The Gambler's Ruin Problem for One-Dimensional Random Walks: Simple Symmetric Random Walk and Some Extensions

Author
CHANG Yung-Hsuan
eiken.sc11@nycu.edu.tw

Advisor
Yuki CHINO
chino@nycu.edu.tw

December 22, 2025

Abstract

The gambler's ruin problem serves as a classical entry point into random walk theory. For the simple symmetric random walk, explicit expressions for hitting probabilities can be derived using difference equations or martingale methods. However, for more general step distributions, the analysis becomes challenging due to the "overshoot" phenomenon. This study explores the problem in three settings: the simple symmetric random walk, the spread-out model where steps are bounded, and the finite-variance case. The methods rest on martingale theory, in particular the optional sampling theorem, whose role we make explicit throughout. We then complement the theory with a Monte Carlo study that visualises the common linear law and probes the sharpness of its constants. This report is expository in nature: all results are drawn from the cited literature, and our aim is a unified, self-contained account rather than the establishment of new theorems.

Table of Contents

List of Figures

1 Introduction

We begin by introducing the fundamental concepts of our model, starting with the random walk, the stopping time, the formulation of the gambler's ruin problem, and the martingale tools used throughout. This work is a study report rather than original research: the statements presented are established results, which we follow from Lawler and Limic (2010) and Koralov and Sinai (2007) rather than claim as new, our contribution being a unified and self-contained exposition across the three settings together with a numerical companion. Unless otherwise stated, the definitions and theorems follow Lawler and Limic (2010).

1.1 Random Walk and Stopping Time

Definition 1 (Random walk). A *random walk* is a sequence $S = (S_0, S_1, S_2, \dots)$, where S_n is the sum of the starting point S_0 and independent, identically distributed steps X_i,

$$S_n = S_0 + X_1 + \cdots + X_n.$$

Throughout this work we focus on the unbiased walk, whose steps have zero mean, $\mathbb{E}(X_i) = 0$.

Definition 2 (Stopping time). The *stopping time* η_r is the first time the walk S exits the interval $(0, r)$,

$$\eta_r = \inf\{n \in \mathbb{N} : S_n \leq 0 \text{ or } S_n \geq r\},$$

where $\mathbb{N} = \{0, 1, 2, \dots\}$ includes 0, so that $\eta_r = 0$ exactly when the walk starts outside $(0, r)$.

For a simple symmetric random walk, the stopping time is finite almost surely, meaning that $\mathbb{P}(\eta_r < \infty) = 1$. This can be understood intuitively through a block argument. In any block of r steps, there is a strictly positive probability $\varepsilon \geq (1/2)^r$ that the walk moves r steps to the right consecutively, guaranteeing an exit eventually.

1.2 The Gambler's Ruin Problem

The problem is formulated as follows. A gambler starts to bet with initial wealth $S_0 = x \in (0, r)$, and we ask for the probability of winning, that is, of reaching the intended wealth r before going broke at 0. We fix this quantity as our central object of study.

Definition 3 (Winning probability). For an initial position $x \in (0, r)$, the *winning probability* is

$$u(x) := \mathbb{P}^x\left(S_{\eta_r} \geq r\right),$$

the probability that the walk exits through the upper boundary, starting from x. We extend u outside $(0, r)$ by the convention $u(z) = 1$ for $z \geq r$ and $u(z) = 0$ for $z \leq 0$, consistent with $\eta_r = 0$ there.

For the simple ± 1 walk the position lands exactly on a boundary, so $u(x) = \mathbb{P}^x\left(S_{\eta_r} = r\right)$; once larger steps are allowed the walk may overshoot, and the event $S_{\eta_r} \geq r$ is the right generalisation. The same symbol $u(x)$ denotes this winning probability in every section below.

1.3 Martingale Tools

Both routes to the gambler's ruin problem run on a single engine, the optional sampling theorem applied to the walk stopped at η_r. We record the facts we use; their proofs may be found in Lawler and Limic (2010) and Koralov and Sinai (2007).

Definition 4 (Martingale). Let $(\mathcal{F}_n)_{n \geq 0}$ be a filtration. An adapted, integrable sequence $(M_n)_{n \geq 0}$ is a *martingale* if $\mathbb{E}(M_{n+1} \mid \mathcal{F}_n) = M_n$ for every n. In particular, if the steps are integrable with $\mathbb{E}(X_i) = 0$, the walk S_n is a martingale for its natural filtration $\mathcal{F}_n = \sigma(X_1, \ldots, X_n)$, since $\mathbb{E}(S_{n+1} \mid \mathcal{F}_n) = S_n + \mathbb{E}(X_{n+1}) = S_n$.

Theorem 5 (Optional sampling). Let (M_n) be a martingale and η a stopping time. If η is bounded, or more generally if the stopped sequence $\left(M_{n \wedge \eta}\right)_{n \geq 0}$ is uniformly integrable — for instance if it is bounded, or bounded in L^2 — then

$$\mathbb{E}\left(M_\eta\right) = \mathbb{E}(M_0).$$

Theorem 6 (Wald's identities). Let (X_i) be i.i.d. with $\mathbb{E}|X_1| < \infty$, and let η be a stopping time with $\mathbb{E}(\eta) < \infty$. Then $\mathbb{E}\left(S_\eta\right) = \mathbb{E}(X_1) \cdot \mathbb{E}(\eta)$. If moreover $\sigma^2 = \mathbb{V}(X_1) < \infty$, then

$$\mathbb{E}\left(\left(S_\eta - \eta \cdot \mathbb{E}(X_1)\right)^2\right) = \sigma^2 \cdot \mathbb{E}(\eta).$$

For the mean-zero steps of this report these read $\mathbb{E}\left(S_\eta\right) = 0$ and $\mathbb{E}\left({S_\eta}^2\right) = \sigma^2 \cdot \mathbb{E}(\eta)$.

2 The Simple Symmetric Random Walk

For a simple random walk, the step probabilities are $\mathbb{P}(X_i = 1) = p$ and $\mathbb{P}(X_i = -1) = q = 1 - p$. In the symmetric case $p = q = 1/2$. Assuming x and r are integers, the walk lands exactly on a boundary, so $u(x) = \mathbb{P}^x\left(S_{\eta_r} = r\right)$.

It is worth noting that if the walk were biased with $p \neq 1/2$, the linear scaling we are about to derive would break down. Solving the same recurrence in the biased case gives

$$u(x) = \frac{1 - (q/p)^x}{1 - (q/p)^r},$$

so the winning probability depends exponentially on x and r through the ratio q/p rather than scaling linearly; for instance, when $p > 1/2$ with x fixed, it converges to $1 - (q/p)^x > 0$ as $r \to \infty$. Since we assume symmetry, however, we may use the following two approaches.

2.1 Difference Equations Approach

This method follows Koralov and Sinai (2007). By conditioning on the first step, we obtain the relation

$$\begin{aligned} u(x) &= \mathbb{P}^x\left(S_{\eta_r} = r \mid X_1 = 1\right) \cdot \mathbb{P}(X_1 = 1) + \mathbb{P}^x\left(S_{\eta_r} = r \mid X_1 = -1\right) \cdot \mathbb{P}(X_1 = -1) \\ &= u(x+1) \cdot p + u(x-1) \cdot (1-p). \end{aligned}$$

In the symmetric case $p = 1/2$ this reduces to the harmonic recurrence

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

with boundary conditions $u(0) = 0$ and $u(r) = 1$ for $x \in \{1, \dots, r-1\}$.

Theorem 7 (Simple symmetric solution). For the simple symmetric random walk, the winning probability is the linear function

$$u(x) = \frac{x}{r}.$$

This is the unique solution of the harmonic recurrence above, as one verifies directly. We can extend the same conditioning philosophy to a countable-step-range case or a continuous-step-range case. For the countable case the relation becomes

$$\begin{aligned}\mathbb{P}^x\left(S_{\eta_r} = r\right) &= \sum_k \mathbb{P}^x\left(S_{\eta_r} = r \,\middle|\, X_1 = k\right) \cdot \mathbb{P}(X_1 = k) \\ &= \sum_k \mathbb{P}^{x+k}\left(S_{\eta_r} = r\right) \cdot \mathbb{P}(X_1 = k),\end{aligned}$$

while for the continuous case we have

$$\begin{aligned}u(x) &= \int_{\mathbb{R}} \mathbb{P}^x\left(S_{\eta_r} \geq r \,\middle|\, X_1 = y\right) f_{X_1}(y)\,\mathrm{d}y \\ &= \int_{\mathbb{R}} u(x+y) f_{X_1}(y)\,\mathrm{d}y \\ &= \int_{\mathbb{R}} u(x+y)\,\mathbb{P}(X_1 \in \mathrm{d}y).\end{aligned}$$

The convention extending u outside $(0, r)$ in our definition of the winning probability makes the integrand meaningful even when the first step y already carries the walk past a boundary.

2.2 Martingale Approach

Alternatively, we employ the tools of Section 1.3. Let $M_n = S_{n \wedge \eta_r}$. The stopped process M_n is a bounded martingale, so the optional sampling theorem gives $\mathbb{E}^x(M_0) = \mathbb{E}^x\left(M_{\eta_r}\right)$. This implies

$$x = r \cdot \mathbb{P}^x\left(S_{\eta_r} = r\right) + 0 \cdot \mathbb{P}^x\left(S_{\eta_r} = 0\right),$$

which recovers the same result

$$u(x) = \frac{x}{r}.$$

3 The Spread-Out Model (Bounded Steps)

In the previous section, the analysis relied on the fact that the walk stops exactly at 0 or r. When the step distribution allows values other than 1 and -1, specifically continuous random variables or larger discrete jumps, the walk will likely overshoot the boundaries. This phenomenon complicates the exact calculation of the winning probability.

We now assume that the steps X_i are independent, identically distributed random variables with mean zero. We further assume the steps are bounded by K almost surely and satisfy a non-degeneracy condition.

3.1 The Overshoot Phenomenon

The stopping time η_r remains the first time the walk exits $(0, r)$. However, at time η_r, the position S_{η_r} is not necessarily 0 or r. Instead, if the gambler wins, the position lands in the interval $[r, r + K]$, and if the gambler loses, it lands in $[-K, 0]$.

Because of this, the relation derived from the optional sampling theorem, $x = \mathbb{E}^x\left(S_{\eta_r}\right)$, no longer simplifies to $x = r \cdot \mathbb{P}^x\left(S_{\eta_r} = r\right)$.

3.2 Deriving the Bounds via Martingale

Despite the overshoot, the martingale property provides useful bounds. If the steps are bounded by K, so that $\mathbb{P}(|X_i| > K) = 0$, the stopped process $S_{n \wedge \eta_r}$ is uniformly bounded in $[-K, r + K]$. The optional sampling theorem then applies and gives

$$x = \mathbb{E}^x\left(S_{\eta_r}\right).$$

We decompose this expectation over the winning and losing events, recalling $u(x) = \mathbb{P}^x\left(S_{\eta_r} \geq r\right)$,

$$x = \mathbb{E}^x\left(S_{\eta_r} \,\middle|\, S_{\eta_r} \geq r\right)u(x) + \mathbb{E}^x\left(S_{\eta_r} \,\middle|\, S_{\eta_r} \leq 0\right)(1 - u(x)).$$

Using the overshoot bounds $r \leq S_{\eta_r} \leq r + K$ on a win and $-K \leq S_{\eta_r} \leq 0$ on a loss, we obtain

$$r \cdot u(x) - K \cdot (1 - u(x)) \leq x \leq (r + K) \cdot u(x) + 0.$$

Theorem 8 (Basic bounds). For a mean-zero walk with steps bounded by K,

$$\frac{x}{r + K} \leq u(x) \leq \frac{x + K}{r + K}.$$

3.3 An Estimate Theorem

While the basic bounds are useful, a sharper estimate exists that preserves the linear scaling seen in the simple symmetric case; we will meet the same form again in Section 4.

Theorem 9 (Bounded-step estimate). Let $K < \infty$. For a one-dimensional random walk whose steps satisfy the boundedness condition $\mathbb{P}(|X_i| > K) = 0$, the mean-zero condition $\mathbb{E}(X_i) = 0$, and the non-degeneracy condition $\mathbb{P}(X_i > 0) > 0$, there exist constants c_1, c_2 depending on K such that

$$c_1 \frac{x + 1}{r} \leq u(x) \leq c_2 \frac{x + 1}{r}.$$

This estimate follows Lawler and Limic (2010, 103–112). It confirms that even with the "fuzziness" of the boundary due to overshoot, the winning probability still scales linearly with the initial position x relative to the intended wealth r.

The appearance of $x + 1$ rather than x in the numerator is deliberate. For $x \geq 1$ the two differ only by a constant factor, but as $x \to 0$ the $+1$ records an *entrance effect:* even a walk starting arbitrarily close to 0 can still win with probability of order $1/r$ by taking a single sufficiently large first step, a possibility guaranteed by the non-degeneracy condition. This is precisely where the estimate sharpens the basic bounds. The upper bound $(x + K)/(r + K)$ is already of order $(x + 1)/r$, so the estimate adds nothing there; the gain is in the lower bound, since $x/(r + K)$ degenerates to 0 as $x \to 0$ whereas the estimate keeps it at order $1/r$. The same linear form, obtained here for bounded steps, will reappear in Section 4 under the weaker finite-variance assumption.

4 Finite-Variance Case

The previous model assumed that the step size was strictly bounded. However, in many applications, steps may follow distributions with infinite support, such as a normal distribution. A natural question is whether the linear scaling of the winning probability survives once the boundedness condition is relaxed. The answer is affirmative, provided the random walk has a finite variance.

4.1 The Uniform Estimate Theorem

We now allow the steps X_i to be unbounded, subject to regularity conditions. The strength of the result is that the constants in the estimate depend only on a set of family parameters K, δ, b, ρ, which we first interpret.

Definition 10 (Family parameters). For a mean-zero random walk S, the *family parameters* are the constants K, δ, b, ρ controlling, respectively, the variance, the non-degeneracy, the local survival, and the global fluctuation,

$$\mathbb{E}(X_i^2) \leq K^2, \quad \mathbb{P}(X_i \geq 1) \geq \delta, \quad \inf_{n\in\mathbb{N}} \mathbb{P}\left(\bigcap_{i=1}^{n^2} S_i > -n\right) \geq b, \quad \text{and} \quad \inf_{n\in\mathbb{N}} \mathbb{P}(S_{n^2} \leq -n) \geq \rho.$$

Here K is now a variance bound rather than the almost-sure step bound of Section 3; the symbol is reused following Lawler and Limic (2010), but its meaning has changed. The parameter δ is the quantitative form of the non-degeneracy condition $\mathbb{P}(X_i > 0) > 0$ used in Section 3. The conditions on b and ρ compare the walk to the diffusive scale $S_{n^2} \sim n$ and hold automatically for any finite-variance, non-degenerate, mean-zero law; we revisit this point numerically in Section 5.

Theorem 11 (Gambler's ruin estimate). Let $\delta, K \in \mathbb{R}^+$ and $b, \rho \in (0, 1)$, and let S be a mean-zero random walk satisfying the four conditions above. Then there exist constants c_1, c_2, depending only on these parameters, such that

$$c_1 \frac{x+1}{r} \leq u(x) \leq c_2 \frac{x+1}{r}.$$

This is the general estimation result, often called the gambler's ruin estimate, following Lawler and Limic (2010, 103–112).

It is crucial to note that the finite variance condition (K) is necessary here. To see that variance is the binding constraint, the counterexample must still respect our standing assumption $\mathbb{E}(X_i) = 0$. A Cauchy distribution does not: it has no finite first moment, so $\mathbb{E}(X_i)$ is undefined and S_n is not even a martingale; it would confound the failure of integrability with that of finite variance rather than isolate the latter. A clean example is instead a symmetric distribution with heavy tails $\mathbb{P}(|X_i| > t) \sim t^{-\alpha}$ as $t \to \infty$ for some $1 < \alpha < 2$. Here $\mathbb{E}(|X_i|) < \infty$, so the mean-zero assumption holds and S_n remains a genuine martingale, yet $\mathbb{E}\left({X_i}^2\right) = \infty$. With such steps the walk can make a single large jump from deep inside the interval to far outside, bypassing the boundary entirely, and the winning probability no longer scales linearly as x/r. The precise scaling in this regime is governed by the index α and is a separate, known result; we illustrate the breakdown in Section 5.

4.2 Sketch of the Derivation

The proof is significantly more involved than the bounded case because the overshoot $S_{\eta_r} - r$ can be very large. We cannot simply bound the position by $r + K$.

For the upper bound estimate, we rely on controlling the overshoot. We utilise an overshoot lemma, which estimates the probability of the walk travelling very fast. The derivation constructs constants from the family parameters, ultimately leading to an upper bound coefficient proportional to the variance bound.

For the lower bound estimate, we return to the martingale identity $x = \mathbb{E}^x\left(S_{\eta_r}\right)$. This identity now requires a different justification than in Section 3: the steps are unbounded, so the stopped process $S_{n\wedge\eta_r}$ is no longer confined to $[-K, r + K]$ and the uniform-boundedness argument used there no longer applies. Instead, the finite-variance and non-degeneracy conditions guarantee $\mathbb{E}(\eta_r) < \infty$, and Wald's second identity (Section 1.3) gives

$$\mathbb{E}^x\left[\left(S_{n\wedge\eta_r} - x\right)^2\right] = \sigma^2\,\mathbb{E}^x[n \wedge \eta_r] \le \sigma^2\,\mathbb{E}^x(\eta_r) < \infty,$$

where $\sigma^2 = \mathbb{E}\big(X_i^2\big) \le K^2$. Hence $\big(S_{n\wedge\eta_r}\big)_{n\in\mathbb{N}}$ is bounded in L^2, therefore uniformly integrable, and the optional sampling theorem applies, yielding $x = \mathbb{E}^x\big(S_{\eta_r}\big)$ as claimed. We then decompose the expectation by the magnitude of the win into three parts: losses ($S_{\eta_r} \le 0$), moderate wins ($r \le S_{\eta_r} \le (1+t_0)r$), and huge wins ($S_{\eta_r} \ge (1+t_0)r$).

Since the loss term is non-positive, we focus on the wins. By choosing a threshold parameter t_0 sufficiently large, we bound the contribution of huge wins,

$$\mathbb{E}^x\big(S_{\eta_r}; S_{\eta_r} \ge (1+t_0)r\big) \le \frac{x}{2}.$$

This forces the moderate wins to carry significant weight,

$$\mathbb{E}^x\big(S_{\eta_r}; r \le S_{\eta_r} \le (1+t_0)r\big) \ge \frac{x}{2}.$$

Since $S_{\eta_r} \le (1+t_0)r$ in this term, we deduce

$$(1+t_0)r \cdot \mathbb{P}^x\big(S_{\eta_r} \ge r\big) \ge \frac{x}{2},$$

and rearranging yields the lower bound

$$u(x) \ge \frac{1}{2(1+t_0)}\frac{x}{r}.$$

This sketch produces the form x/r rather than the $(x+1)/r$ of the theorem; the two agree up to a constant for $x \ge 1$, and the $+1$ reflects the same entrance effect discussed for the bounded estimate in Section 3, which a bare x/r would miss as $x \to 0$. While the existence of c_1 and c_2 is guaranteed, deriving explicit expressions in terms of the family parameters is analytically complex; we instead probe them numerically next.

5 Numerical Simulation

The results above are existence statements: they assert linear scaling and the existence of constants c_1, c_2 without exhibiting them. A direct Monte Carlo experiment complements them in four ways, letting us watch the linear law x/r emerge, see it break when finite variance is dropped, read the four family parameters as measurable quantities, and gauge how informative the constants actually are. Throughout we kill the walk on its first exit from $(0, r)$ with $r = 100$, estimate the winning probability $u(x)$ by the fraction of independent paths exiting through the top, and use of order 10^4 paths per starting point. The accompanying code is provided in `simulation/gamblers_ruin_sim.py`.

5.1 Linear Scaling and the Entrance Effect

Figure 1 collects four mean-zero step laws of finite variance: the simple ± 1 walk, a bounded uniform law on $\{-3, ..., 3\}$, the Gaussian law $\mathcal{N}(0, 1)$, and a bounded mixture making occasional jumps of size 6. For all four the estimated winning probability falls on the diagonal x/r, the ± 1 walk reproducing the exact solution $u(x) = x/r$ of Section 2. The right panel magnifies x near 0, where the three spread-out laws sit strictly above the diagonal. This is the entrance effect of Section 3 made visible, that is, the $x + 1$ rather than x in the numerator: a walk started at $x = 1$ still wins with probability of order $1/r$ by taking a single sufficiently large first step, so $u(1)$ exceeds the bare $1/r$ of the ± 1 walk, and the excess grows with the size of the admissible jumps.

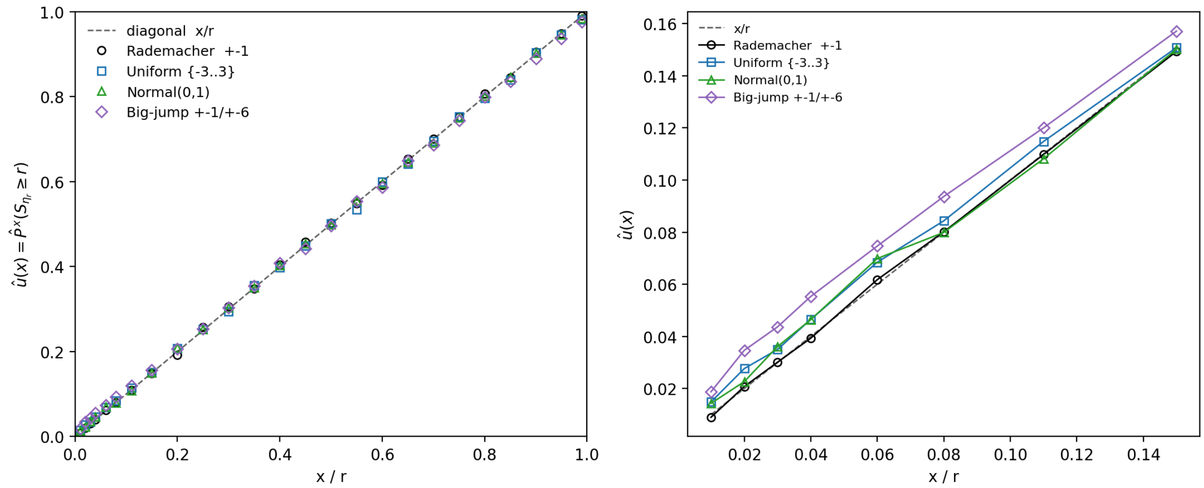

Figure 1: Winning probability $u(x)$ against x/r for four finite-variance step laws (left), with a magnification near the origin (right) exposing the entrance effect, that is, the $x + 1$ in the estimate.

5.2 Failure Under Heavy Tails

To confirm that finite variance is the binding hypothesis, and that its failure is systematic rather than an artefact of one example, we sweep a family of step laws across the variance threshold. Student's t steps on ν degrees of freedom satisfy $\mathbb{E}|X_i| < \infty$ for every $\nu > 1$, so the walk is always a mean-zero martingale, while $\mathbb{E}\left(X_i{}^2\right) < \infty$ precisely when $\nu > 2$. The left panel of Figure 2 shows that as ν falls through 2 the winning probability bends progressively away from the diagonal and is drawn toward $1/2$: from any interior x, a single large jump carries the walk out of $(0, r)$, and the sign of that jump is essentially independent of x, so the linear dependence on the starting point washes out. The laws with $\nu > 2$ stay on the diagonal alongside the Gaussian reference.

To quantify how much a law breaks, the right panel plots the break magnitude $D \coloneqq \max_x [u(x) - x/r]$ against ν. It remains a few hundredths — comparable to the entrance effect — throughout the finite-variance regime, then climbs steadily once $\nu < 2$, reaching roughly 0.1 as $\nu \to 1$. At the finite horizon r the transition is smooth rather than sharp, the estimate being asymptotic; what the experiment isolates is that the deviation is governed by the tail index and becomes substantial exactly where the second moment diverges. This is the quantitative face of the counterexample discussed in Section 4.

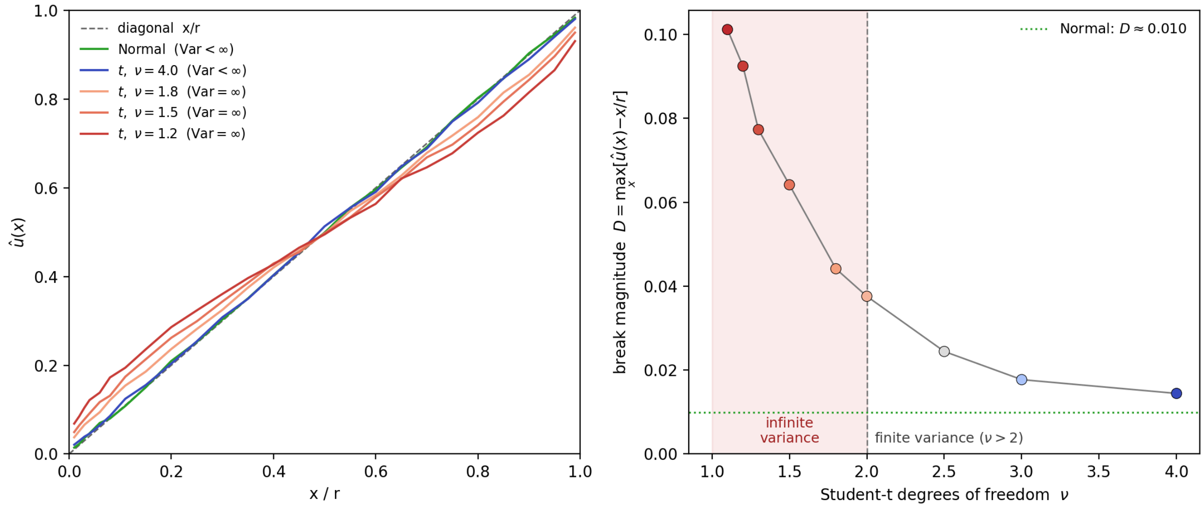

Figure 2: Left: winning probability for Student's t steps at several degrees of freedom ν, bending away from x/r as ν drops below 2, where the variance becomes infinite. Right: the break magnitude $D = \max_x[u(x) - x/r]$ against ν, a few hundredths in the finite-variance regime and growing once $\nu < 2$.

The mechanism is the overshoot $S_{\eta_r} - r$. Figure 3 displays its law conditional on a win. The bounded step laws cap the overshoot at their maximal jump, producing a sharp cliff; the Gaussian overshoot decays quickly; the Student's t overshoot carries a power-law tail extending several orders of magnitude. It is exactly this uncontrolled overshoot that the optional sampling argument of Section 4 must absorb, and that the finite-variance bound on $\mathbb{E}(X_i^2)$ exists to tame.

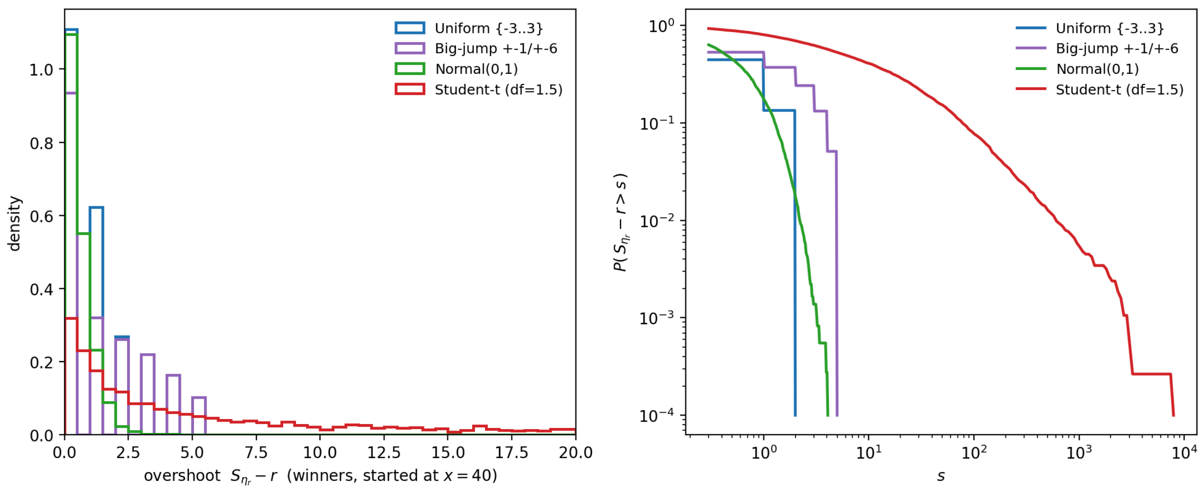

Figure 3: The win-side overshoot $S_{\eta_r} - r$: bounded (cliff), finite-variance (fast decay), and heavy-tailed (power law); the right panel uses log–log axes.

5.3 The Family Parameters as Descriptors

The parameters K, δ, b, ρ read as abstract conditions, but each is a number one can measure from a given step law; they describe the distribution rather than tune the model. The variance bound K and the non-degeneracy $\delta = \mathbb{P}(X_i \geq 1)$ are read off directly. The local-survival b and global-fluctuation ρ both compare the walk to the diffusive scale $S_{n^2} \sim n$, since $\mathbb{V}(S_{n^2}) = n^2\sigma^2$. Figure 4 estimates $\mathbb{P}\left(\bigcap_{i=1}^{n^2} S_i > -n\right)$ and $\mathbb{P}(S_{n^2} \leq -n)$ as functions of n. For the finite-variance laws the former settles to a positive constant — the walk clears the barrier with probability bounded away from 0, giving $b > 0$ — whereas for the Student's t law it decays toward 0, because the fluctuations grow like $n^{4/3}$ and overwhelm the linear barrier $-n$. The heavy-tailed law therefore fails not only the variance bound, whose empirical K diverges, but also the local-survival condition b, and the experiment shows precisely which hypothesis breaks. For the Gaussian law these have clean closed forms, against which the simulation checks out. Here $K = 1$, and since S_{n^2}/n is exactly standard normal, $\rho = \mathbb{P}(S_{n^2} \leq -n) = \Phi(-1) \approx 0.16$ for every n, coinciding with $\delta = \mathbb{P}(X_i \geq 1) = \Phi(-1)$. By the reflection principle the local-survival value is $b = 2\,\Phi(1) - 1 \approx 0.68$. All four parameters are comfortably interior.

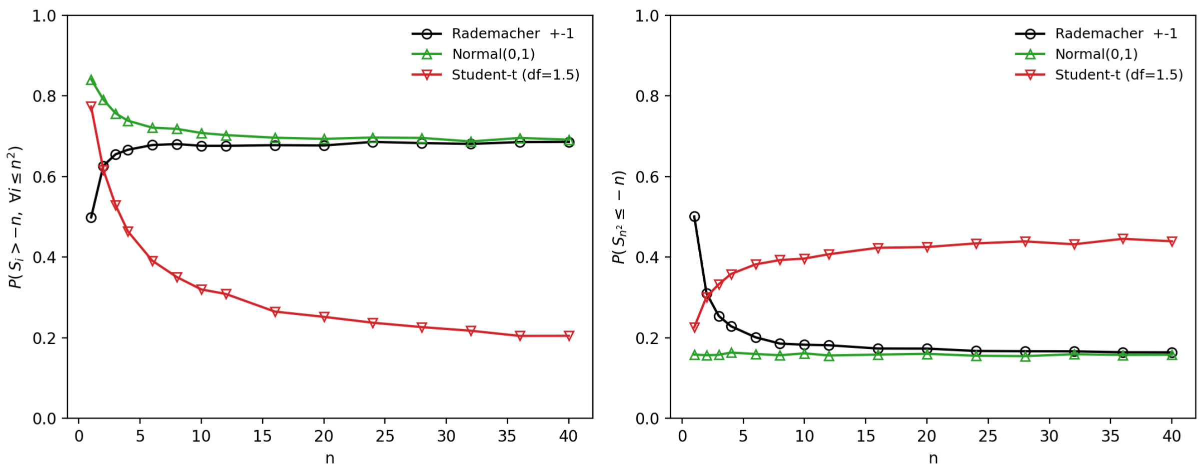

Figure 4: The local-survival probability (left) and the global-fluctuation probability (right) against n. For finite variance both stabilise at positive values; under heavy tails the local-survival probability decays toward 0, violating the condition on b.

5.4 How Sharp Are the Constants?

A uniform estimate $c_1(x+1)/r \le u(x) \le c_2(x+1)/r$ is only as informative as the ratio c_2/c_1: were c_1 near 0 and c_2 enormous, the statement would be true yet empty. We probe this directly through $g(x) := u(x)r/(x+1)$, the quantity the estimate confines to $[c_1, c_2]$. Figure 5 shows that for the well-behaved family g remains within roughly $[0.45, 1.16]$, a ratio near 2.6, with the entire spread concentrated at small x. That factor is structural rather than a defect: it is the cost of the $x+1$ normalisation. Normalising by x alone would pin g to 1 through the bulk but force it to diverge as $x \to 0$, where the entrance effect lives; the $x+1$ form trades a bounded factor for validity uniform down to the boundary. For a fixed, reasonable family the estimate is thus genuinely informative.

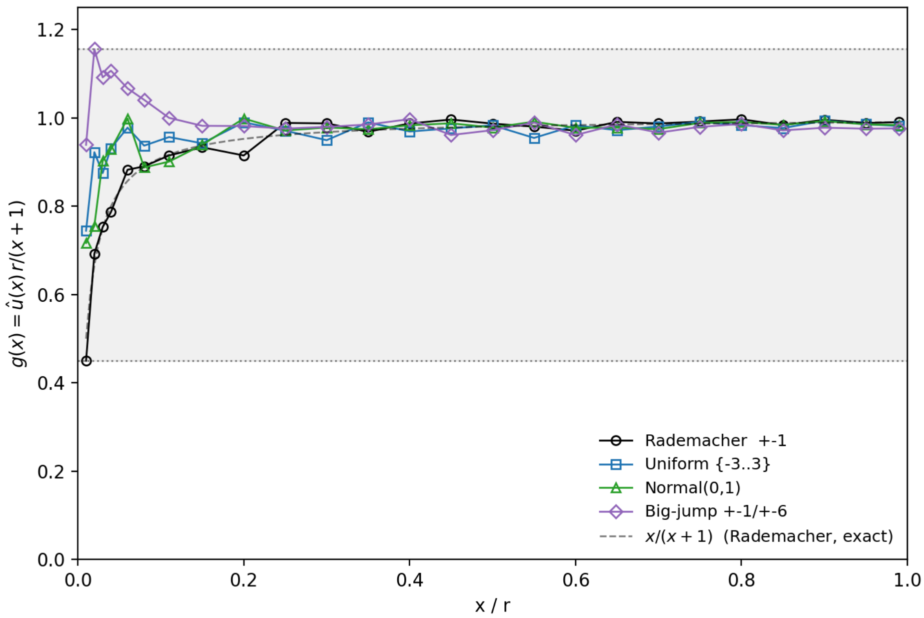

Figure 5: The ratio $g(x) = u(x)r/(x+1)$ remains in a band of width about 2.6 for the finite-variance family, with the gap confined to small x.

The constants are controlled only through the family parameters, however, and degrade as those parameters worsen. Figure 6 fixes r and scales the Gaussian step size, so that $K = \sigma$ and the ratio K/r increases. The upper constant c_2 then climbs from about 1 to nearly 9 while c_1 stays of order 1, and the band $[c_1, c_2]$ widens toward the vacuous. This is the asymptotic nature of the estimate seen from the side: the constants are harmless only when the steps are small relative to the interval, $K \ll r$, and a family with K comparable to r — or, in the limit, of infinite variance — is exactly where the bound loses its force. The simulations thus exhibit the linear law x/r as the common thread running through all three models, while marking out the regime in which it is meaningful.

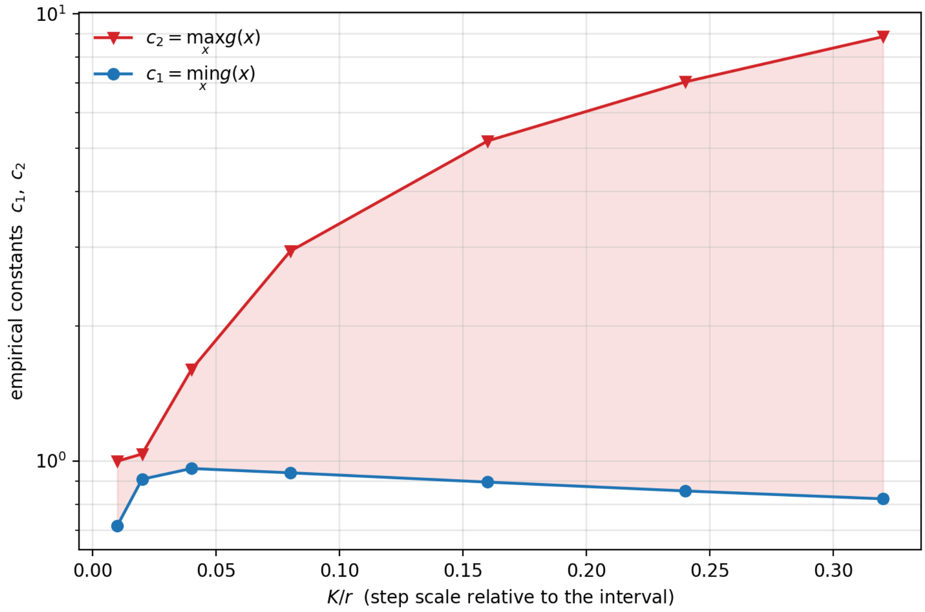

Figure 6: Holding r fixed and increasing K/r, the empirical upper constant $c_2 = \max_x g(x)$ grows without apparent bound while $c_1 = \min_x g(x)$ stays of order 1.

6 Conclusion and Some More Interesting Problems

Across all three models, and confirmed by the simulations, we see the shadow of the linear scaling x/r over each solution and estimate. The gambler's ruin problem is just the tip of the iceberg in random walk theory. Here we briefly introduce three directions for further exploration, focusing on the mathematical structures that emerge in higher dimensions and continuous limits; standard treatments of these topics can be found in Lawler and Limic (2010).

6.1 Connection to Potential Theory (Two-Dimensional Extension)

One of the most beautiful results in probability theory is the connection between random walks and electrostatics. When we extend the problem to two dimensions, for instance a random walk on the grid $\mathbb{Z}^2$, the recurrence for the winning probability $u(x, y)$ satisfies the discrete mean value property

$$u(x, y) = \frac{1}{4}(u(x + 1, y) + u(x - 1, y) + u(x, y + 1) + u(x, y - 1)).$$

This can be rewritten using the discrete Laplacian operator Δ as

$$\Delta u(x, y) = u(x + 1, y) + u(x - 1, y) + u(x, y + 1) + u(x, y - 1) - 4u(x, y) = 0.$$

This is the discrete analogue of Laplace's equation $\Delta V = 0$. In physics, the electric potential in a charge-free region satisfies this property. Thus, calculating the gambler's ruin probability on a grid is mathematically identical to solving for the electric potential on a discrete network of resistors. Theoretically, the shape of the boundary does not change the nature of this connection; the probability is determined by the harmonic measure of the boundary.

6.2 Expected Hitting Time and the Poisson Equation

Besides the probability of winning, another natural question is: how long does the game last? Let $h(x) = \mathbb{E}^{x(\eta_r)}$ be the expected hitting time. For the simple symmetric random walk, we can derive this using a similar difference equation approach. By conditioning on the first step, we have

$$h(x) = 1 + \frac{1}{2} \cdot h(x + 1) + \frac{1}{2} \cdot h(x - 1).$$

The constant term 1 accounts for the single time step taken at the first transition. Rearranging, the second difference is constant,

$$h(x + 1) - 2h(x) + h(x - 1) = -2.$$

This is the discrete analogue of the Poisson equation $\Delta h = f$, with f a non-zero constant. With boundary conditions $h(0) = 0$ and $h(r) = 0$ (the game ends immediately at the boundaries), the solution is the parabola

$$h(x) = x \cdot (r - x).$$

This reveals a profound link to harmonic analysis: while the winning probability corresponds to the Laplace equation (harmonic functions), the expected duration corresponds to the Poisson equation.

6.3 Continuous-Time Limit: Brownian Motion

Finally, if we let the step size $\Delta x \to 0$ and the time step $\Delta t \to 0$ in a specific way where $\Delta t \approx (\Delta x)^2$, the discrete random walk converges to a continuous stochastic process known as Brownian motion, denoted B_t.

In this limit, the discrete difference equation transforms into a differential equation. Expanding $u(x \pm \Delta x) = u(x) \pm \Delta x \cdot u'(x) + \frac{1}{2}(\Delta x)^2 u''(x) + o\left((\Delta x)^2\right)$, the relation $u(x) = \frac{1}{2} \cdot u(x + \Delta x) + \frac{1}{2} \cdot u(x - \Delta x)$ leaves

$$\frac{1}{2}(\Delta x)^2 u''(x) + o\left((\Delta x)^2\right) = 0;$$

dividing by $(\Delta x)^2$ and letting $\Delta x \to 0$ yields

$$\frac{1}{2}u''(x) = 0.$$

The solution to this ordinary differential equation with boundary conditions $u(0) = 0$ and $u(r) = 1$ is simply the linear function

$$u(x) = \frac{x}{r}.$$

This confirms that our discrete result x/r is consistent with the continuous theory. The overshoot problem disappears in Brownian motion because the path is continuous, ensuring it must hit the boundary exactly without jumping over it.

Bibliography

Koralov, Leonid B., and Yakov G. Sinai. 2007. *Theory of Probability and Random Processes.* 2nd ed. Springer. https://doi.org/10.1007/978-3-540-68829-7.

Lawler, Gregory F., and Vlada Limic. 2010. *Random Walk: A Modern Introduction.* Cambridge University Press. https://www.math.uchicago.edu/~lawler/srwbook.pdf.