

Extreme-Target Survival in an Asymmetric Gambling Walk: Exact Ruin Laws, Markov Persistence, and Noncommuting Limits

Po Hu

Independent Researcher

August 2026

Abstract

Consider a gambler with integer capital who gains two units after a win, loses one unit after a loss, and is removed at the first visit to zero. The first observation is elementary but striking: with positive drift, the probability of reaching an arbitrarily high finite target does not tend to zero. It tends to the probability of never being ruined. For independent trials with win probability p , we record the exact ruin law, identify the phase transition at $p = 1/3$, and give a closed formula at criticality whose fixed-capital asymptotic is x/m .

We then keep the one-step win probability equal to $1/2$ and the mean payoff equal to $1/2$, but let successive outcomes form a symmetric two-state Markov chain with persistence parameter θ . An explicit bounded harmonic function gives the ultimate survival probability

$$1 - \frac{1 + r + r^2}{2(1 + r)} r^{x-1}, \quad r = \frac{\sqrt{(1 - \theta)^2 + 4\theta^2} - (1 - \theta)}{2\theta}.$$

This probability decreases strictly with persistence, although every model in the family has the same stationary win rate and the same mean increment. At $\theta = 1$ the survival probability is discontinuous: its left limit is $1/4$, while its endpoint value is $1/2$. Consequently, the high-target limit and the high-persistence limit do not commute. Finally, if $(1 - \theta)x \rightarrow \lambda$, the exact transition profile is $1 - (3/4)e^{-\lambda/2}$. These formulas isolate how early absorption and temporal clustering, rather than the height of a finite target alone, govern extreme-target success.

Keywords: gambler's ruin; skip-free random walk; Markov additive process; correlated random walk; first passage; persistence; noncommuting limits.

MSC 2020: 60J10; 60G50; 91B30.

1 Introduction

Suppose that a gambler starts with one unit of capital. A fair coin is tossed at every round: a win adds two units and a loss subtracts one. Zero is absorbing. The mean increment is positive,

$$\frac{1}{2}(2) + \frac{1}{2}(-1) = \frac{1}{2},$$

but the first loss causes immediate ruin. Nevertheless, the probability of reaching a target of one million units, one billion units, or any other fixed finite amount stays bounded away from zero as the target grows. Its limit is $(3 - \sqrt{5})/2 \approx 0.381966$.

The apparent paradox comes from confusing two limits. On a path that is eventually ruined, only finitely many rounds are played and the maximum capital is finite. On a path that is never ruined, the positive drift forces the capital to infinity, so every finite target

is eventually crossed. Thus the events of reaching increasingly high targets decrease to the event of eternal survival. The target itself is not irrelevant—the waiting time still grows linearly—but it is no longer the main obstruction in the infinite-horizon limit.

The independent walk is a classical skip-free-to-the-left random walk. Such walks have a substantial literature and appear in ruin theory, queueing, and branching processes; see Brown, Peköz, and Ross [2] and Avram and Vidmar [1]. Correlated gambler’s ruin is also classical, beginning at least with Mohan [6]; subsequent treatments include Proudfoot and Lampard [9], Mukherjea and Steele [7], Lal and Bhat [5], and the recent martingale treatment of Pozdnyakov [8]. General boundary crossing for Markov random walks is treated, for example, by Fuh [4].

Accordingly, this note does not claim novelty for gambler’s ruin, skip-free recurrences, or Markov-dependent trials as general subjects. Its purpose is to give a self-contained exact analysis of the concrete asymmetric payoff $\{+2, -1\}$ and to isolate four consequences in a single model:

- (i) extreme-target probability converges to a positive survival mass;
- (ii) persistence changes that mass despite fixed stationary marginals and fixed mean payoff;
- (iii) the fully persistent endpoint is singular;
- (iv) target height and persistence produce noncommuting limits and an explicit joint scaling window.

The exact formula for this asymmetric Markov model and the resulting synthesis do not appear in the sources screened for this paper. This literature positioning is deliberately narrower than a claim of priority for the underlying methods.

2 The extreme-target principle

Let $S_0 = x \in \mathbb{Z}_{>0}$ and let every increment take one of the values $+2$ and -1 . Define

$$\tau_0 = \inf\{n \geq 1 : S_n = 0\}, \quad T_m = \inf\{n \geq 0 : S_n \geq m\}, \quad m > x.$$

The walk is stopped at τ_0 . Since downward motion is by one unit, a path cannot jump across zero without visiting it.

Proposition 2.1 (Extreme-target principle). *Assume that, on $\{\tau_0 = \infty\}$,*

$$S_n \longrightarrow +\infty \quad \text{almost surely.}$$

Then

$$\lim_{m \rightarrow \infty} \mathbb{P}_x(T_m < \tau_0) = \mathbb{P}_x(\tau_0 = \infty). \quad (2.1)$$

In particular, conditional on eternal survival, every finite target is crossed almost surely.

Proof. Put $A_m = \{T_m < \tau_0\}$. The events A_m decrease with m . If ruin occurs, the stopped path has finite length and hence a finite maximum, so it does not belong to $\cap_m A_m$. If ruin never occurs, the assumed divergence of S_n implies membership in every A_m . Therefore

$$\bigcap_{m > x} A_m = \{\tau_0 = \infty\}$$

up to a null set. Continuity of probability from above gives (2.1). $\square$

The proposition concerns probability, not speed. The following inverse law keeps that distinction explicit.

Proposition 2.2 (Time cost of a high target). *Suppose that, on eternal survival,*

$$\frac{S_n}{n} \longrightarrow \mu > 0 \quad \text{almost surely.}$$

Then, on $\{\tau_0 = \infty\}$,

$$\frac{T_m}{m} \longrightarrow \frac{1}{\mu} \quad \text{almost surely.} \quad (2.2)$$

Proof. For every $\varepsilon \in (0, \mu)$, eventually $(\mu - \varepsilon)n \leq S_n \leq (\mu + \varepsilon)n$. At the first crossing, $m \leq S_{T_m} \leq m + 1$, because the upward jump has size two. The usual inverse bounds therefore give $(\mu + \varepsilon)^{-1} \leq \liminf T_m/m \leq \limsup T_m/m \leq (\mu - \varepsilon)^{-1}$. Let $\varepsilon \downarrow 0$. $\square$

3 Independent trials and the critical point

Let the increments be independent, with

$$\mathbb{P}(X_n = 2) = p, \quad \mathbb{P}(X_n = -1) = q = 1 - p, \quad S_n = x + \sum_{j=1}^n X_j.$$

The mean increment is $3p - 1$, so $p_c = 1/3$ is the drift threshold.

Theorem 3.1 (Ultimate ruin for independent trials). *For every $x \in \mathbb{Z}_{>0}$:*

- (i) *if $p \leq 1/3$, then $\mathbb{P}_x(\tau_0 < \infty) = 1$;*
- (ii) *if $p > 1/3$, there is a unique $\rho_p \in (0, 1)$ satisfying*

$$p\rho_p^3 - \rho_p + q = 0,$$

and

$$\mathbb{P}_x(\tau_0 < \infty) = \rho_p^x, \quad \rho_p = \frac{\sqrt{4/p - 3} - 1}{2}. \quad (3.1)$$

Consequently,

$$\lim_{m \rightarrow \infty} \mathbb{P}_x(T_m < \tau_0) = \begin{cases} 0, & p \leq 1/3, \\ 1 - \rho_p^x, & p > 1/3. \end{cases} \quad (3.2)$$

Proof. Let u_x be the ultimate ruin probability. First-step analysis gives

$$u_x = pu_{x+2} + qu_{x-1}, \quad u_0 = 1.$$

The exponential ansatz $u_x = r^x$ yields

$$pr^3 - r + q = (r - 1)(pr^2 + pr - q) = 0.$$

For $p > 1/3$, the quadratic factor has the root ρ_p in $(0, 1)$ displayed in (3.1). Then $\rho_p^{S_{n \wedge \tau_0}}$ is a bounded martingale. On ruin it converges to one, while on survival the strong law gives $S_n/n \rightarrow 3p - 1 > 0$ and hence convergence to zero. Bounded convergence proves (3.1). If $p < 1/3$, the unrestricted walk drifts to $-\infty$; if $p = 1/3$, the nondegenerate zero-mean finite-variance walk oscillates. In both cases downward skip-freeness forces a visit to zero almost surely. Equation (3.2) follows from Proposition 2.1. $\square$

Corollary 3.2 (The fair-coin values). *If $p = 1/2$ and*

$$\alpha = \frac{\sqrt{5} - 1}{2},$$

then

$$\mathbb{P}_x(\tau_0 < \infty) = \alpha^x, \quad \lim_{m \rightarrow \infty} \mathbb{P}_x(T_m < \tau_0) = 1 - \alpha^x.$$

In particular, the limiting extreme-target probabilities are

$$1 - \alpha = \alpha^2 \approx 0.381966 \quad \text{from } x = 1,$$

and

$$1 - \alpha^{20} \approx 0.999933893 \quad \text{from } x = 20.$$

On survival, $T_m/m \rightarrow 2$ almost surely.

At criticality the limiting target probability is zero, but its decay is only of order $1/m$ rather than exponential.

Theorem 3.3 (Exact critical high-target law). *Let $p = 1/3$. For integers $1 \leq x < m$,*

$$\mathbb{P}_x(T_m < \tau_0) = \frac{3x(-2)^m + (-2)^x - 1}{(3m+1)(-2)^m - 1}. \quad (3.3)$$

For fixed x ,

$$\mathbb{P}_x(T_m < \tau_0) \sim \frac{x}{m} \quad (m \rightarrow \infty).$$

Proof. Write $h_x = \mathbb{P}_x(T_m < \tau_0)$. Then

$$h_x = \frac{1}{3}h_{x+2} + \frac{2}{3}h_{x-1}, \quad h_0 = 0, \quad h_m = h_{m+1} = 1.$$

The characteristic polynomial is

$$r^3 - 3r + 2 = (r-1)^2(r+2),$$

so $h_x = A + Bx + C(-2)^x$. Imposing the three boundary values gives (3.3). Dividing numerator and denominator by $(-2)^m$ gives the asymptotic formula. $\square$

For example, from one unit the exact probabilities of crossing targets 4, 5, 10 are respectively 5/23, 11/57, and 341/3527. These values also illustrate why the asymptotic $1/m$ should not be used as an exact finite-target formula.

4 Markov-dependent outcomes with fixed marginals

We now preserve the fair one-step distribution but introduce serial dependence. Let $J_n \in \{W, L\}$ be a stationary Markov chain with transition matrix

$$P_\theta = \begin{pmatrix} \theta & 1-\theta \\ 1-\theta & \theta \end{pmatrix}, \quad 0 < \theta < 1, \quad (4.1)$$

and initial distribution $(1/2, 1/2)$. A win contributes +2 and a loss contributes -1. The parameter θ is the probability that the next outcome repeats the current outcome.

Every θ has the same stationary win probability $1/2$ and therefore the same mean increment $1/2$. What changes is temporal clustering. If $Y_n = 2\mathbf{1}_{\{J_n=W\}} - \mathbf{1}_{\{J_n=L\}}$, then

$$\text{Cov}(Y_0, Y_k) = \frac{9}{4}(2\theta - 1)^k, \quad (4.2)$$

and, for $0 < \theta < 1$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \text{Var} \left(\sum_{j=1}^n Y_j \right) = \frac{9}{4} \frac{\theta}{1 - \theta}. \quad (4.3)$$

Thus persistence increases the long-run variance even though the marginal law does not change.

Theorem 4.1 (Exact Markov ruin law). *For $0 < \theta < 1$, let $r = r(\theta)$ be the root in $(0, 1)$ of*

$$\theta r^2 + (1 - \theta)r - \theta = 0, \quad (4.4)$$

namely

$$r(\theta) = \frac{\sqrt{(1 - \theta)^2 + 4\theta^2} - (1 - \theta)}{2\theta}. \quad (4.5)$$

If the first outcome is sampled from the stationary distribution, then the ultimate ruin and survival probabilities from $x \geq 1$ are

$$\mathbb{P}_x(\tau_0 < \infty) = C(r)r^{x-1}, \quad \mathbb{P}_x(\tau_0 = \infty) = 1 - C(r)r^{x-1}, \quad (4.6)$$

where

$$C(r) = \frac{1 + r + r^2}{2(1 + r)}. \quad (4.7)$$

The same quantity is the extreme-target limit:

$$\lim_{m \rightarrow \infty} \mathbb{P}_x(T_m < \tau_0) = 1 - C(r)r^{x-1}. \quad (4.8)$$

Moreover, on survival, $T_m/m \rightarrow 2$ almost surely.

Proof. Condition on the outcome immediately preceding the next move. Define

$$f_W(x) = \frac{r^{x-1}}{1 + r}, \quad f_L(x) = r^x, \quad x \geq 1,$$

and set $f_L(0) = 1$. Equation (4.4) gives the two harmonic identities

$$f_W(x) = \theta f_W(x + 2) + (1 - \theta)f_L(x - 1),$$

$$f_L(x) = (1 - \theta)f_W(x + 2) + \theta f_L(x - 1).$$

Consequently, $f_{J_{n \wedge \tau_0}}(S_{n \wedge \tau_0})$ is a bounded martingale. By the ergodic theorem applied to J_n , $S_n/n \rightarrow 1/2$ almost surely. On survival the martingale therefore converges to zero; on ruin it equals $f_L(0) = 1$. Bounded convergence shows that f_W and f_L are precisely the ruin probabilities conditional on the current outcome state.

Before the first outcome, a win and a loss each have probability $1/2$. Thus, for every $x \geq 1$ (including $x = 1$, where the first loss is immediate ruin),

$$\begin{aligned} \mathbb{P}_x(\tau_0 < \infty) &= \frac{1}{2}f_W(x + 2) + \frac{1}{2}f_L(x - 1) \\ &= \frac{1}{2}r^{x-1} \left(1 + \frac{r^2}{1 + r} \right) = C(r)r^{x-1}. \end{aligned}$$

Propositions 2.1 and 2.2 finish the proof. $\square$

Corollary 4.2 (Persistence penalty). *For every fixed $x \geq 1$, the survival probability in (4.6) is strictly decreasing in $\theta \in (0, 1)$.*

Proof. Solving (4.4) for θ gives

$$\theta = \frac{r}{1 + r - r^2},$$

whose derivative with respect to $r \in (0, 1)$ is $(1 + r^2)/(1 + r - r^2)^2 > 0$. Hence r increases strictly with θ . The function $C(r)$ also increases strictly, and so does $C(r)r^{x-1}$. The survival probability is its complement. $\square$

For one unit of initial capital the effect is substantial:

θ	outcome pattern	extreme-target limit
0.10	strongly alternating	0.494571
0.25	alternating tendency	0.464816
0.50	independent	0.381966
0.75	persistent streaks	0.305746
0.90	strong streaks	0.270068
0.99	very long streaks	0.251888




Figure 1: All models have stationary win probability $1/2$ and mean increment $1/2$. Survival decreases with persistence on $[0, 1)$, but the fully frozen endpoint jumps from the left limit $1/4$ to $1/2$.

5 The singular persistence edge

The endpoints $\theta = 0$ and $\theta = 1$ are reducible and must be treated separately.

Proposition 5.1 (Endpoint laws). *If $\theta = 0$, outcomes alternate deterministically after the first outcome. Hence*

$$\mathbb{P}_1(\tau_0 = \infty) = \frac{1}{2}, \quad \mathbb{P}_x(\tau_0 = \infty) = 1 \quad (x \geq 2).$$

If $\theta = 1$, the first outcome repeats forever, and

$$\mathbb{P}_x(\tau_0 = \infty) = \frac{1}{2} \quad (x \geq 1).$$

For every fixed $x \geq 1$,

$$\lim_{\theta \uparrow 1} \mathbb{P}_x(\tau_0 = \infty) = \frac{1}{4} \neq \frac{1}{2} = \mathbb{P}_x^{(\theta=1)}(\tau_0 = \infty). \quad (5.1)$$

Proof. The endpoint descriptions are immediate. For the left limit, use $r(\theta) \uparrow 1$ and $C(r) \rightarrow 3/4$ in (4.6). $\square$

This discontinuity has a pathwise explanation. At $\theta = 1$, an initial win is permanent and protects the gambler forever. For every $\theta < 1$, however close to one, switches occur infinitely often almost surely. Over an infinite horizon, a merely rare switch is therefore not equivalent to an impossible switch.

Theorem 5.2 (Noncommuting target and persistence limits). *Let*

$$H_m(\theta) = \mathbb{P}_1(T_m < \tau_0), \quad m \geq 2.$$

Then

$$\lim_{\theta \uparrow 1} \lim_{m \rightarrow \infty} H_m(\theta) = \frac{1}{4}, \quad (5.2)$$

whereas

$$\lim_{m \rightarrow \infty} \lim_{\theta \uparrow 1} H_m(\theta) = \frac{1}{2}. \quad (5.3)$$

Proof. For fixed $\theta < 1$, Theorem 4.1 gives $\lim_m H_m(\theta) = 1 - C(r)$; Proposition 5.1 gives the first iterated limit. For fixed m , a first loss ruins the gambler. Conditional on a first win, the probability of retaining the winning state long enough to cross m tends to one as $\theta \uparrow 1$. Hence $\lim_{\theta \uparrow 1} H_m(\theta) = 1/2$, proving the second limit. $\square$

The discontinuity is resolved by scaling the initial capital with the mean streak length.

Theorem 5.3 (Persistence-capital transition window). *Let $\theta = 1 - \varepsilon$ with $\varepsilon \downarrow 0$, and let $x_\varepsilon \in \mathbb{Z}_{>0}$. If*

$$\varepsilon x_\varepsilon \longrightarrow \lambda \in [0, \infty],$$

then

$$\mathbb{P}_{x_\varepsilon}(\tau_0 = \infty) \longrightarrow 1 - \frac{3}{4}e^{-\lambda/2}. \quad (5.4)$$

If $\varepsilon x_\varepsilon \rightarrow \infty$, the limit is one.

Proof. Expansion of (4.4) at $\theta = 1$ gives

$$r(1 - \varepsilon) = 1 - \frac{\varepsilon}{2} + O(\varepsilon^2), \quad C(r) \longrightarrow \frac{3}{4}.$$

For finite λ , the expansion implies

$$r^{x_\varepsilon - 1} = \exp((x_\varepsilon - 1) \log r) \longrightarrow e^{-\lambda/2},$$

and (4.6) yields (5.4). Moreover, $\log r(1 - \varepsilon)/\varepsilon \rightarrow -1/2$; hence if $\varepsilon x_\varepsilon \rightarrow \infty$, then $(x_\varepsilon - 1) \log r \rightarrow -\infty$ and the ruin probability tends to zero. $\square$

The scale $x \asymp (1 - \theta)^{-1}$ is natural: it is the order of the mean length of a persistent run. Capital much smaller than this scale does not buffer a typical adverse streak; capital much larger than this scale makes ultimate survival overwhelmingly likely.

6 Interpretation, verification, and scope

The mathematical message is not that a small bankroll can create practical wealth from nothing. It is that an infinite-horizon first-passage probability can have a positive mass at “unbounded maximum.” If

$$M = \sup_{0 \leq n < \tau_0} S_n$$

is allowed to take the value $+\infty$, then

$$\mathbb{P}_x(M = +\infty) = \mathbb{P}_x(\tau_0 = \infty).$$

The same surviving paths cross every finite target, so the corresponding hitting events are nested rather than disjoint. They do not form a probability mass function over target levels.

Several modeling assumptions are essential. The horizon is infinite; the payoff rule and transition probabilities never change; capital is integer-valued; transaction costs, discounting, target deadlines, and betting limits are absent; and ruin occurs exactly at zero. In a finite lifetime, the relevant object is $\mathbb{P}(T_m \leq N < \tau_0)$, which can be very small for an astronomical m even when the infinite-horizon limit is positive. Proposition 2.2 quantifies the time issue: in the fair stationary models, a surviving path needs asymptotically two rounds per unit of target capital.

All displayed exact formulas were checked independently in two ways. First, the critical expression (3.3) was substituted exactly into its third-order recurrence and all boundary conditions. Second, the harmonic identities in the proof of Theorem 4.1 were verified over a grid of parameters and capitals, and Monte Carlo first-passage estimates to a high barrier agreed with (4.6). These computations are consistency checks; the proofs above are analytic and cover the full parameter ranges.

A literature-overlap screen was performed in August 2026 using the exact title, the payoff pair $\{+2, -1\}$, the phrases “noncommuting limits” and “persistence-capital transition window,” the displayed characteristic equations, and combinations of gambler’s ruin, correlated random walk, and Markov first passage. The search recovered the classical neighboring work cited in the Introduction but no source presenting this same model and bundle of formulas. This is a topical duplication screen, not a proof of global priority or a substitute for journal similarity software and peer review.

From a risk perspective, the Markov extension makes a precise point: marginal win frequency and expected payoff do not determine ruin probability. The ordering of outcomes matters. Positive serial dependence produces longer loss streaks and a larger long-run variance, and it can substantially reduce the chance that a small reserve survives long enough for the positive drift to be realized.

7 Conclusion

For the asymmetric walk with gains $+2$ and losses -1 , a high target and a low ruin boundary play very different roles. With positive drift, surviving forever is equivalent, up to null events, to eventually crossing every finite target. Hence the extreme-target probability converges to a positive survival mass rather than to zero.

Introducing Markov dependence shows that this mass is not determined by the one-step law. A single persistence parameter, while preserving the stationary win probability and mean payoff, produces an exact and strictly monotone change in survival. The reducible endpoint $\theta = 1$ then creates a jump, two noncommuting limits, and a transition window governed by $x(1 - \theta)$. The resulting model is elementary enough for complete formulas yet rich enough to separate mean reward, temporal clustering, early absorption, target height, and time horizon.

References

- [1] F. Avram and M. Vidmar, First passage problems for upwards skip-free random walks via the scale functions paradigm, *Advances in Applied Probability* **51** (2019), 408–424. doi:10.1017/apr.2019.17.
- [2] M. Brown, E. A. Peköz, and S. M. Ross, Some results for skip-free random walk, *Probability in the Engineering and Informational Sciences* **24** (2010), 491–507. doi:10.1017/S0269964810000136.
- [3] W. Feller, *An Introduction to Probability Theory and Its Applications*, Vol. I, 3rd ed., Wiley, New York, 1968.
- [4] C.-D. Fuh, Uniform Markov renewal theory and ruin probabilities in Markov random walks, *The Annals of Applied Probability* **14** (2004), 1202–1241. doi:10.1214/105051604000000260.
- [5] R. Lal and U. N. Bhat, Some explicit results for correlated random walks, *Journal of Applied Probability* **26** (1989), 757–766. doi:10.2307/3214380.
- [6] C. Mohan, The gambler’s ruin problem with correlation, *Biometrika* **42** (1955), 486–493. doi:10.1093/biomet/42.3-4.486.
- [7] A. Mukherjea and D. A. Steele, Occupation probability of a correlated random walk and a correlated ruin problem, *Statistics & Probability Letters* **5** (1987), 105–111. doi:10.1016/0167-7152(87)90064-2.
- [8] V. Pozdnyakov, Martingale approach to gambler’s ruin problem for correlated random walks, *Journal of Applied Probability* **62** (2025), 1594–1606. doi:10.1017/jpr.2025.25.
- [9] A. D. Proudfoot and D. G. Lampard, A random walk problem with correlation, *Journal of Applied Probability* **9** (1972), 436–440. doi:10.2307/3212813.