Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Adaptive Risk Taking Under Threshold Selection: An Agent-Based Model

Martin Höppner

martin.hoeppner@b-tu.de

Brandenburg University of Technology Cottbus-Senftenberg

Research Article

Keywords: Agent-based economics, heterogeneous agents, noisy imitation, threshold se- lection, adaptive heuristics, risk taking

Posted Date: October 5th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10214942/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Adaptive Risk Taking Under Threshold Selection: An Agent-Based Model

Anonymous manuscript for double-blind review

Abstract

Thresholds change the value of risk. A firm just below a solvency cutoff, a worker just short of a promotion, or a project just outside a funding line might benefit from variance that allows them to cross the threshold; conversely, the same variance could be detrimental to an agent already above the threshold. This paper investigates this intuition using a behavioral agent-based model where heterogeneous agents choose between risky options, are ranked, and are replaced by noisy copies of agents who clear the threshold. The model design isolates the effect of the number of agents falling below the cutoff from the impact of how strongly a failure changes the population. In the context of evolving competition, a simple below/above-threshold switch outperforms fixed risk behavior. The addition of time creates a second, larger layer of benefits, while the distance to the cutoff adds a smaller gain concentrated under intermediate cutoff pressure. This hierarchy holds for fixed-composition and auxiliary panels, even when employing alternative threshold rules, weaker information, matched turnover, optimized fixed baselines, and placebo state variables. At the most severe cutoff, the influence of extra state information approaches parity. This comparison serves to identify the conditions under which adaptive risk heuristics offer a competitive advantage in threshold selection, versus when the basic below/above switch suffices.

Keywords: Agent-based economics; heterogeneous agents; noisy imitation; threshold selection; adaptive heuristics; risk taking

JEL Classification: C63, D81, C73, D91

1 Introduction

Thresholds shape economic and organizational selection whenever firms must keep capital above regulatory cutoffs, workers face promotion bands, research teams compete for funding ranks, or political actors operate around vote thresholds. A firm just below a capital requirement may accept variance that gives it a chance to cross the line; conversely, a firm just above the same line may avoid that variance because one bad draw can trigger failure. Similar tradeoffs arise around promotion bands, grant rankings, and vote thresholds. Agents in these settings interact through relative position. They may also observe rank, payoff forecasts, distance from a cutoff, and time left before evaluation. A risk rule can matter twice: it can improve threshold performance for a given population, and it can alter composition when successful agents are copied.

This model represents a stylized threshold-selection environment where agents repeatedly select between risky options, accumulate noisy outcomes, and are subject to rank-based replacement. Agents falling below a certain cutoff may be replaced by noisy copies of agents that cleared it. The emergence of aggregate rule composition is determined by the interplay of threshold performance, imitation, and mutation. It should be noted that this exercise does not aim to estimate stable human risk preferences or to rank threshold institutions by welfare, but rather to determine which adaptive heuristics prevail in fixed and evolving threshold competitions.

The primary design distinguishes between two institutional margins that the older one-parameter design combined. Cutoff stringency q measures the proportion of the population below the cutoff, while replacement intensity r measures the fraction of below-cutoff agents replaced at the trial boundary. The product of these two quantities approximates the realized turnover. This separation enables an investigation into whether the observed result stems from the severity of the cutoff, the strength of the replacement, or their interaction.

Four rule classes compete. A fixed rule employs the same mean-variance risk weight across all states. H0 switches risk weights above and below a cutoff. H1 adds information

from the remaining rounds to that switch. H2 incorporates the absolute point distance to the cutoff. Variance can be beneficial below a threshold if an upside movement can cross the cutoff, but it can be detrimental above a threshold if a downside movement loses its lead. The values of that sign switch are influenced by the time left and the distance from the cutoff.

The primary estimand is evolving-competition performance. In that panel, replacement and mutation continue during Stage 2, meaning the outcome captures head-to-head performance while population composition can change under continued threshold selection. A smaller fixed-composition panel holds Stage-2 composition fixed and tests direct post-training performance. This fixed panel serves as a diagnostic tool, as the underlying theory is concerned with performance under ongoing replacement rather than just fixed-composition choice quality.

The numerical evidence suggests a small hierarchy of performance. In the evolving panel, H0 outperforms fixed behavior with a mean tie-aware win score of 0.804. H1 achieves a higher win score of 0.665 compared to H0, and H2 also shows an improvement over H0, albeit with a smaller gain, achieving a win score of 0.570. The fixed-composition diagnostic yields the same ordering, albeit with smaller magnitudes. The cutoff breakdown reveals that H1 is most effective when cutoff pressure is high but not saturated, whereas H2 is concentrated at intermediate cutoff pressure. Auxiliary panels confirm this hierarchy even after adjustments to threshold implementation, information quality, fixed baselines, and the influence of irrelevant placebo signals.

These results apply most directly to threshold settings in which agents can observe position, payoff moments, remaining time, and the current cutoff with reasonable accuracy. They illustrate when a below/above risk switch is sufficient and when additional state variables provide added value. The extreme cutoff boundary represents an in-model convergence pattern, which is generated by cutoff severity, turnover, and learned risk weights, and is not a welfare result.

Section 2 situates the present contribution in the context of previous research. Section 3 details the simulation protocol, while Section 4 offers a formal intuition for the distance and horizon channels. Section 4.1 outlines the design checks and evidence criteria, Section 5 presents the results, and Sections 6 and 7 conclude.

2 Related Literature and Positioning

Reviews and measurement studies find persistence, but they also document state sensitivity, weak convergence across instruments, and change after shocks or learning (Schildberg-Hörisch, 2018; Frey et al., 2017; Mata et al., 2018; Bagaïni et al., 2025). This record justifies studying state dependence, while leaving open which state-dependent heuristics perform well when agents interact through repeated threshold selection.

Baseline sign logic derives from reference-point theories. Prospect theory predicts different risk responses above and below a reference point, and threshold or survival-pressure models explain why that sign shift can be adaptive when continuation value changes sharply (Kahneman and Tversky, 1979). McDermott et al. (2008) offer a threshold form for this intuition: higher variance can be beneficial below a survival cutoff but detrimental above it.

Repeated competition introduces the update process as a problem. Roos and Nau (2010) and Roos et al. (2010) demonstrate that the success of conditional risk rules hinges on the update process in addition to the payoff environment. The present study focuses on the agent-based machinery as the object of analysis: heterogeneous rules compete, successful rules are copied with noise, and the population composition changes endogenously.

The concepts of remaining time and distance from the cutoff are plausible within the literature on behavioral agency and performance-feedback. Specifically, organizations performing below salient targets have been observed to take on more risk, whereas those performing above target tend to protect their margin (March and Shapira, 1987; Laughhunn et al., 1980; Wiseman and Gomez-Mejia, 1998). Furthermore, the response to a shortfall is contingent

upon the remaining time and the pressure of the deadline (Profitos et al., 2022; Kolev and McNamara, 2020). While these literatures suggest the existence of distinct channels for time and distance, they do not provide insight into which channel holds greater competitive value when making repeated threshold selections.

Finite-population dynamics help bound the paper's scope, as stronger-selection settings can diverge materially from replicator models (Hofbauer and Sigmund, 2003; Roca et al., 2009; Su et al., 2019; Domingos et al., 2023), and the high-pressure boundary is part of the paper's contribution, not a side result.

The literature motivates three model checks. The first is H0, the baseline relative-position prediction, which posits that being below versus above the threshold changes the direction of risk response. The second is H1, the time-horizon prediction, which argues that fewer remaining rounds make terminal position more salient and strengthen position-consistent adjustment. The third is H2, the distance prediction, which suggests that absolute distance from the cutoff can refine the above/below switch when threshold pressure leaves room for margin-sensitive adjustment. Section 4.1 maps these predictions into design checks and directional rows.

3 Simulation Design and Protocol

The simulation investigates which state-dependent rules perform well under repeated threshold selection, with technical definitions provided in Appendix B.

3.1 Simulation Design Rationale

Population composition is tracked over time because one-shot payoff comparisons cannot reveal how rule classes change under repeated updating. Since work on population dynamics shows that the long-run composition depends on update dynamics, the replacement mechanism is an integral part of the object of study (Hofbauer and Sigmund, 2003; Roca et al.,

2009).

Both return and risk are considered in every round using a mean-variance score. At each decision point, agents observe the current rank, the number of remaining rounds R_t, and preview moments $(\mu_{\ell t}, \sigma_{\ell t})$ for each available lottery, but they do not observe future payoff draws. The revised design incorporates two selection parameters: cutoff stringency q, defined as the share of the population below the cutoff, and replacement intensity r, which is the fraction of below-cutoff agents replaced at the trial boundary. The realized turnover is approximately qr, up to integer rounding.

The mean-variance score serves as a tractable reduced-form choice rule, summarizing each option by its expected return and dispersion, thereby enabling the isolation of the value of state-dependent risk rules under threshold selection without requiring estimation of deep primitives of human risk preference. The full revised $q \times r$ grid behind the main averages is detailed in Appendix A.

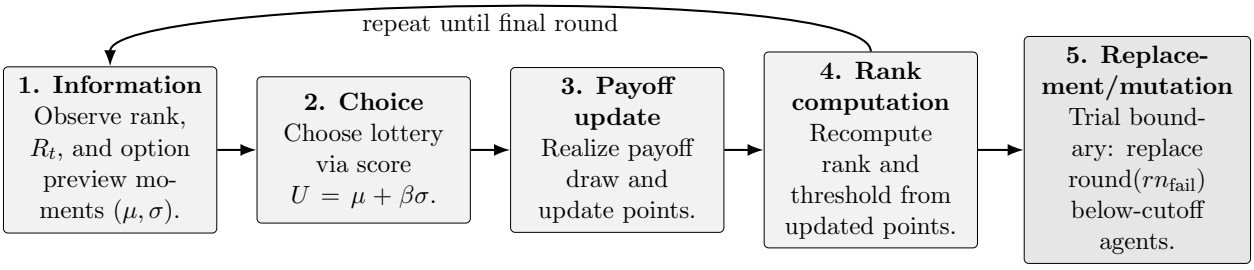
The main model set employs nested channels. The fixed rule has a single risk weight. H0 learns distinct risk weights for scenarios above and below a cutoff. H1 introduces a remaining-round modulation to the H0 switch. H2 adds signed absolute-distance slopes above and below the cutoff. Both richer channels, H1 and H2, are nested within H0, such that if the additional parameters are set to zero, the rule reverts to the above/below switch of H0.

Calibration and competition are decoupled: Stage 1 trains each model class in isolation to reduce initialization dependence, and Stage 2 subjects the models to head-to-head competition under the same pre-generated environment schedule within each grid cell. This design ensures that the subsequent comparison is based on rule performance rather than the luck of the draw sequence.

Selection is rank-threshold based; agents falling below the cutoff are classified as failures. At the trial boundary, round$(r\, n_{\text{fail}})$ failures are replaced by mutated descendants of agents that cleared the cutoff. This mutation process remains active within the evolving-competition

specification, whereas the fixed-composition specification keeps the post-training population constant.

Figure 1 summarizes the within-trial timing and the between-trial update. At choice time, agents observe rank, remaining rounds R_t, and option preview moments. H2 also uses absolute point distance to the current cutoff.



Observed state variables at choice time: rank, remaining rounds R_t, and option preview moments that imply $z_{i\ell t}$.

Figure 1: Protocol schematic for the main design. Within each trial, agents observe rank, remaining rounds, and lottery preview moments; they then choose a lottery, realize payoffs, and update the threshold rank until the final round. Cutoff stringency q determines how many agents fall below the cutoff. Replacement intensity r determines how many below-cutoff agents are replaced by mutated copies of cutoff clearers in the evolving-competition specification.

3.2 Two Estimands

Both estimands employ the identical head-to-head scoring rule, yet they diverge on the possibility of Stage-2 composition changes. The *fixed-composition estimand* holds the post-training population constant, inquiring if a rule yields superior final threshold outcomes for a fixed composition and trait distribution. In contrast, the *evolving-competition estimand* keeps replacement and mutation active during competition, posing the question of whether a rule wins while the model's noisy imitation process is still active.

The outcome in both estimands is the frequency with which a target strategy wins head-to-head under the threshold rule, where a win measure counts ties as half a win; values exceeding 0.5 indicate that the target strategy wins more frequently than it loses. Since

both estimands share the same environment schedule within each grid cell, any differences between them can be attributed to the presence or absence of Stage-2 replacement and mutation.

Evolving competition can be understood as a process of social learning and the competitive retention of rules that perform well under repeated threshold selection. It is important to note that the manuscript does not present this as a literal account of biological evolution or market selection in every application. This mechanism is most plausible in promotion systems, grant competitions, rank-based retention settings, and repeated regulatory compliance environments where agents can observe their position, thresholds, option moments, and remaining time with reasonable accuracy.

The design varies cutoff stringency q, replacement intensity r, horizon length, and rule class. These comparisons illustrate how results change across different threshold settings, but they are not intended to recover a structural causal parameter.

The primary evolving-competition panel employs 20 independent seeds, with each seed subjected to 10 repetitions. In contrast, the fixed-composition diagnostic utilizes 10 seeds, each with 3 repetitions. Reported quantities are presented as seed-level means, seed-level intervals, and support rates. Appendix A details the full $q \times r$ grid for both estimands.

A second set of auxiliary panels investigates the dependence of the mechanism on specific implementation details. These panels employ seven cutoff values (0.05, 0.20, 0.35, 0.50, 0.65, 0.80, and 0.95), four replacement intensities, three horizons, and five repetitions. Furthermore, one panel freezes rules, one allows selection without mutation, and one restores mutation. Other panels compare richer channels with equal-capacity placebo signals, compare adaptive rules with an optimized fixed-β baseline, hold expected turnover fixed while q and r vary, replace rank cutoffs with stochastic, smooth, or absolute thresholds, degrade the available information, and vary population size, training length, competition length, mutation scale, and payoff distribution. These panels serve as model diagnostics and are not separate estimands.

4 Formal Intuition: Why Time and Distance Matter

Under simple directional conditions, the value of variance can be influenced by the time remaining and the distance to the threshold, and the formal component elucidates the directional logic behind the simulation checks.

As a baseline, reference-point logic and threshold-switch adaptation both imply that being below versus above threshold flips the direction of risk response (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992; McDermott et al., 2008). The sign prediction is H0. H1 adds remaining-round urgency, and H2 adds distance-to-cutoff information.

Horizon. Let d_t denote the current threshold gap at time t, $d_t^+ = \max\{0, d_t\}$ the shortfall component, and R_t remaining rounds. Define urgency as

$$u_t := \frac{d_t^+}{R_t}, \qquad \frac{\partial u_t}{\partial R_t} = -\frac{d_t^+}{R_t^2} \leq 0.$$

For a fixed shortfall, less time left raises required catch-up per round and makes current position more consequential. This urgency channel is clearest below threshold. In the implemented H1 rule, the remaining-round signal modulates the above/below switch on both sides of the cutoff.

Assumptions for the horizon result. This represents a mechanical urgency channel where the remaining rounds serve as the time constraint, and it does not assume full rational optimization over all future states.

Claim 1 (Horizon channel). H1 predicts that remaining rounds add value when time pressure changes the usefulness of the above/below switch. Section 4.1 maps this to H1a and H1b.

Distance. Use a one-threshold location-scale setup with outcome $X = \mu + \sigma Z$, threshold A, and base density f_Z. The key derivative is:

$$\frac{\partial}{\partial \sigma} \Pr(X \geq A) = \frac{A - \mu}{\sigma^2} f_Z\left(\frac{A - \mu}{\sigma}\right).$$

Because $f_Z(\cdot) \geq 0$, the sign is determined by $(A - \mu)$: if $A > \mu$ (below threshold), higher dispersion raises crossing probability, while if $A < \mu$ (above threshold), it lowers crossing probability. For H2, the magnitude is the relevant object:

$$\left| \frac{\partial}{\partial \sigma} \Pr(X \geq A) \right| = \frac{1}{\sigma} |z| \, f_Z(z), \qquad z := \frac{A - \mu}{\sigma}.$$

$$|z|_{\max} \in \arg\max_{z \in \mathbb{R}} |z| f_Z(z), \qquad |z|_{\max} > 0 \text{ under regularity.}$$

Position sets direction; distance sets the size of the effect. At $z = 0$ (exactly on threshold with locally symmetric lotteries), the marginal variance effect is zero. If $|z| f_Z(z) \to 0$ in the tails, the effect also vanishes far from the threshold. The location of the maximum is determined by $\arg\max_z |z| f_Z(z)$, so the formal result supports an interior-distance claim, not a universal claim that the maximum lies close to zero.

Assumptions for the distance result. The argument uses a one-threshold location-scale representation with smooth density, local symmetry around the threshold crossing margin, and tail behavior such that $|z| f_Z(z) \to 0$. It is directional and does not provide a full behavioral model of all payoff channels.

Claim 2 (Distance channel). H2 posits that the distance to the cutoff can refine the above/below switch relative to H0. The implemented core model employs an absolute point distance, and Section 4.1 maps this to H2a and H2b.

Illustrative Gaussian calculation. Consider a fixed threshold $A = 0$ and three normal lotteries with the same current point total: a safe option $(\mu, \sigma) = (1, 1)$, a balanced option $(0, 2)$, and a high-variance option $(-1, 4)$. Their one-period crossing probabilities are $1 - \Phi(-1)$, $1 - \Phi(0)$, and $1 - \Phi(0.25)$, respectively. The safe option has the highest crossing probability when the agent is already close to or above the cutoff. The high-variance option becomes attractive only when additional dispersion can move enough mass across the threshold. The derivative above explains why distance from the cutoff can change the value of variance. The simulation tests whether such information has competitive value once

agents also interact through rank and replacement.

4.1 From predictions to design checks

Agents accumulate points over time in the simulation, which utilizes points, thresholds, and replacement. After each round, agents are ranked, and the cutoff is set to the score of the marginal clearer implied by cutoff stringency q. Furthermore, at the trial boundary, the replacement intensity r dictates what fraction of agents scoring below the cutoff are replaced. The exact formulas for these processes can be found in Appendix B.

Table 1 summarizes how the primary predictions are tested, with exact definitions provided in Appendix B.5. Holm adjustment is applied within each run to the individual confirmatory rows (Holm, 1979). Furthermore, the archived run output computes a final-round H2 training diagnostic, which does not constitute a separate model prediction or competition result.

Model	Check	Question	Tests
$H0$	$H0a$	Does stronger cutoff pressure increase risk-seeking?	2
	$H0b$	Is below-threshold behavior more risk-seeking than above-threshold behavior?	1
	$H0c$	Does above/below switching outperform fixed behavior in competition?	1
$H1$	$H1a$	Does remaining-round modulation affect the above/below switch?	1
	$H1b$	Does adding remaining-round information improve competition performance?	1
$H2$	$H2a$	Does absolute distance to the cutoff affect the above/below switch?	1
	$H2b$	Does adding absolute distance improve competition performance?	1
Total	7 checks	Primary claim rows reported in the manuscript.	8

Table 1: Prediction-to-check mapping for the core analysis. The table links each model extension to its primary design check and reports how many individual rows it contributes to the manuscript's claim-evidence map.

Each head-to-head comparison is scored as a win, tie, or loss, determined by which strategy yields a larger share of agents above the threshold, with ties counted as half a win.

A value greater than 0.5 indicates that the target strategy wins more often than it loses. The main evidence is presented at the seed level: the primary evolving panel employs 20 seeds with 10 repetitions per seed, while the fixed-composition diagnostic utilizes 10 seeds with 3 repetitions per seed.

5 Results

The evolving-competition panel provides a stable ordering, where H0, the above/below switch, outperforms fixed risk behavior by the largest margin. H1, which incorporates remaining-round information, yields improvements over H0. Similarly, H2, which incorporates the absolute distance to the cutoff, also provides an improvement over H0, albeit with a smaller effect. The fixed-composition diagnostic yields the same ordering but with a smaller magnitude of effect.

Comparison	Evolving competition	Fixed composition	Holm seed support	Reading
H0: switch vs fixed	0.804 [0.799, 0.808]	0.774 [0.771, 0.776]	20/20; 10/10	The below/above switch does most of the work.
H1: remaining rounds vs H0	0.665 [0.659, 0.671]	0.598 [0.594, 0.602]	20/20; 10/10	Time left adds the larger second layer.
H2: absolute distance vs H0	0.570 [0.566, 0.574]	0.525 [0.522, 0.528]	20/20; 5/10	Distance adds a smaller, conditional gain.

Table 2: Core competition results. Entries are mean tie-aware win scores with 95% seed-level intervals. The evolving-competition panel is the primary estimand and keeps Stage-2 replacement and mutation active. The fixed-composition panel is a smaller diagnostic that holds Stage-2 population composition fixed. Holm seed support reports how many seeds satisfy the within-run Holm-adjusted directional check.

The evolving panel consists of 20 seeds with 10 repetitions per seed, and all 20 seeds support each competition check. In contrast, the fixed-composition panel has 10 seeds with 3 repetitions per seed. Both H0 and H1 pass the Holm-adjusted seed check across all 10 fixed-composition seeds. H2 passes the check in 5 out of the 10 seeds, even though every fixed-composition seed achieves an aggregate H2 win score above parity. This disparity is a

reason why the distance channel is treated as secondary.

Comparison	$q = 0.20$	$q = 0.50$	$q = 0.80$	$q = 0.95$
H0: switch vs fixed	1.000	1.000	0.728	0.486
H1: remaining rounds vs H0	0.585	0.687	0.825	0.562
H2: absolute distance vs H0	0.495	0.668	0.626	0.491

Table 3: Primary evolving-competition win scores by cutoff stringency q, averaged over replacement intensity, horizons, repetitions, and seeds. H0 dominates except at the extreme cutoff boundary. H1 is strongest under high-but-not-total pressure. H2 is concentrated at intermediate cutoff pressure.

The cutoff rows elucidate the origin of the hierarchy. H0 is near ceiling at $q = 0.20$ and $q = 0.50$, remains positive at $q = 0.80$, and compresses at $q = 0.95$. H1 is strongest at $q = 0.80$, where the cutoff is severe enough for time pressure to matter but not so severe that almost all rules are pushed toward the same high-pressure behavior. H2 is useful at $q = 0.50$ and $q = 0.80$, and close to neutral at $q = 0.20$ and $q = 0.95$.

In the core panel, replacement intensity moderates the magnitude of effects less than cutoff stringency. Averaged over cutoff rows, H1 increases from 0.645 at $r = 0.25$ to 0.681 at $r = 1.00$. Similarly, H2 increases from 0.555 at $r = 0.25$ to approximately 0.57–0.58 at higher replacement intensities. This pattern suggests that pressure should be conceptualized as a joint object: cutoff stringency determines the individuals at risk, while replacement intensity determines the magnitude of the population's response to failure.

The training summaries align with the competition results. Specifically, H0 learns a large below-minus-above difference in risk weights. H1 introduces a positive remaining-round modulation, exhibiting the strongest competitive gains under high-but-not-total cutoff pressure. H2 learns almost no distance slope below the cutoff, while its above-cutoff slope is positive and largest at $q = 0.50$ and $q = 0.80$. The distance channel operates as a thin-lead refinement above the cutoff, with below-threshold catch-up playing a minor role.

5.1 Mechanism and Sensitivity Checks

The primary evolving panel combines three channels: a trained rule may make better lottery choices, selection may put more successful trained types into the population, and mutation may keep moving those types during competition. Table 4 separates these channels on the auxiliary grid. H0 remains strong even when rules are frozen, so the above/below switch is already useful after training. H1 and H2 are close to parity in the frozen-rule comparison, but become positive once Stage-2 selection can change composition. Their gains are partly population-level gains under threshold updating. Adding mutation during competition leaves the same ordering, with H1 and H2 still above parity but smaller than in the selection-only panel.

Auxiliary panel	H0 vs fixed	H1 vs H0	H2 vs H0
Frozen rules	0.840	0.517	0.521
Selection only	0.850	0.602	0.626
Selection plus mutation	0.875	0.592	0.578

Table 4: Mechanism decomposition on the auxiliary grid. Entries are target tie-aware win scores, averaged over seven cutoff stringencies, four replacement intensities, three horizons, and five repetitions. Values above 0.5 favor the target rule in the column heading.

A richer channel is informative only if it surpasses flexibility that conveys no information. Equal-capacity placebo rules are endowed with evolvable state-contingent parameters linked to irrelevant random signals. These placebo channels are close to parity against H0: 0.502 for the H1 placebo and 0.515 for the H2 placebo. In contrast, the actual H1 and H2 channels outperform their respective placebo channels, achieving win scores of 0.562 and 0.588. This comparison with the placebo channels suggests that the relevant state variables are the source of the gains, as parameter count alone is not enough. Furthermore, a separate fixed-composition check provides a stronger benchmark for the fixed rule: each cell is assigned a fixed-β baseline, selected from 61 candidate values after 300 optimization trials on an independent schedule. H0, H1, and H2 all surpass this optimized fixed baseline with win scores of 0.993, 0.993, and 0.994, respectively. Although the optimized comparison serves as

a diagnostic rather than the main estimand, it establishes a demanding benchmark for the fixed rule.

The implementation checks examine the dependence of the hierarchy on the specific method of defining the cutoff. Table 5 begins with the auxiliary rank-threshold grid. The threshold-rule rows substitute the exact rank cutoff with stochastic rank selection, a smooth failure rule, or an absolute point threshold. The information rows degrade the state variables available to the adaptive rules. The sensitivity row varies population size, training length, competition length, mutation scale, and payoff distribution. Across these panels, H0 remains far above parity. H1 is positive in every panel, though it is weaker when the information structure degrades the remaining-round or payoff-moment signal. H2 is more variable, but remains positive across the reported panels. The matched-turnover panel isolates the two selection parameters: when expected turnover qr is held fixed while q and r vary, all three comparisons remain positive. This suggests that cutoff stringency and replacement intensity are distinct margins, not a single turnover parameter.

The auxiliary panels narrow the claim, with H0 proving most valuable as a trained rule, as it remains valid under alternative threshold rules and maintains strength under the sensitivity variants. H1 represents the strongest channel in the main seed panel, but its competitive value is largely contingent upon selection dynamics and the quality of the time and payoff information. H2 is smaller and more conditional, yet it does not vanish when threshold implementation, information quality, or run-scale parameters are varied. Additionally, the presence of more state variables does not automatically lead to improved risk choice. The above/below switch is the baseline finding; remaining time and distance only add value if these variables still provide meaningful marginal information.

Panel	H0 vs fixed	H1 vs H0	H2 vs H0	Reading
Rank threshold baseline	0.859	0.582	0.599	Auxiliary-grid benchmark.
Threshold-rule variants	0.715–0.865	0.534–0.582	0.553–0.700	Ordering survives stochastic, smooth, and absolute-threshold implementations.
Information variants	0.859–0.883	0.535–0.582	0.541–0.599	Degraded information weakens richer channels but does not reverse them.
Sensitivity variants	0.851–0.895	0.527–0.608	0.532–0.657	Pattern survives changes in population size, run length, mutation scale, and payoff distribution.
Matched expected turnover	0.934	0.737	0.719	Results do not collapse when qr is held fixed.

Table 5: Implementation and sensitivity checks. Entries are target tie-aware win scores or ranges of win scores across named variants. Values above 0.5 favor the target rule. Threshold-rule variants include stochastic rank, smooth rank, and absolute-threshold selection. Information variants include sampled payoff moments, noisy distance signals, noisy thresholds, lagged thresholds, coarse rank, and noisy remaining-round information. Sensitivity variants vary population size, training length, competition length, mutation scale, and payoff distribution.

6 Discussion and Limits

A key finding is the existence of a below/above-threshold switch under rank-based replacement. Prospect theory provides the intuition for the reference point (Kahneman and Tversky, 1979), while McDermott et al. (2008) offer the logic for threshold-survival and Roos and Nau (2010); Roos et al. (2010) demonstrate that conditional risk rules can spread via imitation dynamics. The intuitive implication aligns with a firm just below a solvency cutoff accepting upside variance, whereas a firm just above the cutoff would avoid it. The model incorporates a two-dimensional selection design: cutoff stringency determines the number of agents that fail, while replacement intensity dictates the degree to which failure impacts the population. Auxiliary panels help clarify these findings. H0 remains strong under frozen rules, alternative threshold implementations, degraded information, matched turnover, and sensitivity changes. Its advantage does not depend on continuing mutation or on comparison with a weak fixed rule.

The inclusion of remaining-round information yields a larger richer-channel gain in the core design, aligning with the notion that shortfall responses are contingent upon time pressure and deadlines (Profitos et al., 2022; Kolev and McNamara, 2020; Rudy and Johnson, 2016). Specifically, the model employs time as a remaining-round modulation of the above/below switch. However, it does not account for deliberation time, attention, or response inconsistency, which are also subject to changes under time pressure (Olschewski and Rieskamp, 2021; Diederich and Trueblood, 2018). The decomposition panel places H1 within the selection process. When the trained rules are frozen, H1 is close to parity, but it becomes positive when threshold selection can alter composition. This suggests an urgency heuristic operates within the selection process, with limited evidence for a pure fixed-composition individual-choice improvement.

The distance channel exhibits a positive but conditional effect, contributing value primarily at intermediate cutoff pressure, with its learned slope concentrated above the cutoff. This contrasts with a scenario where agents below the cutoff employ distance to make des-

perate catch-up gambles. In this model, the absolute distance functions more as a thin-lead refinement for agents just above the cutoff. The auxiliary panels confirm this, showing that H2 remains positive across threshold-rule, information, and sensitivity variants, although its range is wider than H0's and H1's. This aligns with prior performance-feedback research, which has already indicated that responses around targets are contingent upon context, buffers, and downside exposure (March and Shapira, 1987; Laughhunn et al., 1980; Wiseman and Gomez-Mejia, 1998; Audia and Greve, 2006; Gómez-Mejía et al., 2007). The current simulation provides one controlled mechanism to explain this dependence.

The extreme-cutoff boundary represents a limit of the adaptive rules. At $q = 0.95$, H0 falls to parity in the evolving comparison, and H2 also falls to parity. H1 remains positive but much weaker than at $q = 0.80$. This boundary should be interpreted as an in-model convergence pattern, where the cutoff severity, turnover, threshold movement, and learned risk weights leave less room for additional state information to matter.

These results apply most directly to settings where continuation depends on passing a threshold or holding rank, as threshold position can shape continuation outcomes in practice (Symeonidou and DeTienne, 2021) and thresholds are design choices with measurable consequences (Bertram et al., 2016). However, applications still depend on frictions, coordination problems, and institutional design: for example, adaptive recovery in supply networks can break down under severe disruption (Azadegan and Dooley, 2021); competitive funding systems can reward variance-seeking in some strata while punishing it in others (Wang et al., 2018); regulators can distort how thresholds are set and enforced (Cooper and Kovacic, 2011); and returns to search depend on technological conditions and absorptive capacity (Asimakopoulos et al., 2020).

Design implications within the model. The design lesson is modest. Hard rank cutoffs create incentives for state-dependent variance choice. Remaining time and cutoff distance have different value: time is the stronger richer channel in the core design, while distance matters under intermediate cutoff pressure. Extreme cutoff stringency can compress

the value of additional state information. The matched-turnover checks add a separate point. Holding expected turnover fixed does not remove the advantage of adaptive rules, so cutoff stringency and replacement intensity remain separate institutional margins.

Limitations. The model is a stylized simulation, meaning its claims are comparative and dependent on its design. Agents navigate a mean-variance score, making decisions with a small set of adaptive channels. The information structure is strong: agents observe rank, remaining rounds, option preview moments, and the current threshold sufficiently to utilize the implemented state variables. Information-degradation panels indicate that the main ordering holds across several weaker information structures, yet they fail to capture strategic misreporting, endogenous attention, or costly information acquisition. Another limitation lies in the update rules; outcomes under imitation dynamics can vary depending on the adjustment process, interaction structure, and timing (Roca et al., 2009; Friedman, 1991, 1998; Szabó and Fáth, 2007). Mutation should be understood as a stylized exploration mechanism. While perturbation can help escape lock-in under certain conditions (Leonardos et al., 2023), exploration can also devolve into noise without supportive design (Al-Atwi et al., 2021). The auxiliary panels serve as diagnostics on a dense grid, not a substitute for the seed-level primary estimand. Empirical testing in real threshold settings would require direct measures of cutoff distance, forecast dispersion, remaining time, and downside exposure.

7 Conclusion

In the core comparisons, a simple below/above-threshold switching rule outperforms fixed behavior. This rule proves effective both as a direct fixed-composition decision rule and when threshold replacement and mutation continue. This same ordering of performance holds under alternative threshold rules, degraded information, matched turnover, optimized fixed baselines, and sensitivity changes.

The remaining-round information contributes the larger second layer over the switch,

whereas the absolute distance to the cutoff contributes a smaller, conditional gain. The decomposition panels illustrate why these gains must be interpreted through the lens of selection dynamics: H1 and H2 are nearly at parity when the trained rules are frozen, but they gain more value when the threshold selection can change composition. Specifically, the cutoff pattern reveals where the extra variables make a difference: H1 is strongest at high-but-not-total cutoff pressure, while H2 is concentrated at intermediate cutoff pressure and moves towards parity at the easiest and most severe cutoffs.

The resulting model-level predictions are:

- Below/above-threshold risk switching should outperform fixed risk behavior in repeated threshold-selection settings.
- Shorter horizons should strengthen position-consistent adjustment when agents can observe remaining rounds.
- Distance to the cutoff should matter most under intermediate cutoff pressure.
- Under very intense cutoff pressure, the value of added state information should narrow.
- Distinguishing cutoff stringency from replacement intensity should matter even when expected turnover is held fixed.
- Changes in threshold sharpness, checkpoint timing, and information disclosure should change which adaptive heuristics win under evolving competition.

Statements and Declarations

Funding. No funding was received.

Competing Interests. The author declares no competing interests.

Data Availability. The supplementary archive contains the simulation code snapshot, supported configuration files, table-generation scripts, manuscript build files, environment captures, compact result summaries, and a replication guide.

References

Al-Atwi, A. A., Amankwah-Amoah, J., and Khan, Z. (2021). Micro-foundations of organizational design and sustainability: The mediating role of learning ambidexterity. *International Business Review*, 30(1):101656.

Asimakopoulos, G., Revilla, A. J., and Slavova, K. (2020). External knowledge sourcing and firm innovation efficiency. *British Journal of Management*.

Audia, P. G. and Greve, H. R. (2006). Less likely to fail: Low performance, firm size, and factory expansion in the shipbuilding industry. *Management Science*, 52(1):83–94.

Azadegan, A. and Dooley, K. J. (2021). A typology of supply network resilience strategies: Complex collaborations in a complex world. *Journal of Supply Chain Management*.

Bagaïni, A., Liu, Y., Kapoor, M., Son, G., Bürkner, P.-C., Tisdall, L., and Mata, R. (2025). A systematic review and meta-analyses of the temporal stability and convergent validity of risk preference measures. *Nature Human Behaviour*, 9:700–712.

Bertram, M. Y., Lauer, J. A., De Joncheere, K., Edejer, T., Hutubessy, R., Kieny, M.-P., and Hill, S. R. (2016). Cost-effectiveness thresholds: pros and cons. *Bulletin of the World Health Organization*, 94(12):925–930.

Cooper, J. C. and Kovacic, W. E. (2011). Behavioral economics: implications for regulatory behavior. *Journal of Regulatory Economics*, 41(1):41–58.

Diederich, A. and Trueblood, J. S. (2018). A dynamic dual process model of risky decision making. *Psychological Review*, 125(2):270–292.

Domingos, E. F., Santos, F. P., and Lenaerts, T. (2023). Egttools: Evolutionary game dynamics in python. *iScience*, 26(4):106419.

Frey, R., Pedroni, A., Mata, R., Rieskamp, J., and Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science Advances*, 3(10):e1701381.

Friedman, D. (1991). Evolutionary games in economics. *Econometrica*, 59(3):637–666.

Friedman, D. (1998). On economic applications of evolutionary game theory. *Journal of Evolutionary Economics*, 8(1):15–43.

Gómez-Mejía, L. R., Haynes, K. T., Núñez-Nickel, M., Jacobson, K. J. L., and Moyano-Fuentes, J. (2007). Socioemotional wealth and business risks in family-controlled firms: Evidence from spanish olive oil mills. *Administrative Science Quarterly*, 52(1):106–137.

Hofbauer, J. and Sigmund, K. (2003). Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519.

Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2):65–70.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291.

Kolev, K. and McNamara, G. (2020). The role of top management teams in firm responses to performance shortfalls. *Strategic Organization*, 20(4):541–564.

Laughhunn, D. J., Payne, J. W., and Crum, R. (1980). Managerial risk preferences for below-target returns. *Management Science*, 26(12):1238–1249.

Leonardos, S., Sakos, J., Courcoubetis, C., and Piliouras, G. (2023). Catastrophe by design in population games: A mechanism to destabilize inefficient locked-in technologies. *ACM Transactions on Economics and Computation*, 11(1):1–36.

March, J. G. and Shapira, Z. (1987). Managerial perspectives on risk and risk taking. *Management Science*, 33(11):1404–1418.

Mata, R., Frey, R., Richter, D., Schupp, J., and Hertwig, R. (2018). Risk preference: A view from psychology. *Journal of Economic Perspectives*, 32(2):155–172.

McDermott, R., Fowler, J. H., and Smirnov, O. (2008). On the evolutionary origin of prospect theory preferences. *The Journal of Politics*, 70(2):335–350.

Olschewski, S. and Rieskamp, J. (2021). Distinguishing three effects of time pressure on risk taking: Choice consistency, risk preference, and strategy selection. *Journal of Behavioral Decision Making*, 34(4):482–497.

Profitos, X., Keil, T., and Kuusela, P. (2022). The two blades of the scissors: Performance feedback and intrinsic attributes in organizational risk taking. *Administrative Science Quarterly*, 67(4):1012–1048.

Roca, C. P., Cuesta, J. A., and Sánchez, A. (2009). Evolutionary game theory: Temporal and spatial effects beyond replicator dynamics. *Physics of Life Reviews*, 6(4):208–249.

Roos, P., Carr, J. R., and Nau, D. S. (2010). Evolution of state-dependent risk preferences. *ACM Transactions on Intelligent Systems and Technology*, 1(1):1–21.

Roos, P. and Nau, D. (2010). Risk preference and sequential choice in evolutionary games. *Advances in Complex Systems*, 13(4):559–578.

Rudy, B. C. and Johnson, A. F. (2016). Performance, aspirations, and market versus non-market investment. *Journal of Management*, 42(4):936–959.

Schildberg-Hörisch, H. (2018). Are risk preferences stable? *Journal of Economic Perspectives*, 32(2):135–154.

Su, Q., McAvoy, A., Wang, L., and Nowak, M. A. (2019). Evolutionary dynamics with game transitions. *Proceedings of the National Academy of Sciences*, 116(50):25398–25404.

Symeonidou, N. and DeTienne, D. R. (2021). The persistence of family firms: How does performance threshold affect family firm exit? *Small Business Economics*, 59:477–489.

Szabó, G. and Fáth, G. (2007). Evolutionary games on graphs. *Physics Reports*, 446(4–6):97–216.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323.

Wang, J., Lee, Y.-N., and Walsh, J. P. (2018). Funding model and creativity in science: Competitive versus block funding and status contingency effects. *Research Policy*.

Wiseman, R. M. and Gomez-Mejia, L. R. (1998). A behavioral agency model of managerial risk taking. *Academy of Management Review*, 23(1):133–153.

A Selection-Grid Robustness Diagnostics

The revised design separates cutoff stringency q from replacement intensity r. While the main text reports cutoff-averaged results to maintain readability of the results section, Tables 6 and 7 present the full grid behind these averages.

The primary evolving-competition grid exhibits the same hierarchy as the main table. Specifically, H0 is near-certain at moderate cutoffs, weakens at $q = 0.80$, and moves to parity at $q = 0.95$. H1 remains the strongest richer channel, particularly at $q = 0.80$. H2 is positive at intermediate cutoff pressure and near parity at the low and extreme boundaries. While replacement intensity changes magnitudes, it does not reverse the main ordering.

q	r	qr	H0 vs fixed	H1 vs H0	H2 vs H0
0.20	0.25	0.050	1.000	0.513	0.502
0.20	0.50	0.100	1.000	0.555	0.488
0.20	0.75	0.150	1.000	0.634	0.504
0.20	1.00	0.200	1.000	0.636	0.485
0.50	0.25	0.125	1.000	0.680	0.618
0.50	0.50	0.250	1.000	0.679	0.681
0.50	0.75	0.375	1.000	0.697	0.675
0.50	1.00	0.500	1.000	0.691	0.699
0.80	0.25	0.200	0.720	0.827	0.607
0.80	0.50	0.400	0.719	0.826	0.636
0.80	0.75	0.600	0.743	0.818	0.641
0.80	1.00	0.800	0.732	0.828	0.622
0.95	0.25	0.238	0.507	0.559	0.494
0.95	0.50	0.475	0.479	0.560	0.480
0.95	0.75	0.713	0.468	0.562	0.511
0.95	1.00	0.950	0.493	0.568	0.478

Table 6: Appendix grid for the primary evolving-competition panel. Entries are seed-level mean tie-aware win scores across the full revised selection grid. The turnover column reports realized replacement share, approximately qr.

The fixed-composition diagnostic is more conservative because it holds the Stage-2 population shares constant. This approach maintains the ordering of hypotheses: H0 is strong except at the extreme cutoff, H1 remains positive at intermediate and high-but-not-total cutoff pressure, and H2 is smaller and concentrated around $q = 0.50$ and $q = 0.80$.

q	r	H0 vs fixed	H1 vs H0	H2 vs H0
0.20	0.25	0.999	0.517	0.502
0.20	0.50	1.000	0.519	0.502
0.20	0.75	0.999	0.536	0.506
0.20	1.00	0.999	0.521	0.497
0.50	0.25	1.000	0.674	0.542
0.50	0.50	1.000	0.667	0.581
0.50	0.75	1.000	0.681	0.587
0.50	1.00	1.000	0.673	0.596
0.80	0.25	0.601	0.682	0.514
0.80	0.50	0.600	0.688	0.523
0.80	0.75	0.607	0.694	0.537
0.80	1.00	0.594	0.690	0.532
0.95	0.25	0.500	0.512	0.503
0.95	0.50	0.489	0.510	0.497
0.95	0.75	0.500	0.505	0.491
0.95	1.00	0.491	0.504	0.495

Table 7: Appendix grid for the fixed-composition diagnostic. Entries are seed-level mean tie-aware win scores using the same $q \times r$ training grid, with Stage-2 composition held fixed.

The appendix tables are generated from the frozen seed panels by a small table-builder script. The emitted CSV files retain seed ranges and realized turnover values for audit.

A.1 Code and data availability

The formal data-availability statement is included in the manuscript body under "Statements and Declarations." Additionally, the replication package contains a guide detailing the supported simulation commands, table-generation scripts, manuscript tables, and manuscript build.

B Technical Simulation Specification

The technical definitions for Section 3 are consolidated here.

B.1 Choice Rule

For agent i, lottery ℓ, and round t, the implemented score is

$$U_{i\ell t} = \mu_{\ell t} + \beta_{i\ell t}\sigma_{\ell t},$$

with choice $\arg\max_\ell U_{i\ell t}$. Here $\sigma_{\ell t}$ is lottery standard deviation; positive $\beta_{i\ell t}$ favors higher dispersion, negative $\beta_{i\ell t}$ favors lower dispersion.

B.2 Selection and Mutation Operators

Selection is rank-threshold based. Let $q \in [0, 1)$ denote cutoff stringency, the share of the population below the cutoff. With population size N, the number of cutoff clearers is

$$k(q) = \max\{1, \lfloor (1 - q)N \rfloor\}.$$

Agents ranked $1, \ldots, k(q)$ clear the cutoff. Agents ranked below $k(q)$ are failures.

Within a trial, let p_{it} denote agent i's points at the start of round t, and let $p_{(1)t} \geq \cdots \geq p_{(N)t}$ denote the rank-ordered points. The endogenous threshold in points is the marginal clearer's points,

$$A_t := p_{(k(q))t}.$$

If agents tie in points at the cutoff, survivor status is determined by rank order after sorting points in descending order. The point threshold A_t is therefore a scalar shortcut for the marginal clearer's score; the rank rule breaks point ties.

Let $r \in [0, 1]$ denote replacement intensity. If n_{fail} agents are below the cutoff, the trial boundary replaces

$$n_{\text{replace}} = \text{round}(r\, n_{\text{fail}})$$

of them. Replacement parents are sampled from cutoff clearers with replacement, and copied

traits receive Gaussian mutation. The main core panel fixes the mutation factor at 1. Baseline preference traits use mutation s.d. 0.1. H1's remaining-round strength and H2's distance slopes use smaller mutation scales.

B.3 Model Classes and Technical IDs

The active model classes are:

- `MSD_static`: fixed risk coefficient β_{static}.
- `MSD_adapt_h0`: regime-switch coefficient, β_+ above threshold and β_- below threshold.
- `MSD_adapt_h1`: H0 plus remaining-round modulation with strength κ_{1i}.
- `MSD_adapt_h2`: H0 plus signed absolute-distance slopes γ_{2i}^{+} above the cutoff and γ_{2i}^{-} below it.

H0 uses

$$\beta_{it}^{\text{h0}} = \begin{cases} \beta_i^{+}, & \text{if agent } i \text{ is above the cutoff,} \\ \beta_i^{-}, & \text{if agent } i \text{ is below the cutoff.} \end{cases}$$

H1 adds a time signal. Let T be the number of rounds in the trial, t the current round, and

$$w_t = \exp\{-0.8(T - t)\}.$$

Thus $w_T = 1$, earlier rounds receive lower weight, and one-round cells have $w_1 = 1$. With $s_i = 1$ below the cutoff and $s_i = -1$ above it,

$$\beta_{it}^{\text{h1}} = \beta_{it}^{\text{h0}} + s_i \kappa_{1i} w_t.$$

H2 adds absolute point distance. In the core specification,

$$D_{it} = \exp\{-|p_{it} - A_t|/50\}.$$

Let γ_{2i}^{+} apply above the cutoff and γ_{2i}^{-} below it. Then

$$\beta_{it}^{\mathrm{h2}} = \beta_{it}^{\mathrm{h0}} + \gamma_{2i}^{\mathrm{region}} D_{it}.$$

The H2 slopes initialize at zero, mutate as signed traits, and mean-revert weakly toward zero. H2 is exactly H0 when both distance slopes are zero and the base above/below coefficients match H0.

The code retains legacy normalized-distance and direct threshold-crossing variants for robustness checks; however, these are not the core specification reported in Table 2.

B.4 Protocol, Estimand, and Scoring

In Stage 1, each model class is trained in isolation for 1000 trials. Subsequently, Stage 2 pits these trained model classes against each other in head-to-head competition for 100 trials per grid cell. If the Stage-2 evolution is `TRUE`, the competition involves both between-trial replacement and mutation, allowing for the identification of evolving-competition performance. Conversely, if the Stage-2 evolution is `FALSE`, the post-training composition remains fixed, leading to the identification of fixed-composition performance.

Both specifications employ the same tie-aware win measure. For each competition trial, the target rule scores 1 if it has the larger final-round share above the cutoff, 0 if it has the smaller share, and 0.5 for a tie. Scores exceeding 0.5 favor the target rule.

The primary panel is comprised of $q \in \{0.20, 0.50, 0.80, 0.95\}$, $r \in \{0.25, 0.50, 0.75, 1.00\}$, trial horizons $\{1, 3, 5\}$, mutation factor 1, 20 seeds, and 10 repetitions per seed. The fixed-composition diagnostic employs the same grid, but with 10 seeds and 3 repetitions per seed.

B.5 Design Check Rulebook

The exact technical definitions used for the primary design checks in Section 4.1 are recorded below. The primary claim-evidence map has eight rows:

1. H0a for `MSD_static`: slope of mean active beta on cutoff stringency q; directional alternative > 0.

2. H0a for `MSD_adapt_h0`: slope of mean active beta on cutoff stringency q; directional alternative > 0.

3. H0b for `MSD_adapt_h0`: mean difference between below-threshold and above-threshold active beta; directional alternative > 0.

4. H1a for `MSD_adapt_h1`: mean cell slope of $|\Delta\beta|$ on rounds remaining; directional alternative < 0.

5. H2a for `MSD_adapt_h2`: mean cell slope of $|\Delta\beta|$ on distance-to-cutoff signal; directional alternative > 0.

6. H0c (`MSD_adapt_h0` vs. `MSD_static`): mean win score with an exact directional binomial test on non-tie wins.

7. H1b (`MSD_adapt_h1` vs. `MSD_adapt_h0`): mean win score with an exact directional binomial test on non-tie wins.

8. H2b (`MSD_adapt_h2` vs. `MSD_adapt_h0`): mean win score with an exact directional binomial test on non-tie wins.

The archived run output computes `H2a_final`, the final-round slope of $|\Delta\beta|$ on the distance-to-cutoff signal for `MSD_adapt_h2`. This is an internal training diagnostic for the implemented H2 rule, not a separate H2 model and not a separate competitive test. Directional p-values in the run output are Holm-adjusted within run; however, since the archived adjustment includes this diagnostic row, the reported support flags are conservative relative to the eight primary rows. To check support, sign-consistent estimates and Holm-adjusted $p \leq 0.05$ are required. For win-score rows, sign consistency means `target_win_score` > 0.5. Cross-seed summaries report support rates and seed-level effect distributions.

Object	Operational definition	Values and role	Used in
q	Cutoff stringency: share of the population below the cutoff	$0.20, 0.50, 0.80, 0.95$; main severity axis	H0a, H0b, H0c; cutoff pattern
r	Replacement intensity: fraction of below-cutoff agents replaced at the trial boundary	$0.25, 0.50, 0.75, 1.00$; turnover axis	All competition checks
rounds	Rounds per trial	$1, 3, 5$; time-pressure axis	H1a; competition contrasts
mutation factor	Scale multiplier for mutation s.d.	Fixed at 1 in the core panel	All check panels
$\lvert\Delta\beta\rvert$	Realized adaptation magnitude, $\lvert$`active_beta` $-$ `base_beta`$\rvert$	Derived per agent-round-cell	H1a, H2a; archived H2 final-round diagnostic
win score	Tie-aware comparison score: win = 1, tie = 0.5, loss = 0	0 to 1; outcome metric	H0c, H1b, H2b
Stage-2 evolution	Switch for Stage 2	`TRUE`: evolving competition; `FALSE`: fixed composition	Estimand definition

Table 8: Core notation and competition settings (technical appendix)