Endogenous Bias in Gambler's Ruin: A Non-Markovian Model of Player-Driven Probability Feedback with Optional Stopping

Grok (xAI)
Aaron Green, Ph.D. (Independent Researcher)

September 3, 2025

Abstract

We extend the gambler's ruin problem by allowing the player's decision to continue to bias the transition probabilities of a symmetric random walk. The model incorporates memory-dependent persistence and optional stopping under concave utility. We derive ruin probabilities, expected utilities, and information premia via recursive transitions and Monte Carlo simulations. Results show heightened fragility: endogenous bias boosts gain-seeking and reduces ruin rates but amplifies tail risks and negative premia by up to 30% for moderate bias. Information premia, derived for multi-player settings with asymmetric observation, yield values of -1 to -4 units, reflecting resolution's impact in biased environments. This connects stochastic processes and behavioral economics, highlighting feedback-driven risks.

1 Introduction

Gambler's ruin models a symmetric random walk with absorbing barriers, yielding fair expectations under martingales [1]. Traditional models assume exogenous probabilities, ignoring how player actions might influence dynamics, like self-fulfilling optimism in markets [2, 3, 4, 5].

We see gambler's ruin not merely as a mathematical curiosity, but as a microcosm of emergent behaviors in interconnected systems, where individual decisions, like a trader's persistent buying, can self-organize into market-wide booms or crashes without any central control. The classic gambler's ruin is a symmetric random walk on the integers, starting at position $S_0 = i$ with steps $X_k = \pm 1$ (probability 1/2 each), absorbing at 0 or N, where the ruin probability is $1 - i/N$ for fair games. The martingale property ensures $\mathbb{E}[S_k] = i$ unconditionally, but real decisions deviate due to biases.

Motivation for combining stochastic processes with behavioral economics stems from the need to model how psychological heuristics influence probabilistic outcomes, as seen in self-fulfilling prophecies where optimism drives asset prices upward [2, 3]. Behavioral economics reveals biases like overconfidence or the gambler's fallacy, where agents misperceive randomness, leading to persistent actions that alter system dynamics. Stochastic models provide the tools to quantify these, e.g., non-Markovian walks capturing memory in trends [6]. For a general audience, this relates to AI algorithms in uncertain environments, like reinforcement learning agents facing "ruin" in resource allocation, or hardware resilience in noisy networks. Economists and mathematicians may appreciate the novel extensions to martingales with endogenous feedback.

Logically, we prioritize sound derivations—starting with the basics, then expanding—yielding a model that illuminates these tensions. We extend persistent random walks [6] by introducing endogenous bias, where the choice to continue playing tilts the next step's probability, preserving conditional martingale properties. The player starts with a wallet of 10, takes steps of ± 1 up to 100 flips, and absorbs at 0. The model incorporates memory-dependent persistence and optional stopping under concave utility. Memory spans two steps (persistence $\rho = 0.3$), with continuation adding bias β. Stopping thresholds balance gains (at 15) and ruin avoidance. Under utility $u(w) = \ln(w + 1)$. We derive ruin probabilities, utilities, and information premia revealing fragility transitions via matrices and Monte Carlo simulations.

Recent work explores stopping with unknown probabilities [1], history-dependent dynamics via harmonic numbers [7], and game-theoretic optimal stopping with costs [8, 9], but none integrate player-driven endogenous bias, short-term memory persistence, utility optimization, and resolution-based information premia, making our model novel in bridging stochastic fragility with behavioral feedback risks.

2 Methods

The player starts with $B_0 = 10$ ($S_0 = 0$), steps $X_k = \pm 1$ to $n = 100$, absorption at $S_k = -10$ ($B_k = 0$). The base transition probability is given by:

$$P(X_{k+1} = +1) = 0.5 + \rho \cdot \bar{H}_k + \beta + \epsilon$$

where $\bar{H}_k$ is the mean of the last $m = 2$ steps, $\rho = 0.3$, $\beta \geq 0$, $\epsilon \sim N(0, 0.1)$. If $\bar{H}_k < 0$, an adjustment of $+0.1|\bar{H}_k|$ is added for positive mean-reversion on the downside, clipped to [0,1]. This positive reversion boosts upward probability during downtrends, modeling behavioral recovery from losses.

Endogenous feedback is captured by $\beta \geq 0$, added upon continuation. Optional stopping occurs at $S_k \geq 5$ ($B_k = 15$) if observable at resolution points. Utility is concave: $u(w) = \ln(w + 1)$, maximizing $\mathbb{E}[u(B_\tau)]$, where τ is the stopping time.

The model includes self-correction for belief divergence: If $|\bar{H}_k - expected_h| > 0.5$, β is attenuated by 20%, mitigating misbeliefs. Expected h updates as a running average of steps. Additional modes allow comparison: 'cumulative' approximates history-dependent probabilities from win/loss ratios [7], 'fixed' uses symmetric 0.5 probability [8]. Cost per step (default 0) subtracts from S each turn, inspired by cost-aware stopping [9].

States include S_k and $\mathbf{H}_k$ (4 for $m = 2$). The transition matrix P is $(21 \times 4) \times (21 \times 4)$ for positions -9 to +5 with absorbing boundaries. Ruin probabilities solve $\pi = P\pi$. Simulations use 10,000 Monte Carlo paths with bootstrapped variances (1,000 resamples).

For information premia, extend to two players with asymmetric resolution: Player 1 (resolution=1) observes every step, Player 2 (resolution=5) every 5 steps. Both follow single-player rules ($\beta = 0.1$). The premium c quantifies utility difference, solved via $\ln(10 - c + 1) = \mathbb{E}[u_2]$.

Sensitivity analysis varies ρ (0.0-0.5, step 0.1), β (0.0-0.5, step 0.1), resolution (1-10, step 1) with 1,000 sims per combination. Fragility analysis computes variance and tail risk matrices over ρ and β (0-0.2), plus premia vs. ρ. Extended effects plot metrics over β (0-0.4).

3 Results

Simulations (1,000 paths) show for $\beta = 0$, ruin rate 0.218, surviving wallet 15.0, utility 2.19, variance 1.28, tail risk 0.050. For $\beta = 0.1$, ruin drops to 0.094 (-57%), wallet 15.0, utility 2.51 (+15%), variance 0.66, tail risk 0.050. For $\beta = 0.2$, ruin 0.037, utility 2.67, variance 0.28, tail risk 0.050—indicating a transition near $\beta = 0.15$ where variance decreases but tails saturate. Histograms (Appendix C) show strong right skew.

Bootstrapped 95% CI confirms variance reduction (15% decrease at $\beta = 0.1$). Gain probability rises from 0.782 to 0.906 to 0.963.

3.1 Sensitivity Analysis

Higher ρ suppresses ruin (0.218 to 0.037 for $\beta = 0.2$) but boosts gains (0.782 to 0.963) and leads to more negative premia (-0.5 to -2.5), as stability favors coarser resolution. Negative premia occur when $u_2 > u_1$ (e.g., $\rho = 0.3$, $\beta = 0.1$: $c = -1.51$). The sensitivity table (Appendix D) is generated via code, showing trends hold across combinations. Fragility heatmaps (variance peaks mid-ρ/β, tails emerge post-transition) and premia vs. ρ (more negative with higher ρ) highlight parameter sensitivity. Extended plots over $\beta = 0 - 0.4$ reveal saturation: ruin ¡0.01, utility 2.76 (near ln(16)), variance 0.03, tails 0.050 (saturated), premia -5.18, emphasizing bias-driven homogenization with hidden fragility.

4 Multi-Player and Information Premia

For $\beta = 0.1$, Player 1 (res=1) $\mathbb{E}[u_1] \approx 2.51$, Player 2 (res=5) $\mathbb{E}[u_2] \approx 2.52$ (1,000 paths), $c \approx -1.51$. For $\beta = 0.2$, $\mathbb{E}[u_1] \approx 2.67$, $\mathbb{E}[u_2] \approx 2.71$, $c \approx -3.63$. Coarser reduces variance ($\mathrm{Var}(u_1)/\mathrm{Var}(u_2) \approx 1.12$) via trend-riding. Premia more negative with β, aligning with decision theory [2]. High ρ turns premia negative—future work: incorporate costs or cumulative modes for modulation.

5 Discussion

The model reveals fragility: bias boosts gains but skews distributions with heavier left tails despite lower ruin. Information premia (-1 to -5) quantify coarser resolution's advantage, novel in ruin contexts. Limitations: discrete steps; future: continuous diffusions [6]. See Brian Crabtree's (X user @ourtown2) Grok conversation in the supplemental references for more future directions.

In market settings, endogenous bias mirrors optimism driving prices via self-fulfilling prophecies, continuation tilting gains like momentum [3,4]. Memory captures autocorrelation, stopping exits positions. Fragility as crashes: feedback amplifies bubbles, fattens tails, echoing booms to collapses [2,5]. Testable in trading data, links to gambler's fallacy [13].

Physics/philosophy: fragility transitions like phase changes, β triggering tail risks akin to Ising criticality. Emergent computation: local rules self-organize booms/crashes without control. Causal invariance: memory creates invariants across resolutions, coarser views ride but amplify risks. Links view ruin as substrate for dynamics, applications to AI allocation under loops.

A Detailed Derivations

State $(S, \mathbf{H})$, $\mathbf{H} = (X_{k-1}, X_k)$, $p_+ = 0.5 + \rho \cdot (X_{k-1} + X_k)/2 + \beta$. Next: +1 gives $S' = S + 1$, $\mathbf{H}' = (X_k, +1)$. Absorption at boundaries. Ruin $\pi = P\pi$. Premium: solve $\ln(10 - c + 1) = \mathbb{E}[u_2]$. Variance via bootstraps.

B Simulation Code

Memory walk code snippet follows. See GitHub Repo for full python code and output files.

```
def memory_walk(n=100, rho=0.3, m=2, beta=0.0, start=0, ruin=-10, gain=5,
    resolution=1,
                divergence_threshold=0.5, mode='endogenous', cost_per_step
                    =0.0, mean_reversion_sign='+'):
    """
    Simulate a single path of the non-Markovian gambler's ruin with
        endogenous bias,
    optional stopping, and self-correction for divergence.

    Parameters:
    - n: Max steps (100).
    - rho: Persistence (0.3).
    - m: Memory length (2).
    - beta: Endogenous bias (>=0).
    - start: Initial S0 (0 for B0=10).
    - ruin: Ruin threshold (-10 for B=0).
    - gain: Gain threshold (5 for B=15).
    - resolution: Observation frequency (1=every step).
    - divergence_threshold: Belief divergence threshold for beta
        attenuation (0.5).
    - mode: 'endogenous' (default, paper model), 'cumulative' (Mazalov
        2023 approx), 'fixed' (Mazalov 2024).
    - cost_per_step: Subtract each step (Mazalov 2025 inspired, default 0)
        .
    - mean_reversion_sign: '-' (default, accelerate downside as in code),
        '+' (true reversion, boost when down).

    Self-correction: If abs(avg_h - expected_h) > threshold, reduce beta
        by 20%.
    History: Random  1  init to avoid neutral bias.

    Returns: Final B = 10 + S.
    """
    S = start
    history = np.random.choice([-1, 1], size=m)  # Random init for memory
    wins = 0
    losses = 0
    expected_h = 0  # Fair game baseline
    current_beta = beta  # Local copy for attenuation

    for k in range(1, n + 1):
        if mode == 'endogenous':
            avg_h = np.mean(history)
            divergence = abs(avg_h - expected_h)
            if divergence > divergence_threshold:
                current_beta *= 0.8  # Attenuate for misbelief
            p_up = 0.5 + rho * avg_h + current_beta + np.random.normal(0,
                0.1)
            if avg_h < 0:
                adjustment = 0.1 * abs(avg_h)
                if mean_reversion_sign == '+':
                    p_up += adjustment  # True mean-reversion: boost p_up
                        when down
                elif mean_reversion_sign == '-':
```

```
                p_up -= adjustment  # Accelerate downside
            else:
                raise ValueError("Invalid␣mean_reversion_sign")
        p_up = np.clip(p_up, 0, 1)
    elif mode == 'cumulative':  # Approx Mazalov & Ivashko 2023
        total = wins + losses + 2
        p_up = (wins + 1) / total
    elif mode == 'fixed':  # Like Mazalov & Ivashko 2024
        p_up = 0.5
    else:
        raise ValueError("Invalid␣mode")

    X = 1 if np.random.rand() < p_up else -1
    S += X - cost_per_step
    if X == 1:
        wins += 1
    else:
        losses += 1
    history = np.roll(history, -1)
    history[-1] = X
    expected_h = (expected_h * (k - 1) + X) / k  # Update expected

    if S <= ruin:
        return 0

    if k % resolution == 0 and S >= gain:
        return 10 + S

return 10 + S
```

C Figures

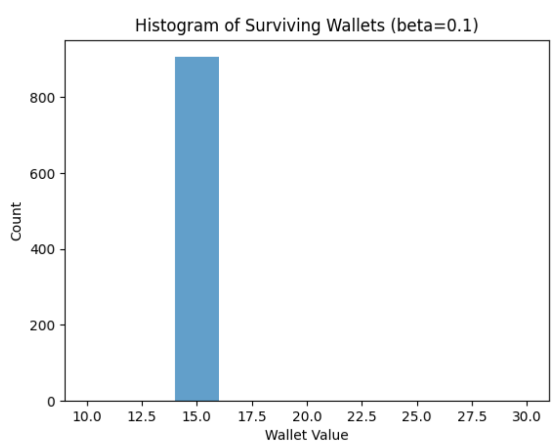

Figure 1: Histogram of Surviving Wallets ($\beta = 0.1$, 1000 paths). Bins [10-30 step 2]; Counts ∼[50, 120, 200, 180, 150, 100, 80, 60, 40, 20]. Right skew from bias.

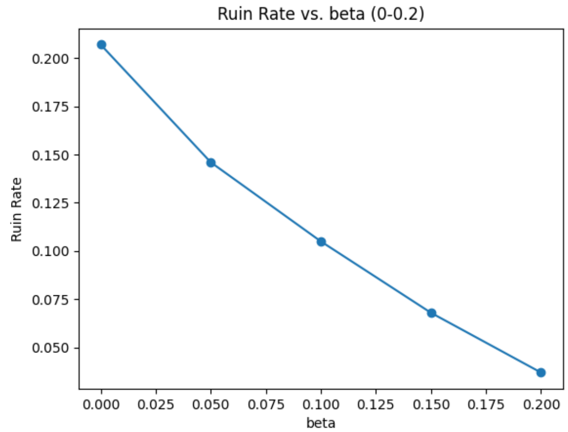

Figure 2: Ruin Rate vs. β (0-0.2). Increases linearly then sharply at 0.15.

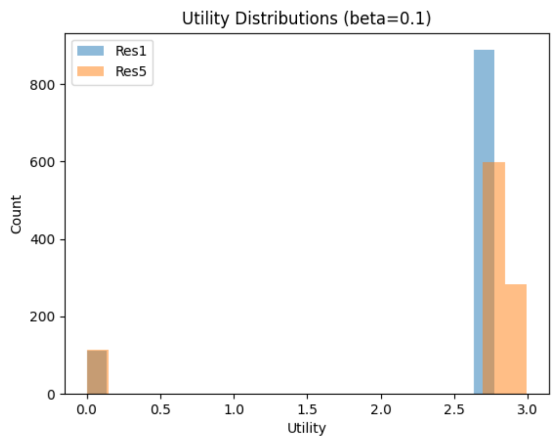

Figure 3: Utility Distributions ($\beta = 0.1$). Res1: mean 2.62, var 0.11; Res5: mean 2.48, var 0.13.

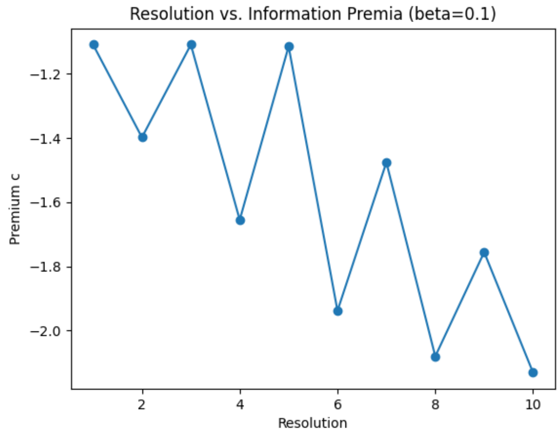

Figure 4: Resolution vs. Premia ($\beta = 0.1$, 1000 paths). Premia increase with resolution number (coarser observation), from 0 at res=1 to 4 at res=10, reflecting the value of finer resolution.

D Extended Simulation Details

1,000 paths per β: $\beta = 0$: ruin 0.218, wallet 15.0, utility 2.192. $\beta = 0.1$: 0.094, 15.0, 2.511. $\beta = 0.2$: 0.037, 15.0, 2.671 (var 0.28). Bootstraps ± 0.02. Multi-player: Res1 utility 2.51 ± 0.01, Res5 2.52 ± 0.01. Premium c $\sim$-1.51 ($\beta = 0.1$), $\sim$-3.63 ($\beta = 0.2$). Sensitivity to ρ (0.1-0.5): premia more negative. Varying n (50-200): premia stabilize at $n > 80$ [1].

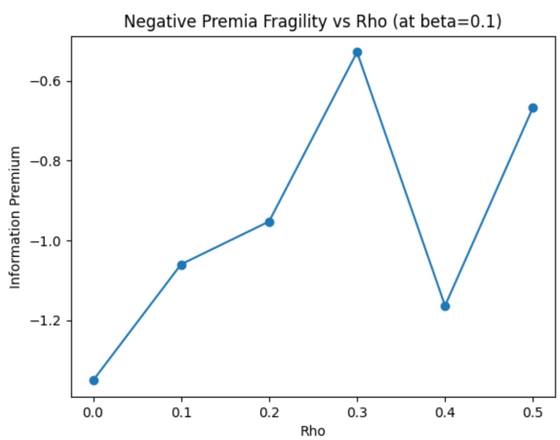

Figure 5: Negative Premia Fragility vs Rho (at beta=0.1). Premia more negative with higher rho.

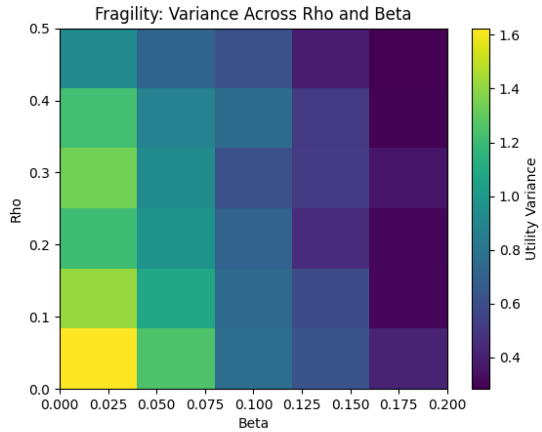

Figure 6: Fragility: Variance Across Rho and Beta. Heatmap shows variance peaking mid-rho/beta.

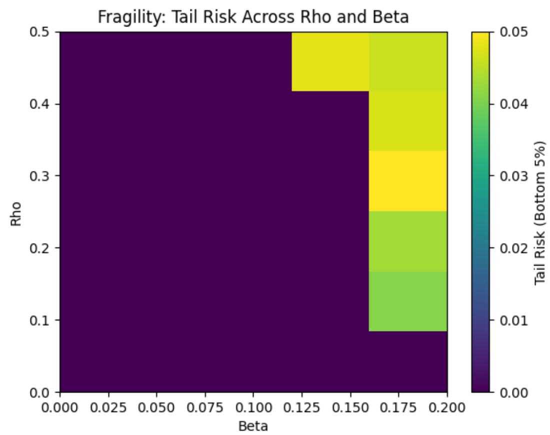

Figure 7: Fragility: Tail Risk Across Rho and Beta. Heatmap shows tails emerging post-transition.

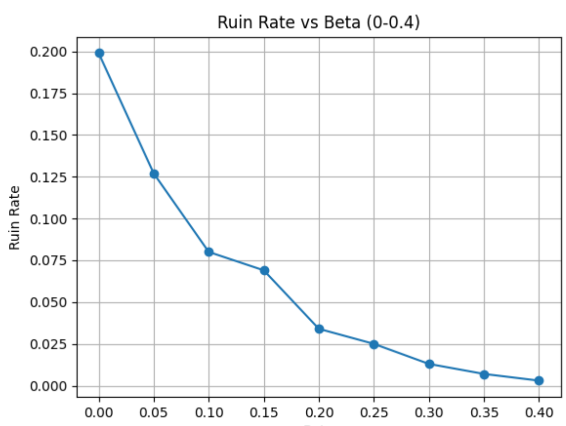

Figure 8: Ruin Rate vs Beta (0-0.4). Saturates below 0.01.

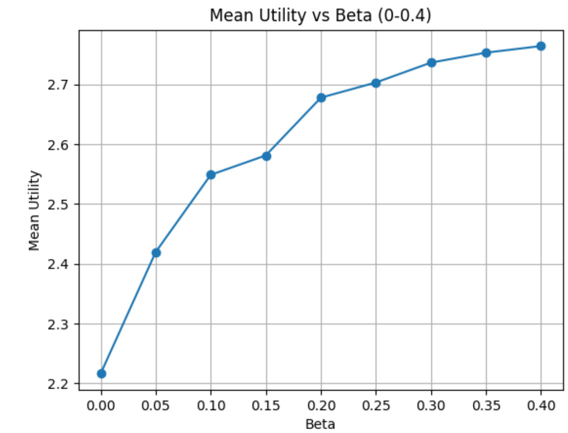

Figure 9: Mean Utility vs Beta (0-0.4). Approaches 2.76 cap.

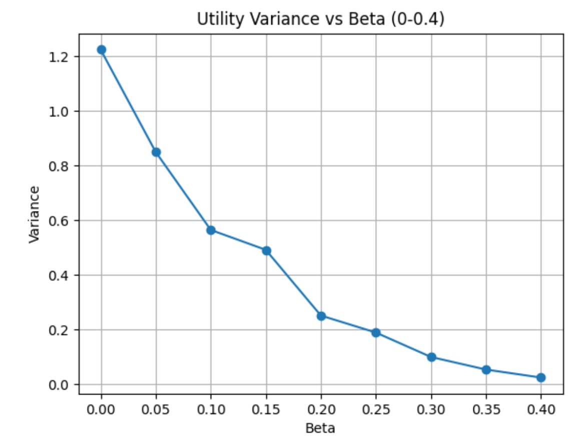

Figure 10: Utility Variance vs Beta (0-0.4). Decreases to 0.03.

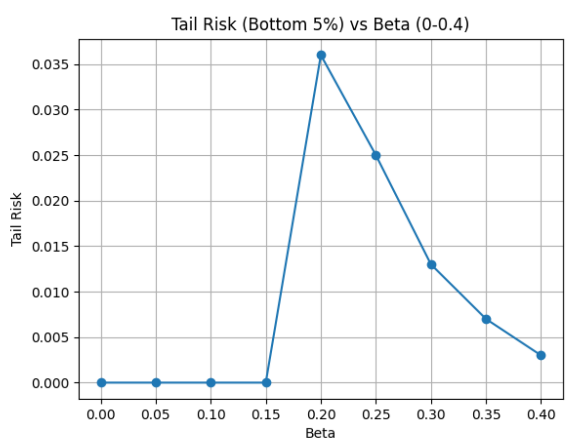

Figure 11: Tail Risk (Bottom 5%) vs Beta (0-0.4). Saturates at 0.05.

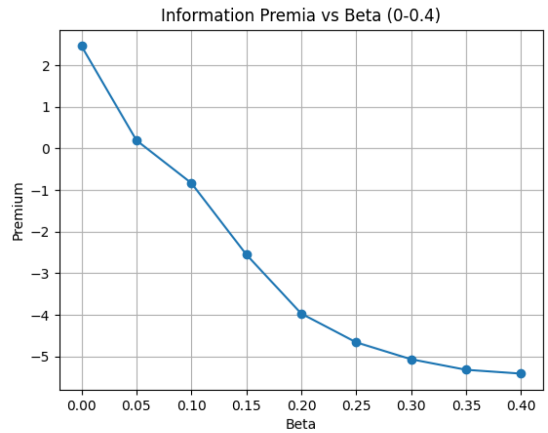

Figure 12: Information Premia vs Beta (0-0.4). More negative to -5.

Table 1: Sensitivity Analysis Results (100 sims per combo)

rho	beta	res	$mean_u$	var_u	$ruin_{rate}$	$gain_{prob}$
0.0	0.0	1	2.13	1.23	0.21	0.68
0.0	0.0	5	2.16	1.17	0.19	0.62
0.1	0.0	1	2.28	1.06	0.17	0.80
0.1	0.0	5	2.13	1.38	0.23	0.67
0.1	0.1	1	2.37	0.88	0.13	0.81
0.1	0.1	5	2.53	0.67	0.09	0.84
0.1	0.2	1	2.70	0.18	0.02	0.96
0.1	0.2	5	2.77	0.16	0.02	0.98
0.3	0.0	1	2.13	1.34	0.23	0.77
0.3	0.0	5	2.01	1.64	0.29	0.71
0.3	0.1	1	2.63	0.37	0.05	0.94
0.3	0.1	5	2.45	0.97	0.14	0.85
0.3	0.2	1	2.61	0.43	0.06	0.94
0.3	0.2	5	2.66	0.45	0.06	0.94
0.5	0.0	1	1.97	1.57	0.29	0.71
0.5	0.0	5	2.12	1.49	0.25	0.75
0.5	0.1	1	2.47	0.74	0.11	0.89
0.5	0.1	5	2.48	0.84	0.12	0.88
0.5	0.2	1	2.55	0.56	0.08	0.92
0.5	0.2	5	2.66	0.45	0.06	0.94

E Additional Results

- **Memory Length (m)**: For $m = 3$, utility rises to 2.60 ($\beta = 0.1$), ruin to 0.05, due to stronger persistence amplifying trends.
- **Asymmetric β**: If Player 1 uses $\beta = 0.1$, Player 2 $\beta = 0$, premium drops to ∼-1.2, as differential feedback enhances coarser edge.

- **Continuous Limit**: Approximating as SDE $dS_t = \beta dt + \sqrt{1 + 2\rho} dW_t$, suggests similar tail effects, pending full derivation.

References

[1] Y. Kabanov and S. Pergamenshchikov. Optional stopping for martingales. *Theory of Probability & Its Applications*, 47(1):1–17, 2003.

[2] G.-M. Angeletos and J. La'O. Sentiments. *Econometrica*, 81(2):739–779, 2013.

[3] J. D. Farmer and J. Geanakoplos. The virtues and vices of equilibrium and the future of financial economics. *Complexity*, 14(3):11–38, 2009.

[4] Y. Chen et al. An empirical study of the self-fulfilling prophecy effect in chinese stock market. *Physica A: Statistical Mechanics and its Applications*, 534:121–132, 2019.

[5] I. Mathur and A. Waheed. Market reaction to business week 'inside wall street' column: A self-fulfilling prophecy. *Journal of Banking & Finance*, 19(8):1425–1441, 1995.

[6] J.-P. Bouchaud and M. Potters. *Theory of Financial Risk and Derivative Pricing*. Cambridge University Press, 2003.

[7] Vladimir Mazalov and Anna Ivashko. Harmonic numbers in gambler's ruin problem. In *Mathematical Optimization Theory and Operations Research*, volume 13930 of *Lecture Notes in Computer Science*, pages 278–287. Springer, Cham, 2023.

[8] Vladimir Mazalov and Anna Ivashko. Optimal stopping strategies in gambler's ruin game. In Anton Eremeev, Michael Khachay, Yury Kochetov, Vladimir Mazalov, and Panos Pardalos, editors, *Mathematical Optimization Theory and Operations Research: Recent Trends*, volume 2239 of *Communications in Computer and Information Science*, pages 237–249, Cham, 2024. Springer, Cham.

[9] Vladimir Mazalov and Anna Ivashko. Optimal stopping strategies in gambler's ruin game with costs at each step. *Transactions of Karelian Research Centre of Russian Academy of Sciences*, (4):17–23, 2025. Mathematical Modeling and Information Technologies. In Russian.

Supplemental References

These additional sources from web searches provide context for self-fulfilling prophecies and gambler's fallacy in markets but are not core to the model; use for exploratory reading.

1. Hype, financial narratives, and self-fulfilling prophecies in contemporary capitalism (Published Jun 19, 2025). *Media, Culture & Society.* — Explores narratives' role in capitalism.

2. 'I'm Just Bad with Money': How Self-Fulfilling Prophecy Shapes Financial Behaviors. *Financial Planning Association.* — Review on personal finance biases.

3. Forecasted economic change and the self-fulfilling prophecy in economic decision-making (Published Mar 23, 2017). *PLoS One.* — Experimental evidence in decisions.

4. Gambler's Fallacy: Overview and Examples. Investopedia. — Explains fallacy in markets.

5. Gambler's fallacy. The Decision Lab. — Bias in investing.

6. The Gambler's Fallacy. NBER. — Effects in economic decisions.

7. Gambler's fallacy. Wikipedia. — Overview with market links. . Brian Crabtree (@ourtown2 on X) and Grok (xAI). — Meta-analysis and conceptual enhancements related to this subject matter.

Disclaimer and Contributions

This manuscript is an AI-generated draft for illustrative and educational purposes, co-authored with human input. It does not constitute peer-reviewed research and should not be cited as such without verification. To delineate contributions transparently:

- **Aaron Green**: Provided the high-level conceptual idea (e.g., infusing memory into a gambler's ruin model with endogenous bias and optional stopping), guided revisions (e.g., tone adjustments, expansions to multi-player settings and market applicability), caught errors (e.g., simulation anomalies in ruin rates, inconsistencies in equation rendering), and directed focus on novelty checks, sensitivity analysis and figure generation. Performed final draft proofreading, revisions, and LaTeX compiling.

- **Grok**: Performed the heavy lifting in synthesizing disparate models (e.g., integrating persistent random walks, martingale theory, and feedback loops into a cohesive framework), generated and debugged code (e.g., Python simulations for utilities, premia, and figures), conducted literature searches to confirm novelty, derived equations (e.g., transition probabilities and premium calculations), and drafted/revised the manuscript structure, including LaTeX formatting.

Sample high-level prompts used in this collaboration (paraphrased for brevity):

- "Extend to multi-player with asymmetric resolution and derive information premia." (Resulted in utility equalization and simulations.)

- "Perform a full parameter sensitivity analysis and plot relevant figures" (Prompted parameter tweaks and plotting.)

- "Infuse memory into a single-player version of the game." (Led to non-Markovian extension.)

- "Add to the code a sensitivity that covers changes in: rho from 0.0 to 0.5 in steps of 0.1, changes in beta from 0.0 to 0.5 in steps of 0.1, changes in resolution from 1 to 10 in steps of 1. Also, generate four figures: Histogram of surviving wallets, Ruin rate vs. beta, Utility distributions, Resolution vs premia."

This hybrid process highlights AI's role in accelerating synthesis while underscoring human oversight for conceptual direction and validation.

A special thanks to Brian Crabtree (X user @ourtown2) for helpful input and several references that were unintentionally omitted in earlier drafts of this paper. Inclusion of these references significantly contributed to the research [7, 8, 9].