirr, MRO)	(0.28, 0.02)	(-0.0008, 0.0239)	Mirr adapts early but weakens; MRO remains steady
(React, MRO) *	(0.36, 0.18)	(0.049, 0.0245)	React adapts dynamically; MRO stabilizes early

Table 7b. Adaptability Dynamics of Match Pairs with Balanced Disparity

Match Pair	AI	Slope	Observation
(Mirr, React)	(3.82, 1.91)	(0.0269, 0.0224)	Mirr adapts actively; React improves steadily
(Mirr, RevMirr)	(1.38, 2.77)	(0.0235, 0.0248)	Mirr adapts actively; RevMirr adjusts consistently
(ProbWt, MRO)	(0.81, 0.71)	(0.0255, 0.0218)	ProbWt adapts smoothly; MRO progresses steadily
(RevMirr, ProbWt) *	(0.41, 0.64)	(0.0246, 0.0224)	RevMirr adapts steadily; ProbWt stabilizes well
(React, ProbWt)	(0.37, 0.46)	(0.0231, 0.0251)	React adapts steadily; ProbWt adapts faster but variably
(Rndm, React)	(0.34, 0.34)	(0.0234, 0.0246)	Both adapt steadily with balanced trends
(Excl, MRO)	(0.13, 0.38)	(0.0288, 0.0339)	MRO adapts faster; Excl stabilizes slower

Table 7c. Adaptability Dynamics of Match Pairs with Low Disparity

Match Pair	AI	Slope	Observation
(Excl, ProbWt)	(0.63, 0.44)	(0.0216, 0.0256)	ProbWt adapts; Excl adjusts smoothly
(Mirr, ProbWt)	(0.78, 0.81)	(0.0106, 0.0104)	Mirr adapts gradually; ProbWt adapts smoothly
(Mirr, Excl)	(0.32, 0.55)	(0.0232, 0.0253)	Mirr adapts steadily; Excl adapts dynamically
(Rndm, RevMirr)	(0.28, 0.55)	(0.0237, 0.022)	Rndm adapts steadily; RevMirr shows dynamic adjustment
(Excl, React)	(0.11, 0.64)	(0.0241, 0.0221)	Excl adapts minimally; React stabilizes with steady adjustments
(Rndm, Mirr)	(0.32, 0.22)	(0.0249, 0.0233)	Rndm adapts steadily; Mirr stabilizes early
(Rndm, Excl) *	(0.41, 0.3)	(0.0249, 0.0242)	Rndm adapts steadily; Excl stabilizes late
(Excl, RevMirr)	(0.5, 0.15)	(0.0228, 0.0238)	Excl adapts actively; RevMirr adjusts minimally
(Rndm, MRO)	(0.2, 0.3)	(0.0245, 0.0231)	Rndm adapts steadily; MRO stabilizes gradually
(Rndm, ProbWt)	(0.19, 0.26)	(0.0239, 0.0241)	Rndm adapts steadily; ProbWt stabilizes moderately

Note: * (Selected for detail analysis).

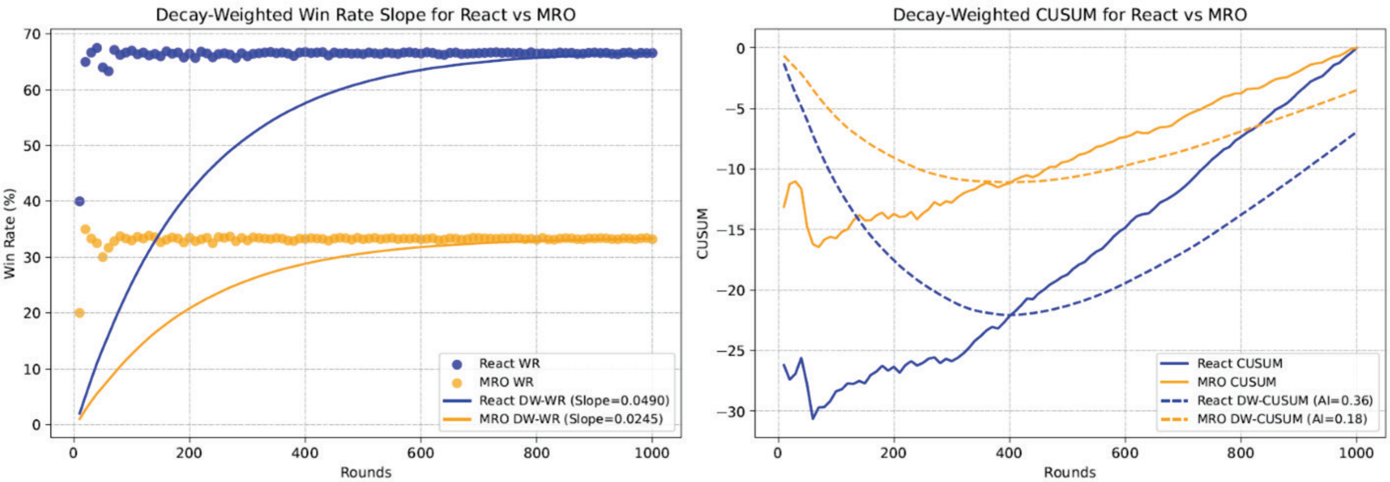

Figure 4a. DW Analysis for React vs MRO

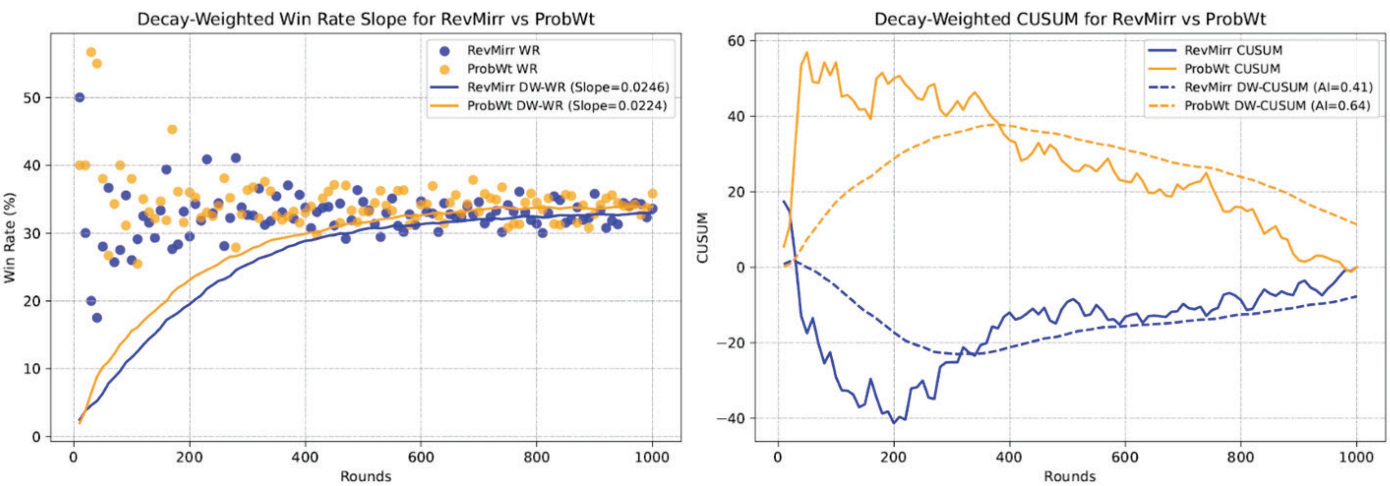

Figure 4b. DW Analysis for RevMirr vs ProbWt.

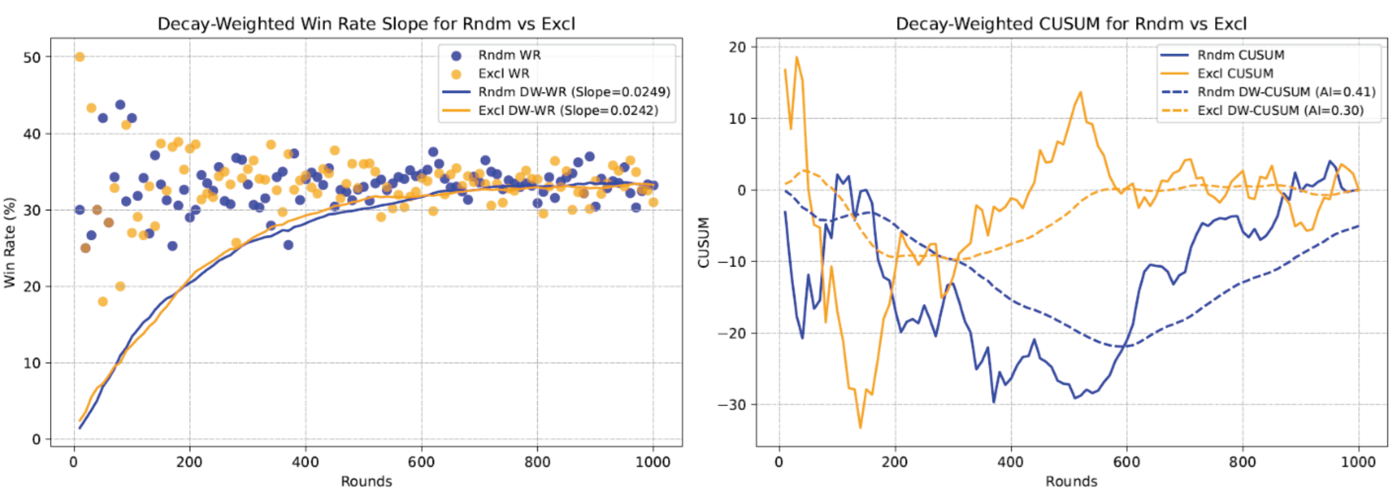

Figure 4c. DW Analysis for Rndm vs Excl.

These three matches provide a deeper understanding of how adaptability dynamics play out across different strategies. By analyzing *S* and AI, the evolution and adjustment of strategies are observed over both short and long terms, revealing distinctive patterns of reactivity and stabilization.

The next section shifts focus to individual strategy adaptability, analyzing performance trends through WR *S* and AIs. Table 8 summarizes how each strategy adjusts and stabilizes over time, providing insights beyond match pair interactions.

Table 8. Adaptability Dynamics of Strategy

Strategy	AI	Slope	Observation
Rndm	0.1	0.2414	Stable performance with minimal adaptability shifts
Mirr	0.5	0.1332	Gradual adaptability with moderate variability
Excl	0.07	0.276	Strong adaptability with minimal variation
React	0.66	0.2787	High adaptability with consistent shifts
RevMirr	0.51	0.2547	Moderately adaptable with steady improvements
ProbWt	0.2	0.1967	Adaptation is gradual, stabilizing effectively
MRO	**0.71**	**0.3567**	Strong adaptability with consistent progression

To maintain consistency across perspectives, the focus now shifts to the individual dynamics of the MRO and Excl strategies, previously analyzed within match pairs. This analysis delves into their adaptability, utilizing DW-WR and DW-CUSUM trends to evaluate performance over time. Figures 5a, 5b provide detailed visualizations, offering deeper insights into their adaptability and long-term performance patterns.

The MRO strategy shows a robust improvement in win rates, reflected in a high slope of 0.3567 for its DW-WR. This upward trajectory indicates a fast stabilization phase, supported by a relatively high AI of 0.71 in the DW-CUSUM chart, showcasing its ability to respond dynamically to changes during the early and middle phases. However, the consistent decline in CUSUM toward the later rounds suggests some challenges in maintaining adaptability over extended gameplay.

In contrast, the Excl strategy exhibits a more tempered progression, with a moderate DW-WR slope of 0.2760. This indicates a steady but less aggressive adaptation curve compared to MRO. The DW-CUSUM reveals minimal shifts, with a notably low AI of 0.07, suggesting that Excl prioritizes stability over dynamic responsiveness. The flat trajectory in both DW-WR and CUSUM during later rounds highlights its ability to maintain consistent performance without significant variability.

The two strategies show different adaptability styles: MRO responds quickly to an early edge, while Excl follows a steady path, reducing risk and variation. This highlights how different strategies affect long-term performance in changing situations.

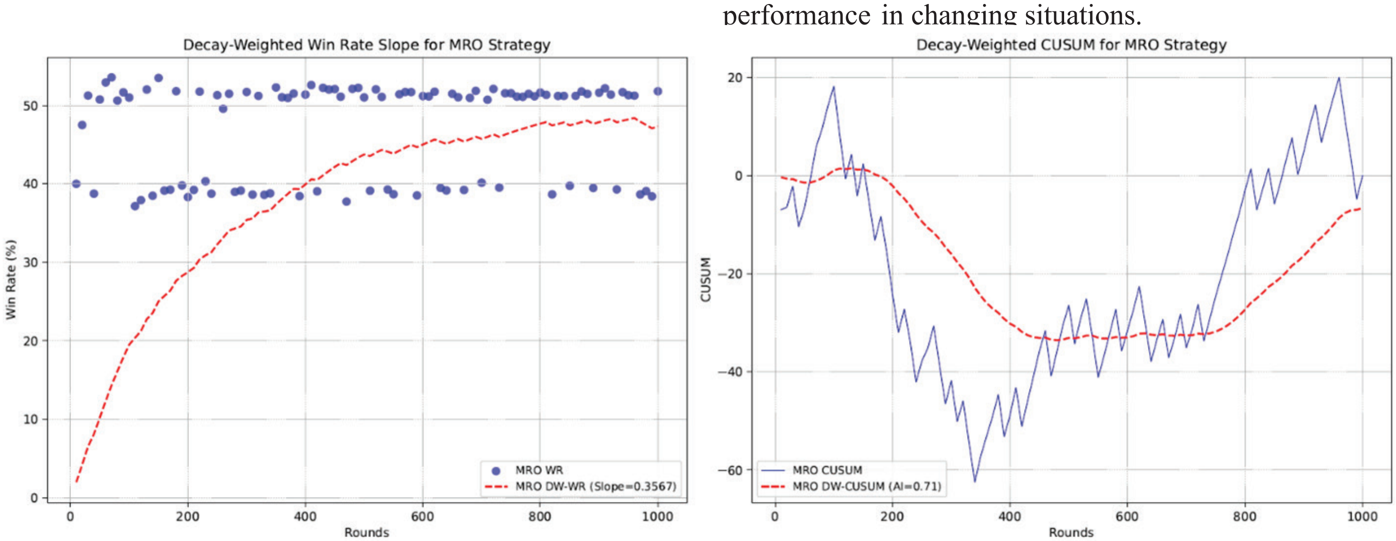

Figure 5a. DW Analysis for MRO Strategy.

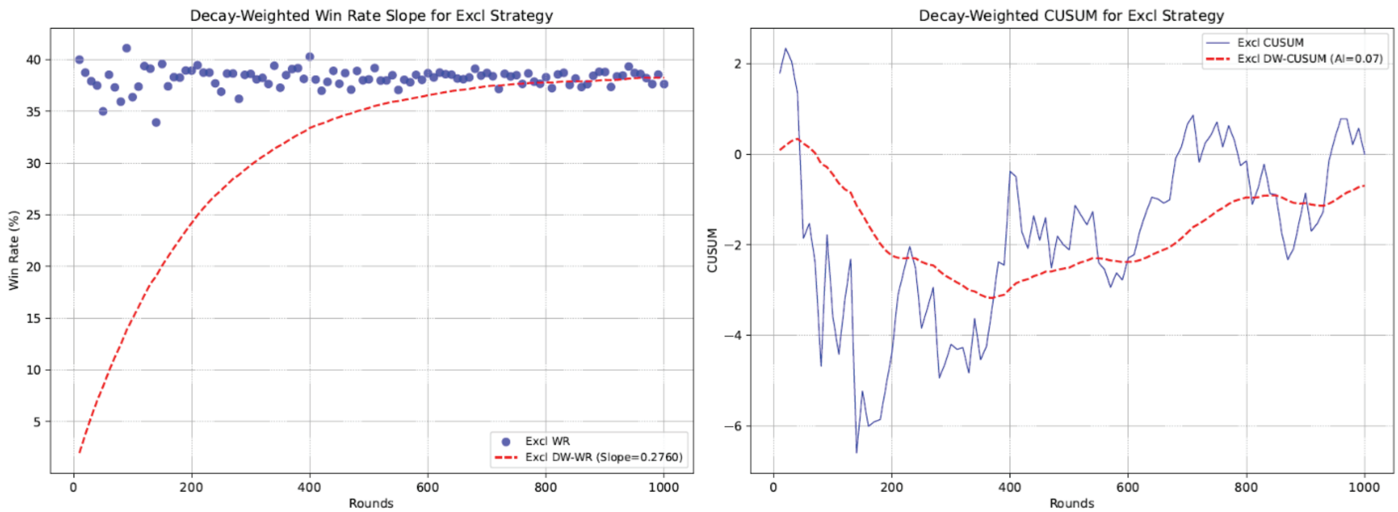

Figure 5b. DW Analysis for Excl Strategy.

CONCLUSION

This research investigates nine RPS strategies, encompassing random, reactive, and adaptive approaches, to evaluate their performance, stability, and adaptability through simulated matchups over progressively increasing rounds. Using statistical tools such as MA, CUSUM Control Charts and DW Metrics, these strategies were analyzed across early, middle, and late phases of gameplay. Key metrics, including WR Mean, MASI, AI, and *S*, were employed to uncover trends in both aggregated results and specific match pairs.

The findings reveal clear patterns in how strategies perform. MRO stood out as the best performer, with the highest WR Mean (46.98%), AI (0.71), and *S* (0.3567). These findings show that MRO stands out for its ability to adapt quickly and stay effective over the long term, with the shortest ramp-up time of all strategies. Match pair analysis revealed that MRO beat other adaptive and reactive strategies but struggled against React, whose sharp short-term adjustments countered MRO's focus on recent trends. However, MRO has trouble staying stable over time. Its MASI CUSUM of 16.81 is the second highest among all strategies, and its performance in the early phase showed a lot of variation before stabilizing later. React's sharp, short-term adjustments worked well against MRO's focus on recent trends. Stability analysis revealed that all strategies became more stable over time, as shown by declining MASI values across phases. Excl, although classified as reactive, showed outstanding stability with the lowest MASI values across all phases (early: 0.47, middle: 0.07, late: 0.08), making it the most reliable performer. Its steadiness resulted in closely matched games with minimal disparities against other reactive and adaptive strategies. On the other hand, ProbWt, while adaptive, struggled because it relied on all past moves. This caused it to overreact to older trends and respond too slowly to recent changes, reflected in its lower AI (0.2). The Rndm strategy, which lacks patterns, served as a baseline, with a WR Mean of 33.31% and AI of 0.1, offering a point of comparison for evaluating other strategies.

This research shows how strategies work in competitive situations and highlights why adaptability is so important. MRO shines because it can quickly respond to recent trends while avoiding overreacting, making it a strong performer. However, its struggles against React point to a weakness—it needs to combine long-term adaptability with quick, short-term reactions. These findings remind us that no single strategy is perfect. The best approach often depends on the situation, the competition, and how well a strategy matches up against its opponent.

The principles of MRO are reflected in real-world scenarios where adapting to recent changes is essential. For example, the Digital SAT adjusts question difficulty based on a student's answers to previous questions.

By focusing on recent performance trends without overreacting to individual mistakes, it ensures a fair and accurate assessment of abilities, much like MRO balances responsiveness and stability. Similarly, Algorithmic Trading Systems analyze recent market data, such as stock prices and trading volumes, to make quick buy or sell decisions. These systems prioritize current trends while considering historical data to avoid overreacting to short-term fluctuations, enabling steady performance even in volatile markets. These examples show how MRO's way of adapting to what's happening now while learning from the past works well in changing environments.

While this research offers valuable insights, it has some limitations. The simulation settings, like MRO's fixed observation window of 3 rounds and the absence of multiplayer dynamics or player-specific traits, simplify real-world scenarios. The round-robin pairing system, while thorough, doesn't mimic the unpredictability of actual matchups. The strategy designs also lack complexity in detecting advanced patterns. For instance, MRO performs well with recent trends but struggles in situations requiring deeper analysis. Similarly, ProbWt's use of all past data leads to overreactions, showing the need for better balance between short- and long-term focus.

Lastly, the statistical tools used, such as the decay rate ($\lambda = 0.05$) and moving average window size ($N = 20$), may not always be the best fit. Exploring other decay models or introducing a decay-weighted stability index could provide clearer insights into how strategies adapt and stabilize. Addressing these areas could make future research even more impactful.

Future work could make MRO even better by addressing its challenges and trying new ideas. One key improvement is balancing short- and long-term trends. Right now, MRO excels at adapting to recent changes but sometimes overreacts. Using weighted observation windows or smoothing techniques could help it stay steady while remaining flexible. Adding dynamic observation windows that adjust to changing conditions, like shrinking during volatile moments and expanding during calm periods, could make it even more powerful. Another exciting direction is better understanding opponents. MRO focuses on recent moves, but learning long-term patterns or repeated behaviors could make its decisions even sharper. Testing MRO against multiple opponents would also show how it adapts in different situations, helping fine-tune settings like decay rates. RPS is a great testing ground for these ideas because it's simple yet rich enough to reveal how strategies evolve. Exploring these enhancements could take adaptive strategies like MRO to the next level, showing how to "Adapt to Win" in all kinds of challenges.

ACKNOWLEDGEMENT

The author expresses heartfelt gratitude to the anonymous reviewers for their insightful comments and suggestions that significantly enhanced the quality of this paper.

REFERENCES

1. Seale DA, Phelan SE. Bluffing and betting behavior in a simplified poker game. *J Behav Decis Mak*. 2009; 23 (4): 335–352. https://doi.org/10.1002/bdm.658.
2. Bhatia M. Intelligent system of game-theory-based decision making in the smart sports industry. *ACM Trans Intell Syst Technol*. 2021; 12 (3): 1–23. https://doi.org/10.1145/3447986.
3. Zaytsev A, Kuznetsova T, Shirin A. Formation of investment behavior strategy using the game-theoretical method. *TEM J*. 2021; 10 (2): 673–680. https://doi.org/10.18421/TEM102-23.
4. Luce RD, Raiffa H. Games and decisions: Introduction and critical survey. Wiley; 1957. ISBN: 978-0471553410 (original edition), 978-0486659435 (Dover reprint).
5. Axelrod R. The evolution of cooperation. *Princeton University Press*; 1984. ISBN: 978-0465021215.
6. Koller D, Megiddo N. The complexity of stochastic games. *J Comput Syst Sci*. 1992; 44 (1): 85–106.
7. Friedman D, Oprea R. A continuous prisoner's dilemma experiment. *J Econ Behav Organ*. 2012; 81 (1–2): 1–16.
8. Fudenberg D, Levine DK. The theory of learning in games. *MIT Press*; 1998. ISBN: 978-0262061940.
9. Camerer CF. Behavioral game theory: Experiments in strategic interaction. *Princeton University Press*; 2003. ISBN: 978-0691090399.
10. Nowak MA, Sigmund K. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature*. 1993; 364 (6432): 56–58. https://doi.org/10.1038/364056a0.
11. Law AM, Kelton WD. Simulation modeling and analysis. 4th ed. *McGraw-Hill*; 2007. ISBN: 978-0073294414.
12. Wilkinson L. The grammar of graphics. 2nd ed. *Springer*; 2005. ISBN: 978-0-387-24544-7.
13. Cleveland WS. Visualizing data. *Hobart Press*; 1993. ISBN: 978-0963488404.
14. Brockwell PJ, Davis RA. Introduction to time series and forecasting. 3rd ed. *Springer*; 2016. ISBN: 978-3-319-29852-8.
15. Hyndman RJ, Athanasopoulos G. Forecasting: Principles and practice. 2nd ed. *OTexts*; 2018. ISBN: 978-0987507112.

https://doi.org/10.32614/CRAN.package.fpp2.
16. Montgomery DC. Introduction to statistical quality control. 6th ed. *Wiley*; 2009. ISBN: 978-0470169926.
17. Woodall WH. Controversies and contradictions in statistical process control. *J Qual Technol*. 2000; 32 (4): 341–350. https://doi.org/10.1080/00224065.2000.11980013.
18. Layer E, Tomczyk K. Measurements, modelling and simulation of dynamic systems. *Springer*; 2010. ISBN: 978-3-642-04587-5 (Hardcover), 978-3-642-04588-2 (eBook). https://doi.org/10.1007/978-3-642-04588-2.