Munich Personal RePEc Archive

Ruin Probabilities for Strategies with Asymmetric Risk

Whelan, Karl

University College Dublin

July 2025

Online at https://mpra.ub.uni-muenchen.de/126349/
MPRA Paper No. 126349, posted 20 Oct 2025 07:19 UTC

Ruin Probabilities for Strategies with Asymmetric Risk

Karl Whelan*

University College Dublin

September 2025

Abstract

We analyze sequential investment strategies that expose capital to asymmetric payoff structures–scenarios in which losses are small and frequent, while gains are large but rare. This setup generalizes the classical gambler's ruin problem, traditionally framed as a symmetric game, into a framework for modeling repeated financial decisions under uncertainty. Payoff asymmetry is common in domains such as venture capital, tail-risk hedging, and derivative strategies. We consider cases where each investment has positive, zero, or negative expected return, and derive analytic results for ruin probabilities, expected final wealth, and game duration. Our findings show that increasing asymmetry–higher potential rewards but lower success probability–raises the likelihood of ruin in positive-return settings and mitigates it when returns are negative. For zero-return strategies, we establish bounds on ruin probabilities and show that convergence to terminal outcomes is faster when payoffs are skewed. The results have implications for portfolio risk management and capital allocation in repeated-risk environments.

Keywords: Gambler's Ruin, Asymmetric Risk, Stopping Problems
JEL Classifications: C61, G11

*karl.whelan@ucd.ie.

1. Introduction

Many real-world investment strategies involve repeated exposure to asymmetric risk–scenarios in which losses are frequent and limited, while gains, though rare, can be substantial. For example, venture capital portfolios typically lose money on most individual investments, with overall returns driven by a few outsized successes; empirical work by Kerr, Nanda, and Rhodes-Kropf (2014) documents this power-law pattern in venture outcomes. In derivatives markets, traders often buy out-of-the-money options that pay off only under extreme conditions, with Broadie, Chernov, and Johannes (2009) showing such strategies typically have negative average returns but large occasional gains. Tail-risk hedging strategies similarly involve paying regular premiums to insure against rare market crashes; see Ilmanen et al. (2021) for evidence on the performance of such approaches across market regimes.

These strategies share a common structure: they involve repeated decisions to risk capital in exchange for rare, asymmetric rewards. Understanding how wealth evolves under such dynamics–and how likely it is to be depleted before reaching a target level–is essential for portfolio allocation, risk management, and product design.

This paper extends the classical gambler's ruin model–originally framed as a symmetric game of equal-sized wins and losses–to accommodate asymmetric payoff structures.[1] We model each investment decision as a probabilistic "game" in which the investor either loses a fixed amount or gains a multiple of that amount with lower probability. This framing allows us to study the cumulative behavior of wealth under repeated exposure to such opportunities.

We analyze three settings: (i) zero expected return, (ii) positive expected return, and (iii) negative expected return, and examine how the probability of ruin, expected final wealth, and expected game duration vary with payoff asymmetry. In each case, we derive both analytic results and simulation-based validation.

Our findings include:

- For positive-return strategies, higher payoff asymmetry (involving larger rewards at lower probabilities) increases the risk of ruin and reduces expected final wealth, even when expected value per play is held constant.
- For zero-return strategies, long-run ruin probabilities remain close to the classical formula, but outcomes converge more rapidly when payoffs are highly skewed due to increased return variance.
- For negative-return strategies, higher asymmetry improves outcomes: ruin becomes less likely and expected wealth improves, because rare large gains occasionally overcome the long-run

[1] For the historical origins of the problem, which dates from an exchange between Pascal and Fermat, and developments in deriving solutions, see Takacs (1969), Edwards (1983), and Song and Song (2013).

disadvantage.

To illustrate the magnitude of these effects, consider an investor allocating 1% of capital per trade to a repeated strategy with an expected return of 1%, targeting a tripling of wealth. When the winning payoff equals the loss (i.e. symmetric), the probability of ruin is 13% and expected final wealth is 2.6 times the initial amount. If the winning payoff is twice the capital at risk, the ruin rate rises to 34% and expected wealth falls to 2.0 times the starting value. With a 20-to-1 payoff, ruin becomes likely (64%) and expected wealth drops to just 1.1 times initial capital. The skewed payoff increases variance, raising the chance of hitting ruin early and preventing the investor from realizing long-run gains.

Conversely, for a negative-return strategy (with expected return of -1%), outcomes improve with greater payoff asymmetry. In the symmetric case, ruin is near-certain (98%) and expected wealth falls to 5% of the original amount. If the winning payoff is 20 times the risked amount, ruin falls to 71% and expected wealth rebounds to 90% of the original level. Here, the higher variance occasionally allows the investor to escape the losing game before ruin.

We also consider how varying the stake size–i.e., the amount of wealth placed at risk per round–interacts with payoff asymmetry. We find that smaller stakes reduce ruin risk in positive-return environments, echoing the logic of the Kelly criterion for optimal betting under uncertainty (Kelly, 1956; Breiman, 1961). Conversely, larger stakes are preferable in negative-return settings, consistent with recommendations for "bold play" in adverse conditions (Dubins and Savage, 1965). However, the marginal impact of stake size diminishes when payoff asymmetry is high.

Our work builds on extensions of the gambler's ruin problem by Feller (1950), Harper and Ross (2005), and Hunter et al. (2008), but contributes a more systematic treatment of asymmetric strategies relevant to financial decision-making. Though Feller (1950) briefly considered cases with multi-unit gains and losses, and Harper and Ross (2005) discussed asymmetric variants of the ruin problem, prior work did not fully characterize how key outcomes, such as ruin probabilities, expected wealth, and game duration, varied with increasing payoff asymmetry. Hunter et al. (2008) generalized the gambler's ruin problem to include probabilities of jumping straight to target wealth or zero wealth using the same solution methods that we use. They noted that introducing these elements led to ruin rates for games with zero expected profit being a nonlinear function of initial wealth, something that will also apply in our analysis of games with two asymmetric payoffs. Our results complement this literature by providing clear analytic characterizations and simulation-based validation.

The remainder of the paper is structured as follows. Section 2 reviews the symmetric gambler's ruin model. Section 3 describes the generalization to asymmetric payoffs. Section 4 presents analytic results and illustrations. Section 5 extends the model to varying stake sizes. Section 6 concludes.

2. Symmetric Payoffs

Here we provide a short description of the classic version of the gambler's ruin in which the size of potential wins and losses are the same, as presented for example by Feller (1950).

2.1. Positive or negative value expected profit games

An investor starts with wealth of $W_0 = n$ and wealth changes over time according to

$$W_t = W_0 + \sum_{i=1}^{t} X_i \tag{1}$$

where X_i is the profit from the ith playing of the game, which is an i.i.d. Rademacher variable

$$X_i = \begin{cases} 1 & \text{with probability } \ p \\ -1 & \text{with probability } \ 1-p \end{cases} \tag{2}$$

so $E(X_i) = 2p - 1$. The investor keeps playing the game until they reach a target level of wealth T or they lose all their wealth.[2] Denoting D_n as the duration of play, the game ends with either $W_{D_n} = 0$ or $W_{D_n} = T$. Define the probability of reaching the target level of wealth starting from $W_0 = n$ as

$$P_n = \text{Prob}\left(W_{D_n} = T \mid W_0 = n\right) \tag{3}$$

This probability satisfies a difference equation of the form

$$P_n = pP_{n+1} + (1-p)\,P_{n-1} \tag{4}$$

The associated characteristic equation for this difference equation is

$$pr^2 - r + 1 - p = 0 \tag{5}$$

which has roots $r_1 = \frac{1-p}{p}$ and $r_2 = 1$. Here, we consider the case $p \neq 0.5$ where there is either a positive or negative expected profit and discuss the zero profit case $p = 0.5$ below. With $p \neq 0.5$, the general solution is of the form

$$P_n = A\left(\frac{1-p}{p}\right)^n + B\,(1)^n = A\left(\frac{1-p}{p}\right)^n + B \tag{6}$$

[2] Some presentations define this as a two-player problem where one player has wealth of n, the other has wealth of m and they both keep playing until one of them has wealth of $T = n + m$ and the other has zero. We will just assume the investor can set any target level of wealth and there is a willing counter-party on the other side who agrees to keep playing the game.

Specific solutions are obtained from using the boundary conditions, $P_0 = 0$ (you have no money left) and $P_T = 1$ (the target wealth is reached) which imply a solution of the form

$$P_n = \frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^T - 1} \tag{7}$$

When $p < 0.5$, so the expected profit of each play is negative, $\frac{1-p}{p} > 1$ and the denominator gets ever larger as the target level of wealth T increases, meaning the probability of success goes to zero and the probability of ruin becomes one, which is the classic gambler's ruin result.

When the game has a positive expected profit ($p > 0.5$), one might imagine that as the target wealth level increases, so the investor can play for ever longer amounts of time, the probability of ruin might go to zero but this is not the case. When $p > 0.5$, then $\frac{1-p}{p} < 1$ so the denominator in equation 7 tends towards -1 as T gets larger. This means that for large T the probability of success tends towards

$$P_n \approx 1 - \left(\frac{1-p}{p}\right)^n \tag{8}$$

For example, if $n = 100$ and the probability of winning is $p = 0.505$—implying an expected profit on a one unit stake of 0.01—then for large values of target wealth, the probability of success tends to

$$P_n \approx 1 - \left(\frac{0.495}{0.505}\right)^{100} = 0.865. \tag{9}$$

In other words, no matter how high the target wealth is, there is still a 13.5% chance of being ruined when staking 1% of your initial wealth each time on this favorable game.

We will also be interested in the expected final amount of wealth and the duration of play. The expected terminal wealth is simply the probability of reaching the target, multiplied by the target wealth

$$E\left(W_{D_n}\right) = P_n T = \left(\frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^T - 1}\right) T \tag{10}$$

The expected duration of the game starting from $W_0 = n$ can be calculated from Wald's first identity, which states that for independent identically distributed variables X_i and a randomly distributed natural number N

$$E\left(\sum_{i=1}^{N} X_i\right) = E\left(X_i\right) E\left(N\right) \tag{11}$$

In this case, $E(X_i) = 2p - 1$, so the expected duration of play can be calculated from

$$E(W_{D_n} - n) = (2p-1)E(D_n) \Longrightarrow E(D_n) = \frac{1}{2p-1}\left(\left(\frac{\left(\frac{1-p}{p}\right)^n - 1}{\left(\frac{1-p}{p}\right)^T - 1}\right)T - n\right) \tag{12}$$

2.2. Zero profit games

For the game with zero expected profit ($p = 0.5$), the characteristic equation has a double root at one. This means there is a different kind of solution, which takes the form

$$P_n = An + B \tag{13}$$

The boundary conditions $P_0 = 0$ and $P_T = 1$ imply the solution is

$$P_n = \frac{n}{T} \tag{14}$$

so the probability of success is the ratio of starting wealth to target wealth. Again, as $T \to \infty$ the probability of success $P_n \to 0$. Even though the expected profit and loss from the game is zero, the investor ends up losing all their money. This is because their wealth follows a random walk without drift and eventually it will reach zero with probability one.

An intuitive way to understand this formula is the optional stopping theorem which states that for any martingale sequence bounded by a stopping rule, the expected value of the martingale at the stopping time equals its initial value.[3] In this case, the wealth sequence is a martingale so

$$E(W_{D_n}) = P_n T + (1 - P_n)(0) = n \Rightarrow P_n = \frac{n}{T} \tag{15}$$

The zero expected profit means it is not possible to use Wald's first identity to calculate the expected duration of the game. Instead, we can use Wald's second identity, which states that for independent identically distributed random variables X_i and a randomly distributed natural number N

$$E\left[\left(\sum_{i=1}^{N} X_i\right)^2\right] = E\left(X_i^2\right)E(N) \tag{16}$$

For the zero expected-profit game, $E\left(X_i^2\right) = 1$, so expected duration can be calculated as

$$E(D_n) = E(W_{D_n} - n)^2 = P_n(T-n)^2 + (1-P_n)n^2 = n(T-n) \tag{17}$$

The expected duration of the game equals the initial wealth times the potential gain.

[3] This theorem was first proved by Doob (1953), page 300.

3. Asymmetric Payoffs

We now consider the case where the profit from each round of a game takes the form

$$X_i = \begin{cases} K & \text{with probability } \ p \\ -1 & \text{with probability } \ 1-p \end{cases} \tag{18}$$

where K is a positive integer. The investor again sets a target wealth level of T but in this case it is possible for them to end up with wealth of up to $T+K-1$. This means there are $K+1$ possible endings to the game: Wealth could equal zero, reach the target T or it could end up at one of the $K-1$ feasible values of wealth that are above the target. We will write the set of possible outcomes for wealth as $S = \{0, T, T+1, ..., T+K-1\}$.

Feller (1950) and Harper and Ross (2005) also considered cases where the gambler could lose more than one unit or gain multiple different positive amounts. Here we restrict ourselves to the two-outcome case to retain simple comparisons with the traditional gambler's ruin problem and to allow for a full general characterization of the properties of the solutions.

3.1. Solution method for games with positive or negative expected profit

Each round of the game has zero expected profit if $p = \frac{1}{K+1}$. We will first consider the case $p \neq \frac{1}{K+1}$. For each $j \in S$ define

$$P_n^j = \text{Prob}\left(W_{D_n} = j \mid W_0 = n\right) \tag{19}$$

All $K+1$ of these sets of probabilities can be characterized by a difference equation of the form

$$pP_{n+K}^j - P_n^j + (1-p)\, P_{n-1}^j = 0 \tag{20}$$

which implies the characteristic equation

$$pr^{K+1} - r + (1-p) = 0 \tag{21}$$

For all of the values of K considered here, calculations showed that, for $p \neq \frac{1}{K+1}$, this equation has $K+1$ distinct roots, r_i where $i = 1, 2, ..., K+1$. It can be easily seen that one of the roots equals 1. So, for each possible outcome, the general solution is of the form

$$P_n^j = A_1^j r_1^n + A_2^j r_2^n + ... + A_K^j r_K^n + A_{K+1}^j r_{K+1}^n \tag{22}$$

and specific solutions can be found from the boundary conditions for $j \in S$ of the form

$$P_n^j = \begin{cases} 1 & \text{if } \ n = j \\ 0 & \forall x \in S, x \neq j \end{cases} \tag{23}$$

These probabilities can be obtained by solving for the coefficients of the specific solution separately for the $K+1$ different sets of boundary conditions. However, the underlying difference equation is the same for each case. Harper and Ross (2005) showed this feature could be exploited using a computationally convenient method to solve for all $K+1$ sets of coefficients simultaneously using matrix algebra. This works as follows. Define the $(K+1) \times (K+1)$ matrix of coefficients A, with entries $a_{ij} = A_i^j$. We define i to index the root (associated with exponent base r_i) and j to index the outcome $j \in S$. Also define the $(K+1) \times (K+1)$ matrix D such that $d_{1j} = 1$ and $D_{ij} = r_j^{T+i-2}$ for $i > 1$. Then the coefficients can be obtained as the solution to

$$DA = I_{k+1} \Longrightarrow A = D^{-1} \tag{24}$$

To give a concrete example, suppose $K = 2$ and $T = 8$, then the boundary conditions can be written as

$$\begin{pmatrix} 1 & 1 & 1 \\ r_1^8 & r_2^8 & r_3^8 \\ r_1^9 & r_2^9 & r_3^9 \end{pmatrix} \begin{pmatrix} A_1^0 & A_1^8 & A_1^9 \\ A_2^0 & A_2^8 & A_2^9 \\ A_3^0 & A_3^8 & A_3^9 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

For each value of initial wealth n, this allows us to calculate the probabilities $P_n^0, P_n^T, ..., P_n^{T+K-1}$ that the investor ends up in each of the $K+1$ possible end states and from this to calculate the expected final wealth as

$$E\left(W_{D_n}\right) = \sum_{v=0}^{K-1} P_n^{T+v}\left(T+v\right) \tag{26}$$

The expected profit for each round of the game is $E\left(X_i\right) = p\left(K+1\right) - 1$ so Wald's identity implies the expected duration of the game is

$$E\left(D_n\right) = \frac{\left(\sum_{v=0}^{K-1} P_n^{T+v}\left(T+v-n\right) - nP_n^0\right)}{p\left(K+1\right)-1} \tag{27}$$

3.2. Solution method for games with zero expected profit

For the zero expected profit case of $p = \frac{1}{K+1}$, the solution method is again different. The function

$$F(r) = \frac{1}{K+1} r^{K+1} - r + \left(1 - \frac{1}{K+1}\right) \tag{28}$$

has the property that $F(1) = 0$ and $F'(1) = 0$ which is the condition for 1 to be a double root. For all of the values of K considered here, calculations showed that with $p = \frac{1}{K+1}$, the other $K-1$ roots are distinct. This means the general solution takes the form

$$P_n^j = A_1^j r_1^n + A_2^j r_2^n + ... + A_K^j r_K^n + n A_{K+1}^j \tag{29}$$

In this case, we can again calculate the coefficient matrix as $A = D^{-1}$ where the entries of D are

$$d_{1j} = \begin{cases} 1 & \text{if } \ j \neq K+1 \\ 0 & \text{if } \ j = K+1 \end{cases} \tag{30}$$

and for $i > 1$

$$d_{ij} = \begin{cases} r_j^{T+i-2} & \text{if } \ j \neq K+1 \\ n r_j^{T+i-2} & \text{if } \ j = K+1 \end{cases} \tag{31}$$

Because the expected value of profits is zero, we again use Wald's second identity to calculate the expected duration of the game. The expected squared profit (which is also the variance) for each round is

$$E\left(X_i^2\right) = \frac{K^2}{K+1} + \left(1 - \frac{1}{K+1}\right)(-1)^2 = K \tag{32}$$

and the expected square of the final profit is

$$E\left(W_{D_n} - n\right)^2 = P_n^0\left(n^2\right) + \sum_{v=0}^{K-1} P_n^{T+v}\left(T + v - n\right)^2 \tag{33}$$

so the expected duration can be calculated as

$$E\left(D_n\right) = \frac{E\left(W_{D_n} - n\right)^2}{E\left(X_i^2\right)} = \frac{1}{K}\left[P_n^0\left(n^2\right) + \sum_{v=0}^{K-1} P_n^{T+v}\left(T + v - n\right)^2\right] \tag{34}$$

3.3. Markov Chain method

The solution methods just described were how the long-run probabilities in this paper were calculated. Because these methods sometimes involve inverting large matrices including entries that are high powers of complex numbers (many of the roots for these problems are complex), there is the

potential for there to be numerical inaccuracies. As a check on our results, we also calculated the probabilities by expressing the models as Markov chains and raising the transition matrices to very higher powers. For example, for $K = 2$ and $T = 4$, wealth process can be expressed as a Markov Chain with transition matrix

$$M = \begin{array}{c} \\ 0 \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 1 & 2 & 3 & 4 & 5 \end{array} \\ \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1-p & 0 & 0 & p & 0 & 0 \\ 0 & 1-p & 0 & 0 & p & 0 \\ 0 & 0 & 1-p & 0 & 0 & p \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \end{array} \tag{35}$$

where the rows indicate the value of wealth at time t and the columns show the probabilities of being in the other states at time $t + 1$.

The numbers provided by the solution method described here were confirmed by the Markov Chain method and also by averages of final outcomes from simulations. Another benefit from writing the models as Markov chains is that we can characterize the speed of convergence to the ultimate outcome probabilities: The matrix M^h tells us the probabilities of being in the various states from each possible starting point after h periods.

4. Properties of Games with Asymmetric Payoffs

Here we describe the properties of games with asymmetric payoffs, first looking at the zero expected profit case, then the case of positive expected profits and finally negative expected profits.

4.1. Zero expected profit

With zero expected profit and asymmetric payoffs, wealth is still a martingale so the optional stopping theorem again applies meaning

$$E\left(W_{D_n}\right) = \left(\sum_{v=0}^{K-1} P_n^{T+v}\left(T+v\right)\right) + \left(1 - P_n^0\right)(0) = n \tag{36}$$

Expected final wealth for these games is the same as in the symmetric case but the asymmetric games include cases where the investor ends up with more than T. For expected final wealth to still equal n, the probabilities of ruin now need to be higher to offset the higher expected final wealth contingent on winning.

Figure 1 provides some illustrations of how the probabilities P_n^{T+v} typically get smaller as v increases—reaching the target T is more likely than reaching the highest possible value $T+K-1$ because wealth is more likely to reach $T-K$ than it is to reach $T-1$. Given this pattern, a tractable lower bound approximation on the probability of success comes from assuming the probabilities of reaching all feasible successful outcomes are equal, meaning $P_n^{T+v} = \frac{P_n^S}{K}$ where P_n^S is the combined probability of any successful outcome starting from wealth of n. This lower bound can be calculated as follows:

$$\begin{aligned} n = E\left(W_{D_n}\right) &\leq \frac{P_n^S}{K}\left(\sum_{i=0}^{K-1}\left(T+i\right)\right) + \left(1-P_n^S\right)(0) && (37)\\ &\leq \frac{P_n^S}{K}\left(KT + \frac{(K)\,(K-1)}{2}\right) + \left(1-P_n^S\right)(0) && (38)\\ \Rightarrow \quad P_n^S &\geq \frac{n}{T+\frac{K-1}{2}} && (39) \end{aligned}$$

Because the probability of success has to be lower than for the symmetric problem, we can put both upper and lower bounds on it as follows

$$\frac{n}{T+\frac{K-1}{2}} \leq P_n^S \leq \frac{n}{T} \tag{40}$$

These bounds imply that as long as K is small relative to T, the ruin rate formula derived in the symmetric case will work well.[4] In the case where K is high relative to T, the potential deviations

[4]Harper and Ross (2005) note that for the case $K=2$, $n=5$ and $T=9$, the probability of ruin is pretty close to $\frac{4.5}{9.5}$

from the standard ruin formula will rise as n gets bigger.

Figure 2 shows ruin probabilities as a function of initial wealth for two different values of the target, $T = 150$ and $T = 1000$, for a range of probabilities of success p, each associated with a different value of K so that $p = \frac{1}{K+1}$. The upper panel (for $T = 150$) shows deviations from the symmetric case formula can be large when K is high relative to T and n is also close to T. For example, the probability of ruin when starting from $n = 140$ with symmetric payoffs ($K = 1$) is 0.075. When $K = 10$, the ruin probability becomes 0.093; when $K = 50$ it is 0.156; when $K = 100$ it is 0.306. For the higher value of $T = 1000$ displayed in the bottom chart of Figure 2, the range of ruin probabilities for even high values of n is much smaller.

The intuition for these results is fairly simple. When investors are close to their target level of wealth and the probability of winning a round of the game is relatively high, then they are very likely to reach their target, particularly if T is high so they are very far from ruin. However, when the payoff from a round of the game is high but occurs with a low probability, the investor that is close to target has a higher probability of going on a long losing streak and eventually being ruined. The methods used here can be used to replicate the numerical examples reported by Harper and Ross (2005) but while their examples generally showed ruin probabilities close to the symmetric case, we have seen here that there are some instances where this is not the case.

Figure 3 further illustrates the differences between the symmetric and asymmetric versions of zero expected profit games. It fixes the value of initial wealth at $n = 100$, so the investor is staking 1 percent of their wealth on each play, but varies both the target T and the probability of success in each round p. The largely flat lines in the upper panel show that the symmetric case formula for the probability of ruin is accurate until the probability of a winning a round falls below 0.1. The bottom panel shows that the expected duration of play falls as p declines. For the symmetric case of $K = 1$, $p = 0.5$ and $T = 400$, the expected duration of this game is $n\,(T - n) = 30,000$. Changing to $K = 99$ and thus $p = 0.01$, the expected duration is 934 periods.

The lower expected duration of the asymmetric games as K increases means it takes less time for ruin rates to converge on the rates that we have calculated analytically. The upper panel of Figure 4 uses Markov Chains to illustrate ruin rates for various lengths of game play for different values of the probability of success in each play for $n = 100$ and $T = 400$. Unlike the traditional "coin toss" game, which lasts a very long time in this case before either ruin or success occurs, zero expected-profit games with high values of K can approach their long-run ruin rates quite quickly. For example, for the game with $K = 99$, the ruin probability after 500 periods is 0.66, which is most of the way to the long-run ruin rate of 0.77. The lower panel shows the same calculations for the probability of success. These also converge faster for higher values of K, though in this example, fast rates of convergence

which comes from replacing T in the standard symmetric formula with the mid-point of the possible successful outcomes for wealth. This equals the upper bound of $\frac{T+\frac{K-1}{2}-n}{T+\frac{K-1}{2}}$ derived here.

generally only occur when p is less than 0.1. These results show the high probability of eventual ruin for games with zero expected profit changes from being almost a theoretical curiosity for some version of the symmetric games (is anyone really going to play a coin toss game 30,000 times?) to being relevant for the more asymmetric games.

The explanation for these results is relatively simple. Wealth follows a random walk with increments X_i. In this case, the increments have zero mean and variance K, so the variance of cumulated profits at time t is

$$\text{Var}\left(W_t - W_0\right) = Kt \tag{41}$$

For any specific value of t, a higher K means the distribution of wealth has lower probability frequencies for values close to the mean of n and more weight in the tails, so outcomes like $W_t = 0$ and $W_t = T$ become more likely. The calculations above show that the relative probability of eventually reaching zero or T is generally not much affected by the amount that can be won in each round but the probability of reaching them by any specific point of time increases with K.

Figure 1: Probabilities of reaching successful wealth amounts T to $T+K-1$ for initial wealth $n = 100$ for two values of T and various values of the profit from winning a round, K

(a) $T = 300$

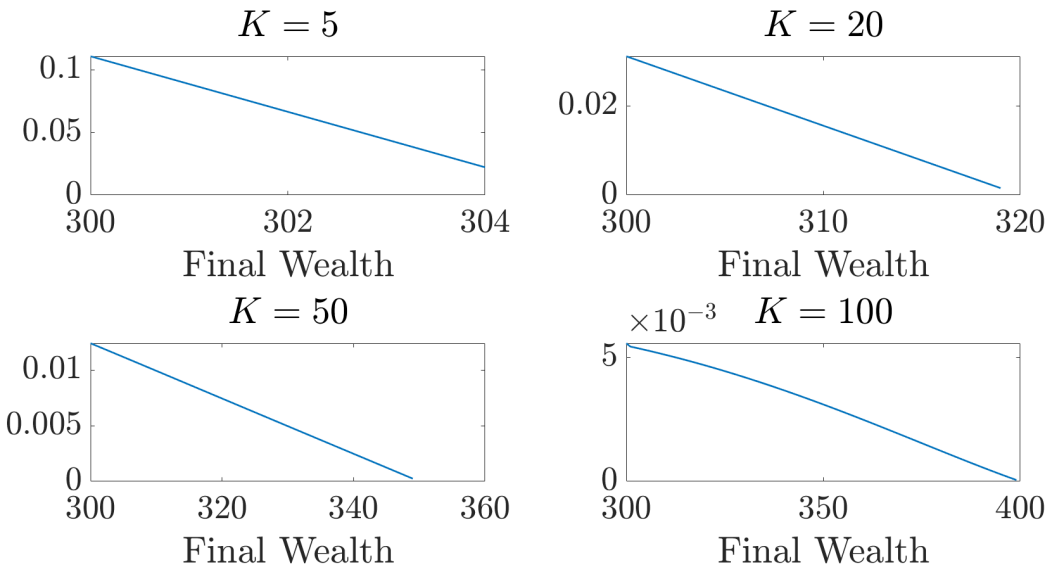

(b) $T = 600$

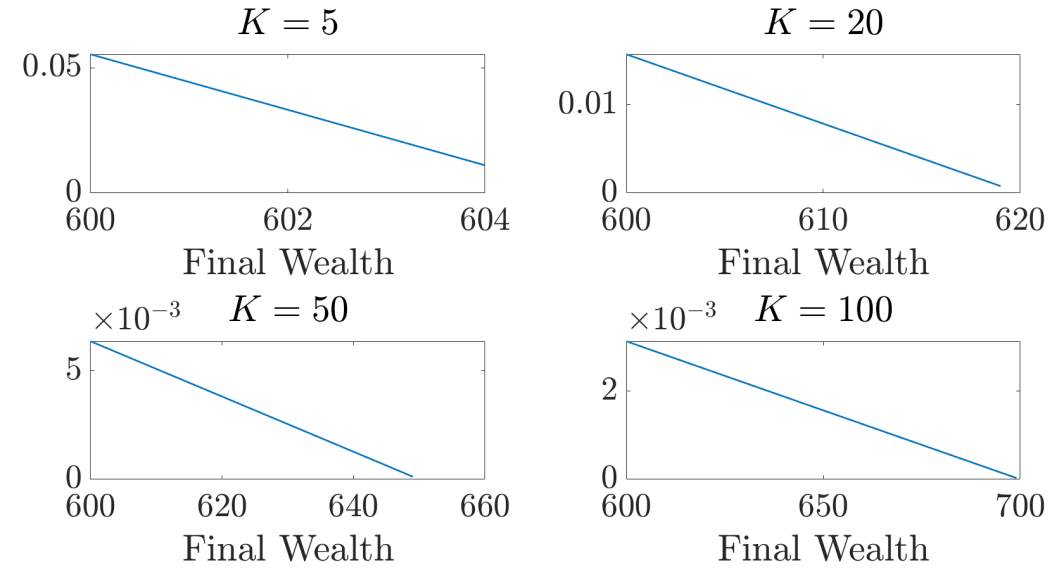

Figure 2: Comparing ruin probabilities for two values of target wealth T and various values of the winning profit K with initial wealth of $n = 100$

(a) $T = 150$

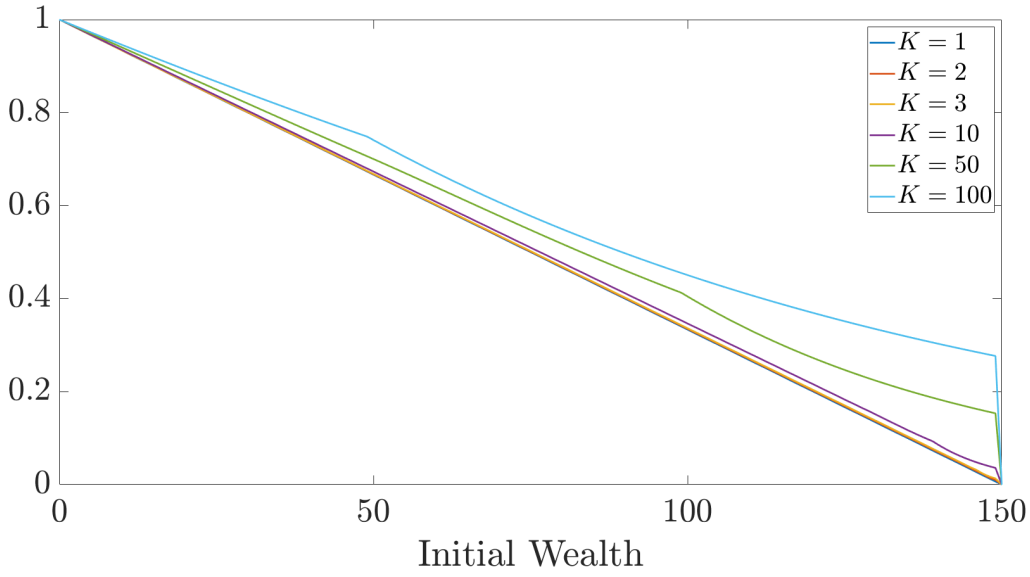

(b) $T = 1000$

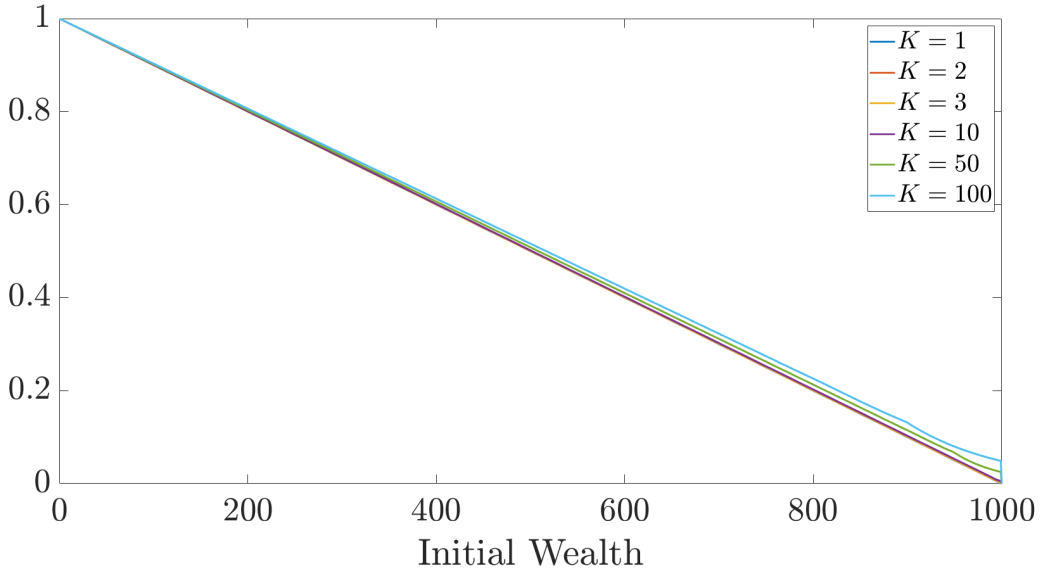

Figure 3: Probability of ruin and expected duration of the play for games with zero expected profit with initial wealth $n = 100$ with different values of the win probability per play (p) and target wealth (T)

(a) Probability of Ruin

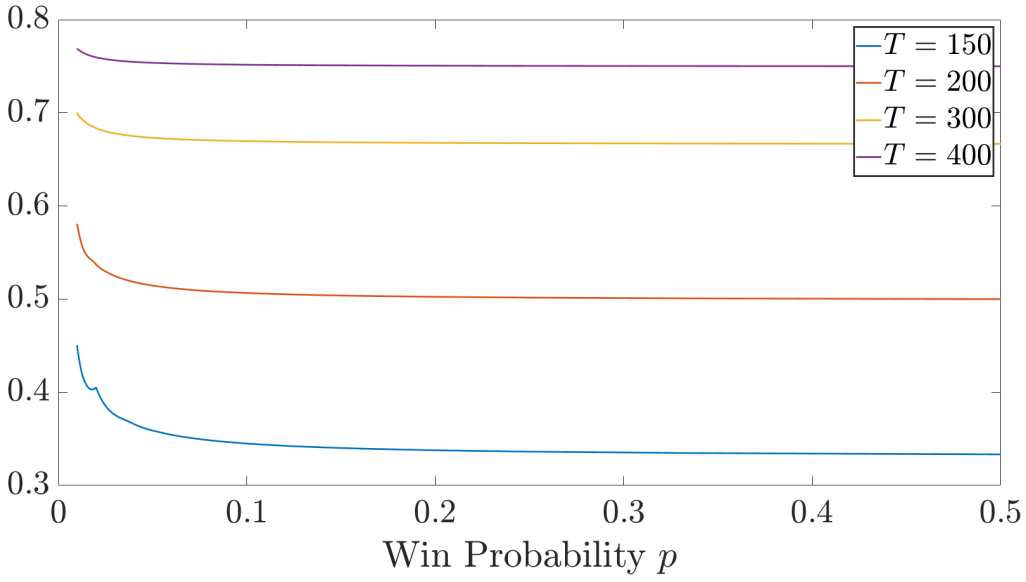

(b) Expected Duration

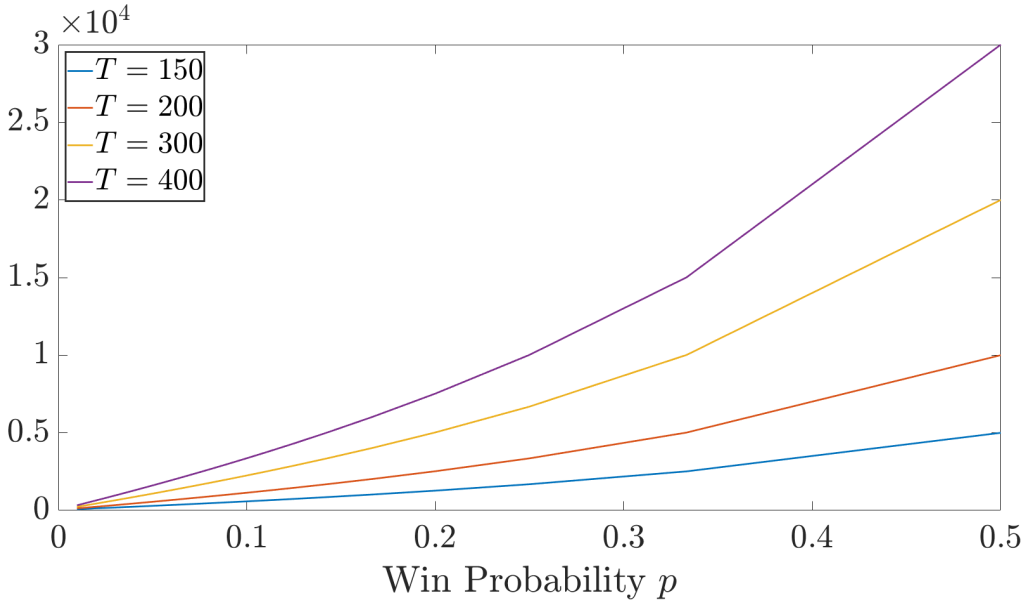

Figure 4: Ruin and success probabilities for games with zero expected profit after certain amounts of time for $n = 100$, $T = 400$ and various values of the win probability per play (p)

(a) Probabilities of Ruin

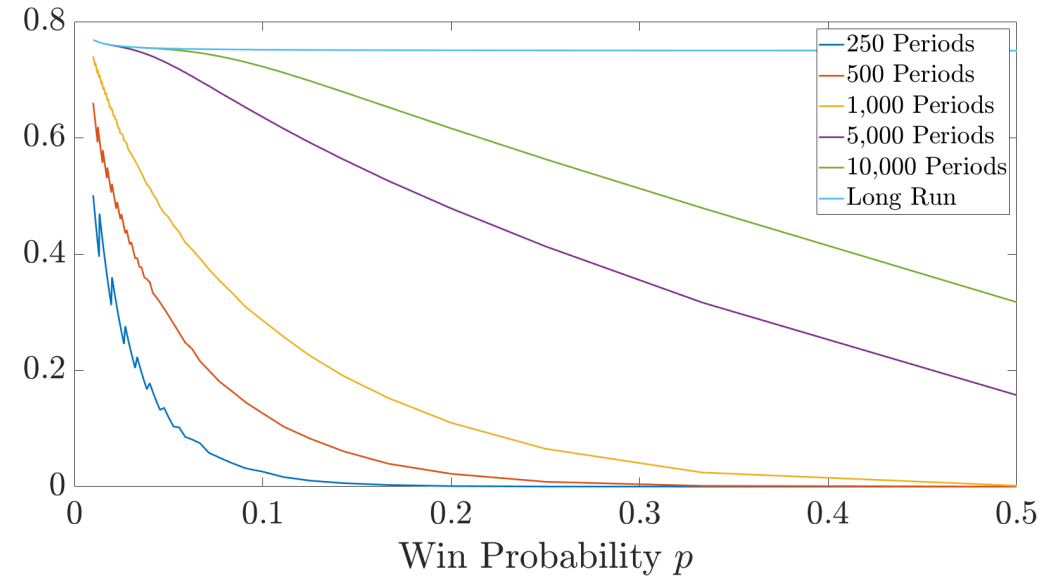

(b) Probabilities of Success

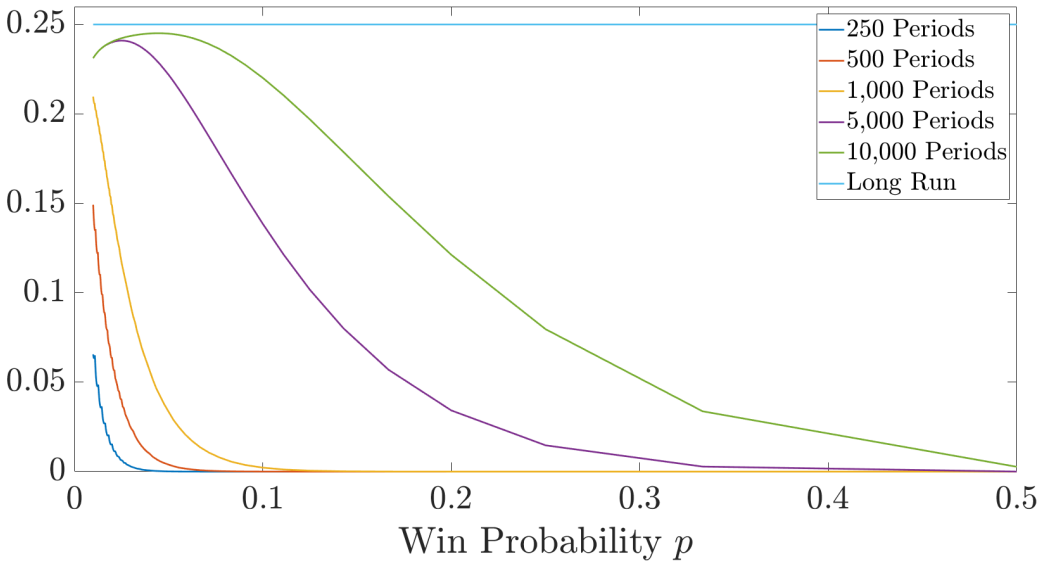

4.2. Positive expected profit

We now consider the case where the game has a positive expected profit. We calibrate this by assuming that

$$p = \frac{1+\mu}{K+1} \tag{42}$$

which gives an expected profit per round of $E(X_i) = \mu$. This means that as long as the game continues, wealth follows a random walk with positive drift of μ. In this case, wealth is a sub-martingale and a variant of the optional stopping theorem tells us that $E(W_t) > n$, which means the ruin rate will be lower and expected final wealth higher than when the expected return is zero.[5]

Figure 5 illustrates outcomes with $\mu = 0.01$ and $n = 100$ for different values of K (and thus p) and for different values of T. The upper panel shows that, for each value of T, the ruin rate depends negatively on p. The effect is large and does not only apply to games with low values of p. For $T = 300$, as we move from the coin toss situation of $K = 1$ to $K = 2$, the probability of ruin rises from 0.13 to 0.34. By $K = 20$, the ruin probability is 0.64. As K gets higher, the ruin probability tends towards the zero expected-profit rate. Even with a positive expected return, these very high risk investments result in ruin almost as often as games with zero expected profits.

The middle panel shows the ratio of final expected wealth to initial wealth. The higher ruin rates for bigger values of K translate into lower expected final wealth, so the possibilities of final wealth levels above T do not offset the higher chance of ending up with zero. For the most extreme asymmetric games, expected final wealth is effectively the same as for games with zero expected profit, so the game's positive expected value does them almost no good in terms of their bottom line. The bottom panel shows, as expected, that expected duration also falls as K rises. Duration is longer for higher values of T and this longer time playing a game with a positive expected return means expected final wealth depends positively on T.

The key factor driving these results is again the variance of profits increasing with K. In this case, the variance of profits is

$$\text{Var}(X_i) = (1+\mu)(K+1) - 1 - \mu^2 \tag{43}$$

so the variance of cumulated profits is

$$\text{Var}(W_t - W_0) = \left((1+\mu)(K+1) - 1 - \mu^2\right) t \tag{44}$$

which again rises with K.

In the previous case of zero expected profit, if stopped sequences were allowed to continue, they were just as likely to decrease as increase, so the average wealth of these truncated paths was the same as the untruncated alternatives. However, in this case, uninterrupted wealth sequences have a

[5] Again, see Doob (1953), page 300 and onwards.

positive drift and getting stopped at $W_t = 0$ means investors miss out on the positive expected future profits generated by this drift which may have pushed them towards the wealth of T. This is why the higher variance profits associated with larger values of K produce lower expected final wealth.

There are both short-run and long-run patterns driving these results. In the short-run, the distribution of wealth with high values of K has fatter tails than for symmetric games. Figure 6 uses Markov Chains to show the expected distribution of wealth for various game lengths starting from $n = 100$ with a target of $T = 200$. For $t = 250$, the distribution of wealth from the coin toss case of $K = 1$ is essentially a bell curve and almost nobody has either reached target or ruin. In contrast, for $K = 20$, the distribution at $t = 250$ features about 20 percent of players reaching either target or ruin. At $t = 250$, the positive expected return of $\mu = 0.01$ does not have much impact on the relative probabilities of either success or ruin, so about equal amounts end up in the two outcomes. As t gets larger, the upward drift in the mean of the distribution of uninterrupted wealth sequences means that the vast majority of outcomes for the symmetric case end up with the investors reaching their target but as K rises, the share of successful outcomes declines.

These outcomes are not merely artifacts of short-run dynamics. In the long-run, the distributions of wealth outcomes end up with all weight at either zero or at T and above. But one might imagine that, for high values of t, the distribution of an uninterrupted wealth series moves sufficiently to the right to all but rule out ruin outcomes once t gets large and there has not been a stop. This is not necessarily the case. The X_i are independent identically distributed series so the Lindberg-Levy Central Limit Theorem would hold for uninterrupted wealth sequences, albeit convergence would be relatively slow for high values of K. The asymptotic distribution of the mean of profits for such uninterrupted sequences is

$$\frac{1}{t}\sum_{i=1}^{t} X_i \overset{a}{\sim} N\left(\mu, \sqrt{\frac{(1+\mu)(K+1) - 1 - \mu^2}{t}}\right) \tag{45}$$

However, the level of wealth of uninterrupted sequences depends not on the mean of profits but on the accumulated sum and this has asymptotic distribution

$$W_t - W_0 \overset{a}{\sim} N\left(\mu t, \sqrt{[(1+\mu)(K+1) - 1 - \mu^2]\, t}\right) \tag{46}$$

This asymptotic distribution of uninterrupted cumulated profits has a mean that increases multiplicatively with t and a standard deviation that increases with the square root of t. Higher values of t increase the asymptotic mean of $W_t - W_0$ which reduces the probability of cumulated profits being below $-W_0$. But higher t also raises the asymptotic standard deviation which increases the probability of cumulated profits being below $-W_0$, with this latter effect increasing with the size of K. For large enough values of t, the mean effect wins out over the standard deviation effect and the probability of wealth going below W_0 goes to zero. But this takes an extremely long time to happen,

particularly for high values of K.

Figure 7 illustrates this by showing the probability of having lost at least $W_0 = 100$ for uninterrupted wealth series drawn from $N\left(\mu t, \sqrt{[(1+\mu)(K+1) - 1 - \mu^2]\, t}\right)$ distributions for various values of t and K with $\mu = 0.01$. This probability rises at first as t increases, reaching a peak at $t = 10,000$ in this case, before declining slowly towards zero. The peak of the probability of losing at least 100 in these uninterrupted sequences is higher for higher values of K and the reduction towards zero is slower.

The peak value of t for these probabilities is the same for each value of K. This is because we are charting the cumulative distribution of $W_t - n$ at $-n$ which depends negatively on the z-score $\frac{-n-\mu t}{\sqrt{[(1+\mu)(K+1)-1-\mu^2]t}}$. The term $\frac{1}{\sqrt{(1+\mu)(K+1)-1-\mu^2}}$ multiplies this value but does not have an impact on which value of t attains the minimum value. One can show this minimum value occurs at $t = \frac{n}{\mu}$, which is 10,000 in the case shown in Figure 7.

Of course, these calculations do not show the probability of ruin once stopping is incorporated. Some of the uninterrupted sequences that reach $W_t = 0$ considered in this case would have paths that have previously gone below zero or above T, so these distributions are not the same as for uninterrupted sequences when stopping rules are in place. The distribution of uninterrupted sequences once stopping is incorporated is instead illustrated by the non-extreme values of the distributions shown in Figure 6. But these calculations show that even without incorporating a stopping rule, it is quite possible to lose all your money placing small fractions of it (such as the 1 percent stake assumed here) on high variance investments, even when playing games with a 1 percent expected return for a very long time.

These results relate to the literature on the Kelly criterion (Kelly, 1956, Breiman, 1961) for the optimal strategy when a gambler has an edge. The Kelly criterion predicts that the log of expected wealth is maximized by setting the share of wealth allocated to a gamble equal to "edge over odds" where the edge is the expected return of the gamble and the odds are the fractional odds equivalent of the payoff (K in our terminology).[6] This means that in our case, the Kelly criterion predicts the share of wealth allocated to gambles with an expected return of μ and a winning profit of K should be $\frac{\mu}{K}$. So, in the case of $n = 100$, an initial stake of 1 percent of wealth is consistent with the Kelly criterion but its suggested stakes for gambles with higher values of K are lower. This is consistent with the relatively poor outcomes illustrated in Figure 5 from staking 1 percent of wealth with $\mu = 0.01$ on games with high values of K. For these values of K, the shares of wealth staked here (1 percent to start, on average a bit less afterwards) are on the "over betting" side of the Kelly-optimal strategy.

[6] See Whelan (2025) for a compact derivation of this result.

Figure 5: Outcomes for games with positive expected profit ($\mu = 0.01$) with different values of the win probability per play (p) and $n = 100$

(a) Probability of Ruin

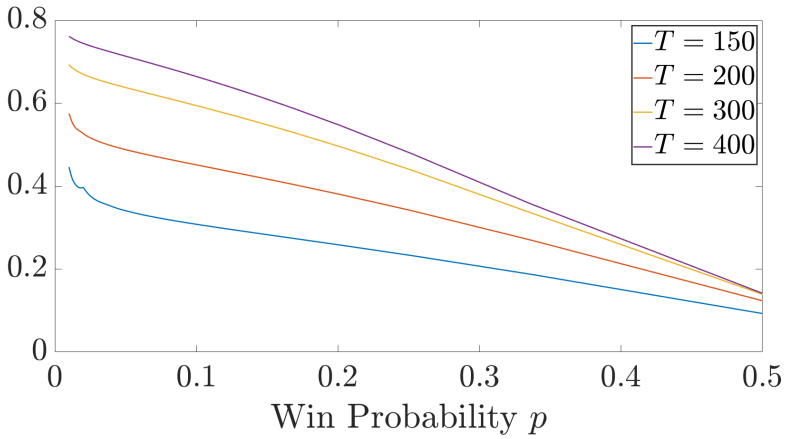

(b) Expected Ratio of Final Wealth to Initial Wealth

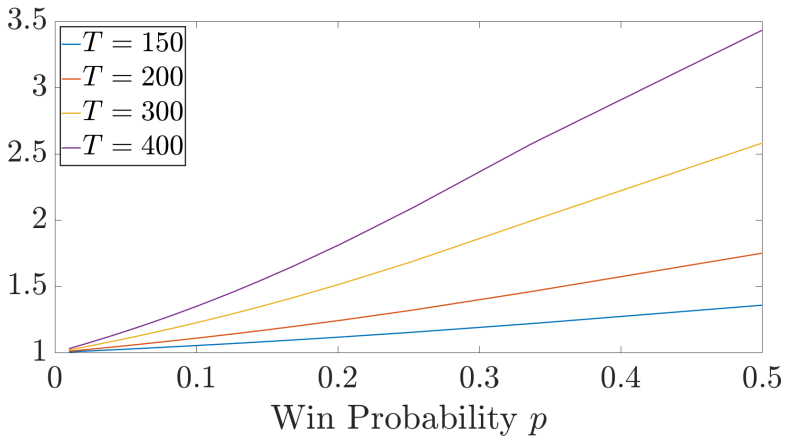

(c) Expected Duration

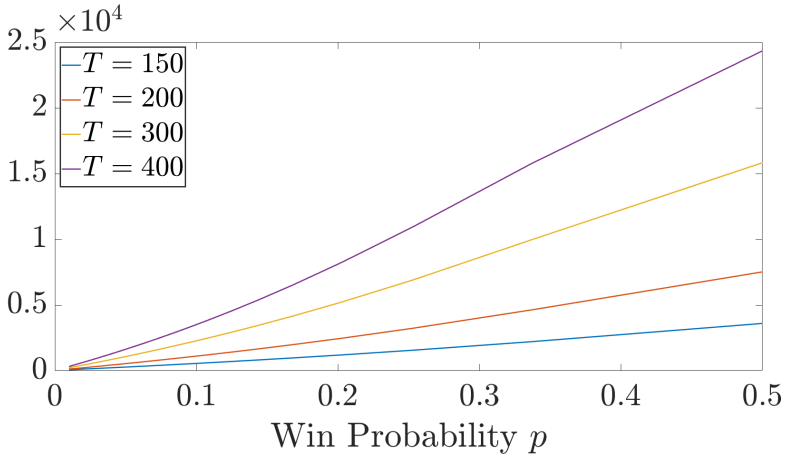

Figure 6: Distribution of outcomes for different lengths of game play with positive expected profit ($\mu = 0.01$) with different values of winning payoff K, initial wealth $n = 100$ and target wealth $T = 200$

(a) $K = 1$

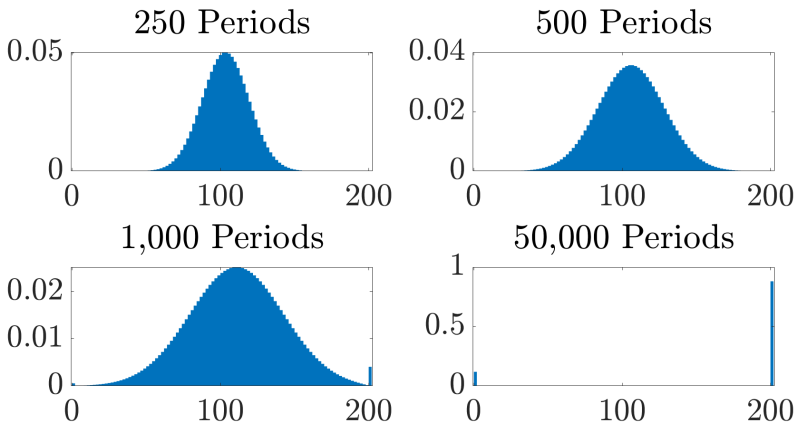

(b) $K = 5$

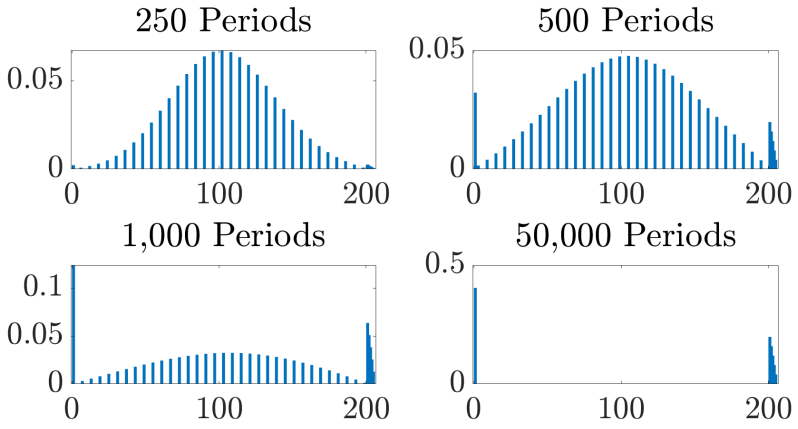

(c) $K = 20$

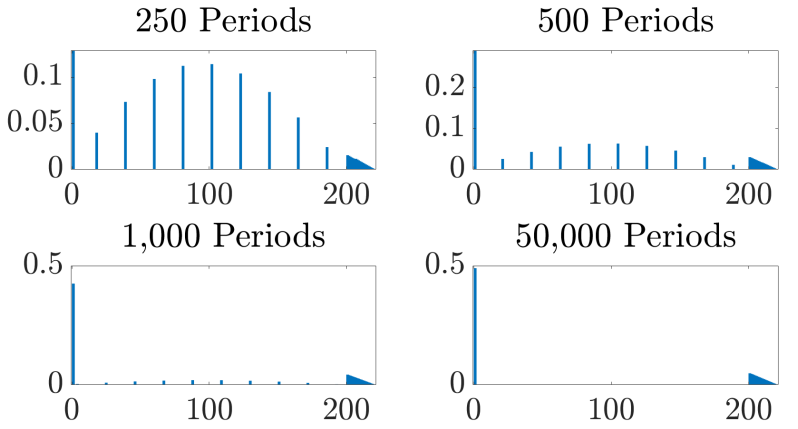

Figure 7: Probabilities of $W_t - W_0 \leq -100$ for values of $W_t - W_0$ drawn from $N\left(\mu t, \sqrt{\left[(1+\mu)(K+1) - 1 - \mu^2\right] t}\right)$ distributions for various values of t and K and an expected return of $\mu = 0.01$

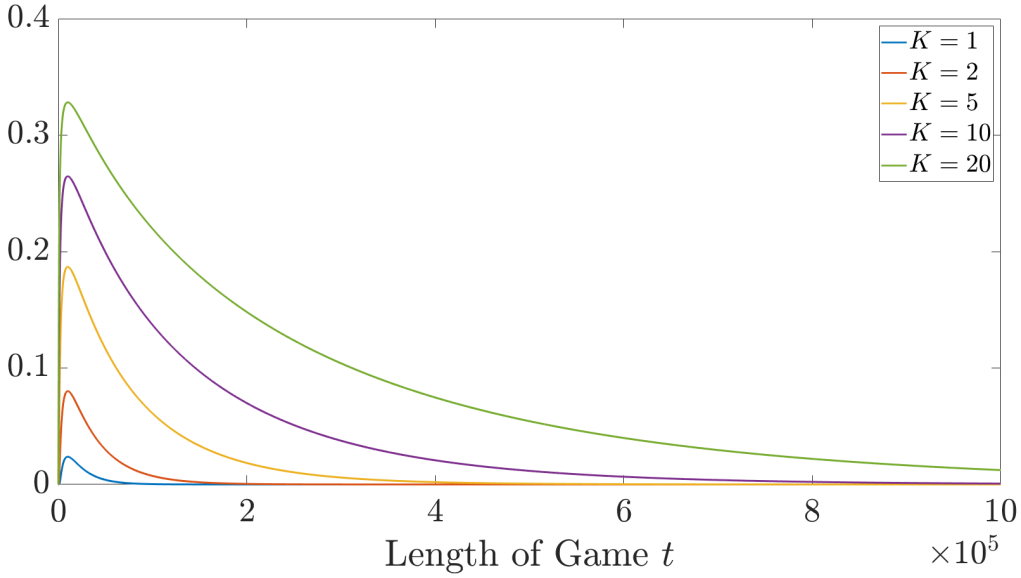

4.3. Negative expected profit

Figure 8 repeats our previous analysis with $n = 100$ but this time for $\mu = -0.01$, so the game has a negative expected return. The upper and lower panels show that ruin rates and durations are higher as target wealth T rises. In contrast to the positive expected return case, ruin rates generally fall as p declines, with only a slight uptick for the very lowest rates shown here (the lowest shown is $p = 0.01$). Expected final wealth (shown in the middle panel) follows the pattern suggested by the ruin rates with the highest values being obtained for the lowest values of p (highest values of K). In this case, investors do better taking high-variance investments with low probabilities of success.

The explanations for these results are similar to those for positive-value games. Figure 9 repeats the analysis of distributions over time, this time for $\mu = -0.01$. The bottom chart, for $K = 20$, looks very similar to the same chart in Figure 6 so the outcomes with $\mu = -0.01$ are fairly similar to those with $\mu = 0.01$. The difference in μ values has huge implications for very long sequences of games but makes little substantive difference to short-run outcomes when K is large. And since the majority of outcomes for these games are settled before the positive or negative expected value has made a big difference to the mean of wealth for unstopped sequences, the final outcomes are also very similar. For $K = 20$, $n = 100$ and a target of $T = 300$, for $\mu = 0.01$, the ruin rate is 0.64 and the expected ratio of final wealth to initial wealth is 1.11. For $\mu = -0.01$, the ruin rate is 0.71 and the expected ratio of final wealth to initial wealth is 0.9. For more extreme values of K, there is almost no difference between the outcomes for the two different values of μ.

In contrast, for games with low values of K, expected duration is longer and so games rarely end with reaching the target early. Instead, games mostly end after the cumulative disadvantage of the negative expected return has a predominant influence. Repeatedly making negative expected return investments ends up badly, particularly if you have set a high target. For example, with $n = 100$ and $T = 400$ and $K = 1$, expected final wealth with $\mu = -0.01$ is just below 1 percent of initial wealth.

The longer duration of games with negative expected values when you choose "timid" play (i.e. low K in our case) is well known from Freedman (1967). However, Freedman's characterization of this result as "timid play is optimal" only applies to maximizing the length of the game rather than expected final wealth. And, pretty clearly, the optimal strategy in this case is simply to avoid these investments.

Figure 8: Outcomes for games with negative expected profit ($\mu = -0.01$) with different values of the win probability per play (p) and $n = 100$

(a) Probability of Ruin

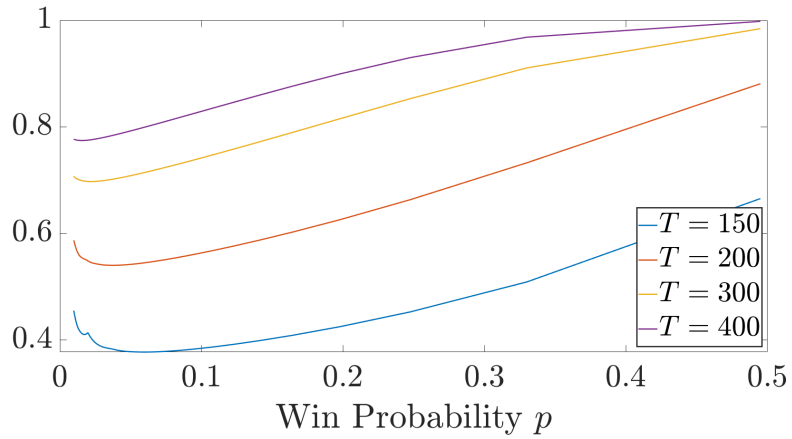

(b) Expected Ratio of Final Wealth to Initial Wealth

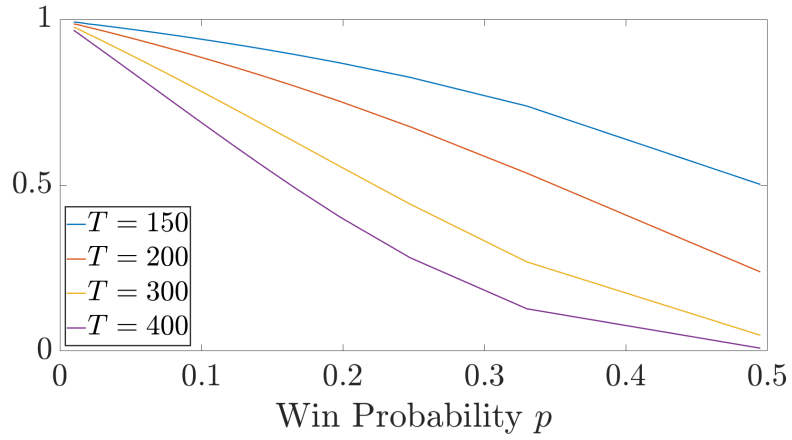

(c) Expected Duration

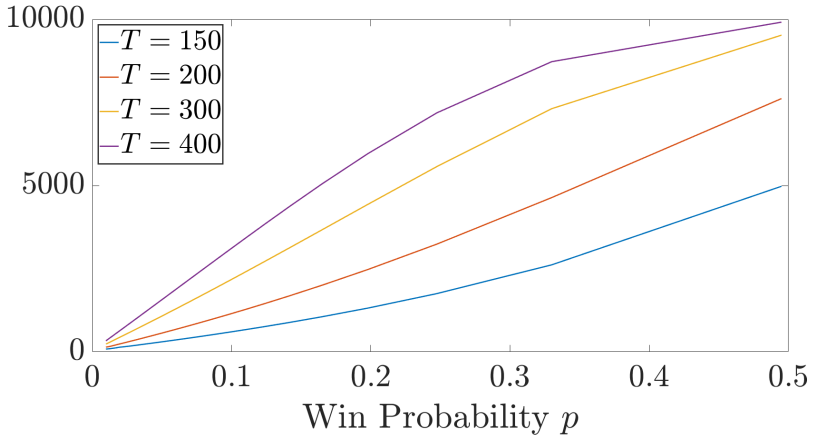

Figure 9: Distribution of outcomes for different lengths of game play with negative expected profit ($\mu = -0.01$) with different values of winning payoff K, initial wealth $n = 100$ and target wealth $T = 200$

(a) $K = 1$

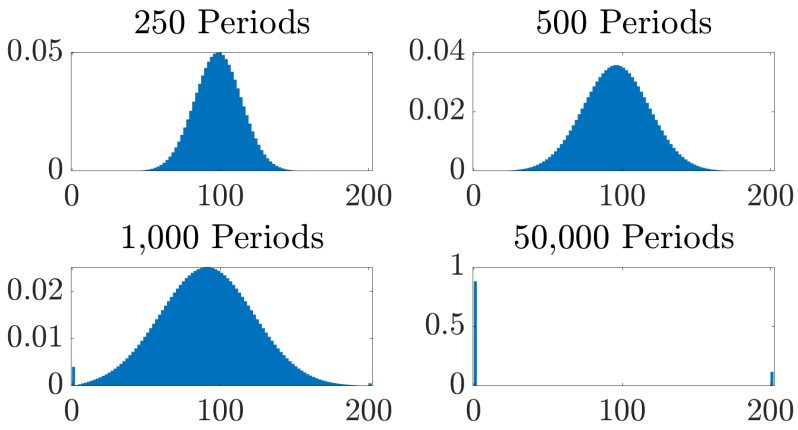

(b) $K = 5$

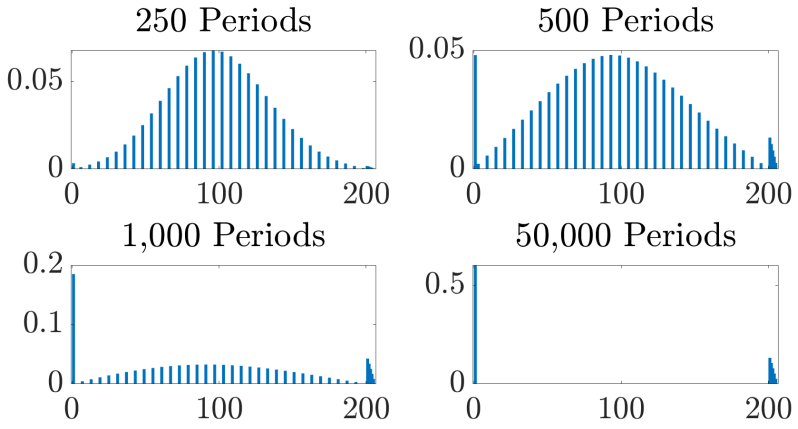

(c) $K = 20$

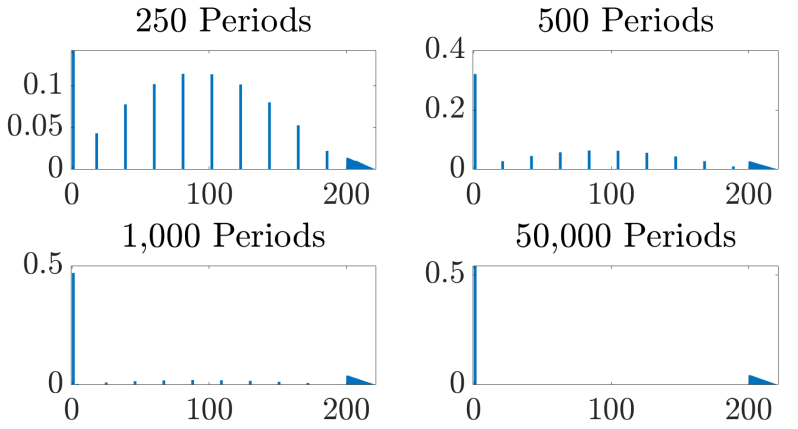

5. Varying Stake Size

So far, we examined a stake size of 1, meaning for example that the case of $n = 100$ could be understood as staking 1 percent of wealth. We now consider the case where the stake size is independent of initial wealth. We introduce both varying stake size and varying payoffs for unit bets by allowing the profit process to be

$$X_i = \begin{cases} Ks & \text{with probability } \ p \\ -s & \text{with probability } \ 1-p \end{cases} \tag{47}$$

where the different possible values of wealth are $0, s, 2s, ..., n-s, n, ..., T-s, T, T+(K-1)s$ where n and T are assumed to be integer multiples of s.

The symmetric case $K = 1$ was analyzed by Feller (1950). The probability of ruin for $p = 0.5$ is as before. For $p \neq 0.5$, the probability of success is

$$P_n = \frac{\left(\frac{1-p}{p}\right)^{\frac{n}{s}} - 1}{\left(\frac{1-p}{p}\right)^{\frac{T}{s}} - 1} = \frac{z^{\frac{n}{s}} - 1}{z^{\frac{T}{s}} - 1} \tag{48}$$

where $z = \frac{1-p}{p}$.

When $p > 0.5$ we have $z > 1$. Increasing s reduces $z^{\frac{n}{s}}$ but reduces $z^{\frac{T}{s}}$ by more because $T > n$. Thus, increasing s reduces the probability of success and raises the probability of ruin. To minimize the chance of ruin when the game has a positive expected profit, the best strategy is to repeatedly make small bets that minimize the chance of losing all your wealth due to bad luck. The reverse applies when $p < 0.5$ and thus $z < 1$. Increasing stakes reduces the probability of ruin. If you are playing a game where you are at a disadvantage, then repeated small stake bets just makes your opponent's win inevitable. The best strategy is to go big and raise the chance of a quick positive outcome. A formal proof of the optimality of higher stakes bets when $K = 1$ and the game has a negative expected profit can be found in Isaac (1999). The theme of "bold play" being the best strategy in games with negative expected profits was also discussed in a variety of contexts by Dubins and Savage (1965).

This finding that higher stakes produce worse outcomes when the game has a positive expected profit, and better outcomes when they do not, generalizes to higher values of K. The driving force behind these results is the same as the results on the effect of the size of a winning payoff. The variance of the expected profit when $E(X_i) = \mu$ is

$$\text{Var}(X_i) = s^2\left((1+\mu)(K+1) - 1 - \mu^2\right) \tag{49}$$

Both higher winning payoffs (K) and higher stakes (s) raise the variance of profits but the form of

these effects is different. The variance of profits depends on the square of s while the size of a winning payoff has a linear effect.

Figures 10 and 11 illustrate the impact of varying both the win probability per play (p) and the fraction of wealth being staked, first for $\mu = 0.01$ and then for $\mu = -0.01$ where the target is to multiply initial wealth by 5. As expected, the higher variance from placing larger stakes reduces expected duration. As also expected, higher stakes raises ruin rates and lowers expected final wealth when $\mu = 0.01$ and generally does the opposite when $\mu = -0.01$, with the notable exception of games with lowest values of p when $\mu = -0.01$ where the ruin rate at higher stakes ticks up more than when the stakes are low. The slightly different behavior of ruin rates for negative μ and low probabilities of success was evident earlier in the upper panel of Figure 8.

Interestingly, however, we can see that the impact of higher stakes gets smaller as the value of K rises. Focusing on the bottom line, expected final wealth, we can see that for the games closely resembling coin tosses, stake size has a huge impact on the expected outcome. However, stake size has little impact for the more extreme high values of K. This reflects a "diminishing marginal impact" of adding more variance. These charts illustrate that the stake size effect and winning prize effect on outcomes are different from each other and also interact in complex ways.

Figure 10: Outcomes for games with positive expected profit ($\mu = 0.01$) with different values of the win probability per play (p) where the target is to multiply wealth by 5

(a) Probability of Ruin

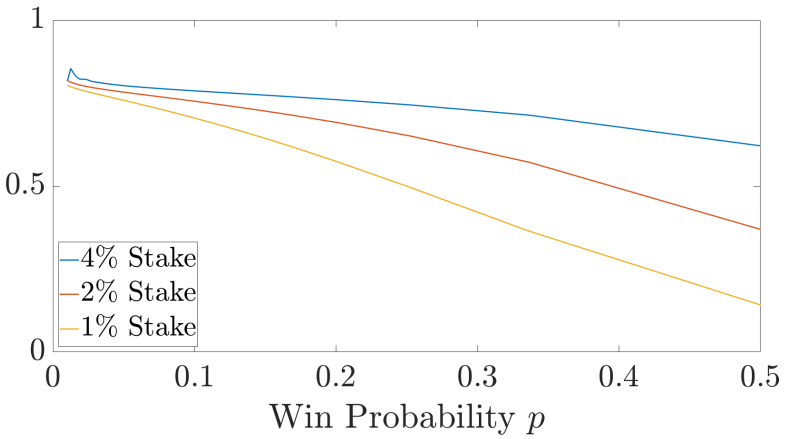

(b) Expected Ratio of Final Wealth to Initial Wealth

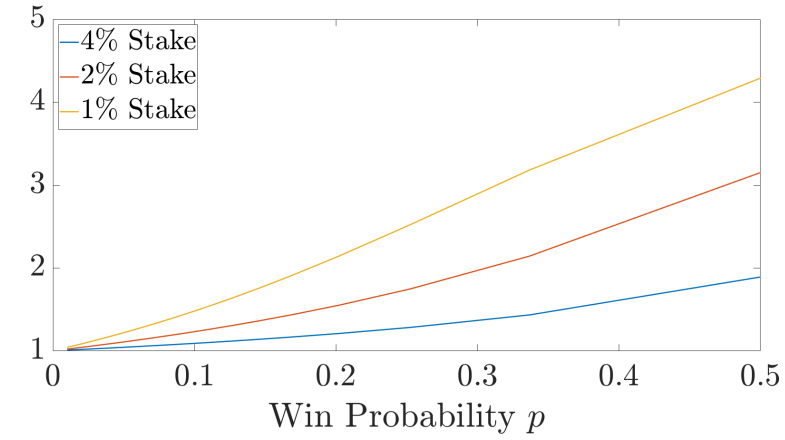

(c) Expected Duration

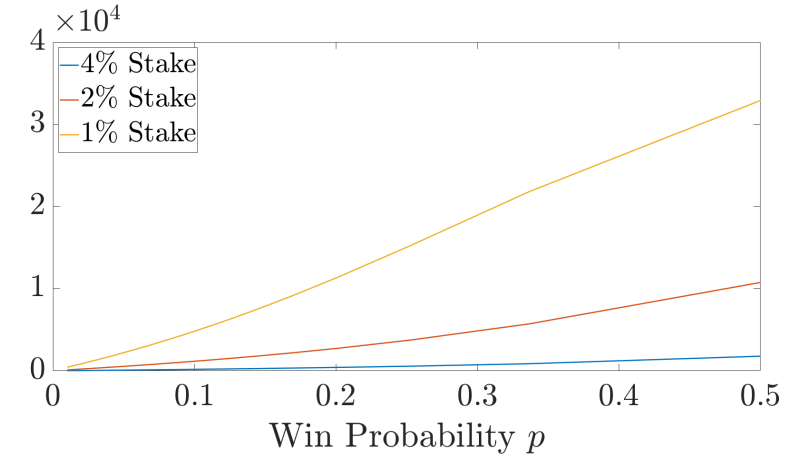

Figure 11: Outcomes for games with negative expected profit ($\mu = -0.01$) with different values of the win probability per play (p) where the target is to multiply wealth by 5

(a) Probability of Ruin

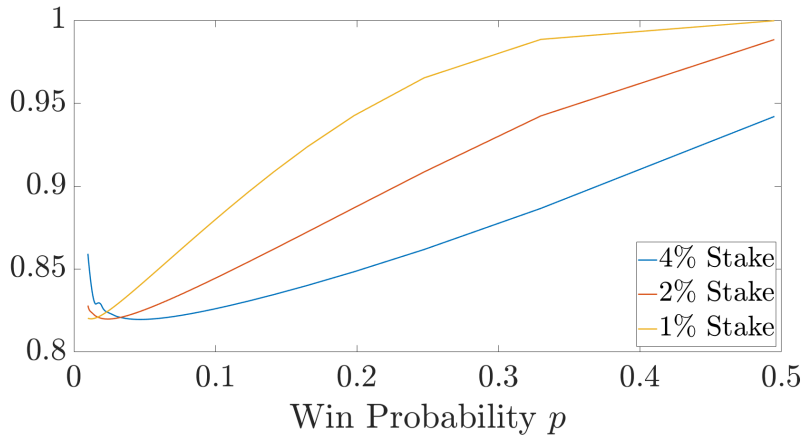

(b) Expected Ratio of Final Wealth to Initial Wealth

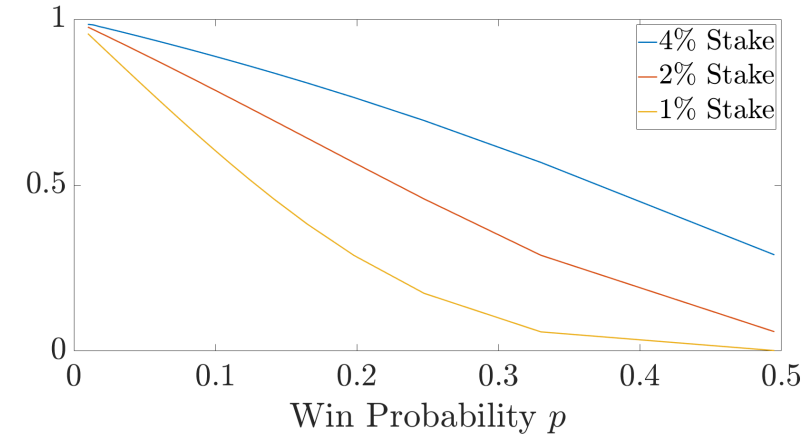

(c) Expected Duration

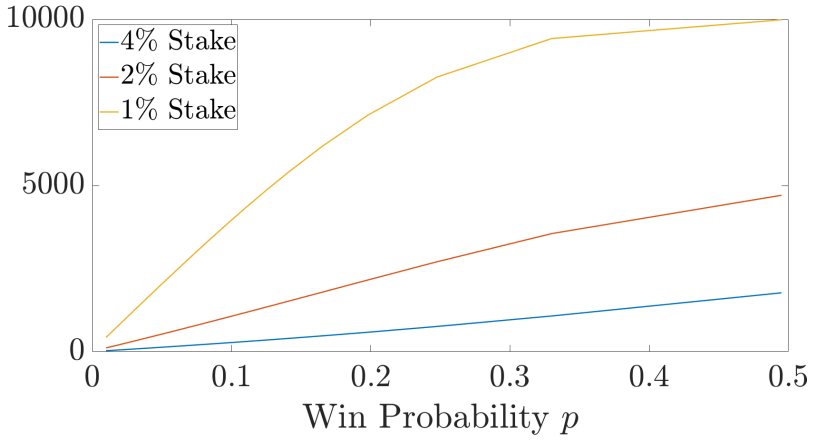

6. Conclusions

The gambler's ruin remains a foundational problem in probability theory, with enduring relevance far beyond its classical gambling origins. In this paper, we extended the traditional model to accommodate asymmetric payoffs, where the gains from success are large multiples of the stake at risk and occur with low probability. These types of payoff structures are common in modern finance—in options trading, early-stage investing, and tail-risk strategies—where investors repeatedly face decisions about capital allocation under skewed risk-reward profiles.

Our results show that even when expected returns are held constant, increasing asymmetry in payoffs can significantly alter the probability of ruin, the distribution of final wealth, and the speed at which outcomes are realized. These findings have direct implications for portfolio construction, risk management, and the design of investment strategies involving rare but high-impact outcomes.

In addition to providing new theoretical insights into the dynamics of ruin under asymmetric risk, the framework developed here offers a tractable and flexible tool for evaluating the sustainability of such strategies over time. Future work could extend this analysis by incorporating transaction costs, dynamic bet sizing, or empirical calibration to real-world investment strategies that rely on asymmetric return profiles.

Funding

No funding was received for this research.

Conflicts of Interest

The authors declare that they have no conflicts of interest in relation to this research.

References

[1] Breiman, L. (1961) Optimal Gambling Systems for Favorable Games. *Fourth Berkeley Symposium on Mathematics, Statistics and Probability*, 1, pages 65-78.

[2] Broadie, M. Chernov, M. Johannes, M. (2009). "Understanding Index Option Returns," *Review of Financial Studies*, Volume 22, pages 4493–4529.

[3] Doob, J.L. (1953). *Stochastic Processes*, Wiley.

[4] Dubins, L. and L.J. Savage (1965). *How to Gamble If You Must: Inequalities for Stochastic Processes*, Dover.

[5] Edwards, A.W.F. (1983). "Pascal's Problem: The 'Gambler's Ruin'," *International Statistical Review*, Volume 51, pages 73-79.

[6] Feller, W. (1950). *An Introduction to Probability Theory and Its Applications*, Wiley.

[7] Freedman, D. (1967). "Timid Play is Optimal," *Annals of Mathematical Statistics*, Volume 38, pages 1281-1283.

[8] Harper, J.D. and K.A. Ross (2005). "Stopping Strategies and The Gambler's Ruin," *Mathematics Magazine*, Volume 78, pages 255-268.

[9] Hunter, B., A.C. Krinik, C. Nguyen, J.M. Switkes, H.F. von Bremen (2008). "Gambler's Ruin with Catastrophes and Windfalls," *Journal of Statistical Theory and Practice*, Volume 2, pages 199-219.

[10] Ilmanen, A. Thapar, A. Tummala, H. Villalon, D. (2021). "Tail Risk Hedging: Contrasting Put and Trend Strategies," *Journal of Systematic Investing*, Volume 1, pages 111–124.

[11] Isaac, R. (1999). "Bold Play is Best: A Simple Proof," *Mathematics Magazine*, Volume 72, pages 405-407.

[12] Kelly, J.L. (1956). "A New Interpretation of Information Rate," *Bell System Technical Journal*, Volume 35, pages 917–926.

[13] Kerr, W.R. Nanda, R. and Rhodes-Kropf, M. (2014). "Entrepreneurship as Experimentation," *Journal of Economic Perspectives*, Volume 28, pages 25–48.

[14] Song, S. and J. Song (2013). "A Note on the History of the Gambler's Ruin Problem", *Communications for Statistical Applications and Methods*, Volume 20, pages 157-168.

[15] Takacs, L. (1969). "On Classical Ruin Problems," *Journal of the American Statistical Association*, Volume 64, pages 889-906.

[16] Whelan, K. (2025). "On Optimal Betting Strategies With Multiple Mutually Exclusive Outcomes," *Bulletin of Economic Research*, Volume 77, pages 67-85.