

The Kelly Criterion with an Estimated Edge

Harry Crane*

September 14, 2026

Abstract

The Kelly criterion prescribes the stake that maximizes the long run growth rate of bankroll when the bettor's edge is known. In most practical situations, the edge is not known, but estimated. Based on the estimated edge, many practitioners opt to stake a fraction of the amount prescribed by the Kelly criterion. Common choices for fractional Kelly range from 0.10 to 0.50 of the full Kelly stake, with the choice guided primarily by folklore and heuristics.

We reframe bet sizing as a constrained optimization depending on two parameters the bettor controls: (i) a *stop loss* E , the fraction of bankroll the bettor is willing to lose before abandoning the strategy, and (ii) a *premature stopping probability* π , the probability of hitting the stop loss even though the assumed edge is correct.

We show that when the edge is accurately estimated, the probability of hitting the stop loss is $(1 - E)^{\theta^*}$, where θ^* is the Cramér root of the per-round log-return. The constraint that this probability not exceed π determines the optimal stake in closed form: expressed as a fraction of the Kelly stake,

$$\bar{\alpha} = \frac{2 \log(1 - E)}{\log(1 - E) + \log \pi}$$

in the small-edge limit. Closed forms follow for the probability of reaching a target multiple of bankroll before the stop loss and for the degradation of the guarantee under an overestimated edge together with its robust correction.

Keywords: Kelly criterion, bet sizing, fractional Kelly, stop loss, drawdown, likelihood ratio, growth-optimal investment, risk management.

1 Introduction

In gambling, the bettor is prone to two types of error:

- (I) betting when there is no advantage, and
- (II) not betting when there is an advantage.

Error (I) arises when the bettor incorrectly believes that he has an advantage when he does not. Error (II) arises when the bettor incorrectly believes that he has no advantage. The primary challenge is to decide how much money to wager in order to achieve the bettor's objective while mitigating the above two sources of error.

The question of bet sizing is foundational to the theory of gambling [6, 8]. Under the Kelly criterion [12], the bettor risks the amount that maximizes the expected growth rate of bankroll.

*Department of Statistics, Rutgers University. Email: harry.crane@rutgers.edu.

The Kelly criterion is a widely cited theoretical solution to the problem of optimal bet sizing for a given edge and odds. The Kelly criterion, however, assumes that the bettor's edge is known, and therefore does not account for the possibility of errors (I) and (II). Few professional bettors use the Kelly criterion exactly as prescribed, because its assumption of a known edge is not realistic in practice.

The question Kelly addressed is how much a bettor should stake on a favorable wager. The practicing bettor faces a different question, namely how much to stake on a wager that is perceived to be favorable. In practice the edge is estimated, not known, and the bettor must account for this uncertainty when sizing his bet.

To address this uncertainty, bettors typically make *ad hoc* adjustments to the Kelly criterion, such as staking a fraction αf^* of bankroll, for $0 < \alpha < 1$ and f^* denoting the Kelly fraction under the estimated edge, rather than risking the full Kelly fraction f^* . Folklore suggests α in the range $0.10 \leq \alpha \leq 0.50$. But beyond this heuristic range, there is no established theoretical basis to guide the bettor in choosing α .

This paper proposes a framework in which the fractional value α is derived based on the bettor's preference for controlling the risk of errors (I) and (II) when betting with an uncertain edge. The framework presented here allows the bettor to mitigate the impact of both errors through parameters that are interpretable and within the bettor's control. Before discussing the framework in detail, we first review the classical setting in which the Kelly criterion is derived. Section 1.2 reviews Thorp's practical case for fractional Kelly, and Section 1.3 outlines the proposed framework and the contributions of this paper.

1.1 The classical setup and the Kelly criterion

A bettor with initial bankroll $B_0 > 0$ faces a sequence of identical wagers. Each wager pays b -to-1 ($b > 0$) on an event with true probability $p \in (0, 1)$ of occurring; that is, a stake of one unit returns a net profit of b units with probability p and loses the unit with probability $q := 1 - p$. The *edge* of each wager is given by

$$e \equiv e(p, b) := \mathbb{E}[\text{profit per unit staked}] = pb - q = p(b + 1) - 1. \quad (1)$$

Throughout, the symbol e is reserved for the edge. To avoid confusion, the exponential function is always written as $\exp(\cdot)$.

On each round n , the bettor stakes an amount $0 \leq s_n \leq B_{n-1}$, where B_{n-1} is the bettor's bankroll after round $n - 1$ (or, equivalently, immediately preceding round n). The stake s_n on round n may depend on previous rounds through any measurable function of the outcomes of rounds $1, \dots, n - 1$. For each $n = 1, 2, \dots$, define

$$X_n := \begin{cases} 1, & \text{bettor wins round } n, \\ 0, & \text{otherwise.} \end{cases}$$

The bettor's bankroll trajectory $(B_n)_{n \geq 0}$ is given by

$$B_n = B_{n-1} + ((b + 1)X_n - 1)s_n.$$

The above setup has been considered by Kelly [12], Thorp [20, 21], Ziemba [14], Epstein [7], and many others throughout the literature on probability and gambling theory. Under the formal mathematical treatment, $X_1, X_2, \dots$ are assumed to be independent, identically distributed

(i.i.d.) Bernoulli random variables with success probability p , so that $(B_n)_{n \geq 0}$ is a random walk whose jumps are determined by the outcomes of the i.i.d. sequence $(X_n)_{n \geq 1}$.

It is within the above framework that Kelly considers his classic question of bet sizing.

Question 1. Assuming $e > 0$, how much should the bettor stake on each round?

Suppose the bettor stakes a fixed fraction $f \in [0, 1]$ of the current bankroll on each round. That is, on each round n the bettor chooses $s_n = fB_{n-1}$ for some $0 \leq f \leq 1$. Under this strategy, the bankroll increases by factor of $1 + bf$ on each win, and decreases by a factor $1 - f$ on each loss. On each round, therefore, the bankroll is multiplied by the random factor $(1 + bf)^{X_i}(1 - f)^{1-X_i}$, and the bankroll after n rounds can be expressed as

$$B_n = B_0 \prod_{i=1}^n (1 + bf)^{X_i} (1 - f)^{1-X_i} = B_0 (1 + bf)^{W_n} (1 - f)^{n-W_n}, \quad (2)$$

where $W_n := \sum_{i=1}^n X_i$ is the number of rounds won in the first n . Writing

$$g(f) := p \log(1 + bf) + q \log(1 - f) \quad (3)$$

for the expected per-round log-growth, the expected log-wealth after n rounds is

$$\begin{aligned} \mathbb{E} \log(B_n) &= \log(B_0) + \mathbb{E}(W_n) \log(1 + bf) + \mathbb{E}(n - W_n) \log(1 - f) \\ &= \log(B_0) + n g(f). \end{aligned} \quad (4)$$

The Kelly criterion is the stake that maximizes (4), or equivalently the growth rate $g(f)$.

Proposition 2 (Kelly criterion [12]). *The function g is strictly concave on $[0, 1)$. If $e > 0$, $g(f)$ is uniquely maximized at*

$$f^* = \frac{e}{b} = \frac{p(b+1) - 1}{b}, \quad (5)$$

and if $e \leq 0$ it is maximized at $f^* = 0$. The quantity f^* above is called the Kelly fraction.

Proof. Differentiating (3) with respect to f gives a first derivative

$$g'(f) = \frac{pb}{1 + bf} - \frac{q}{1 - f}$$

and second derivative

$$g''(f) = -\frac{pb^2}{(1 + bf)^2} - \frac{q}{(1 - f)^2} < 0.$$

Therefore, g is strictly concave. Since $g'(0) = pb - q = e$, the maximizer is interior if and only if $e > 0$, in which case setting $g'(f) = 0$ gives $pb(1 - f) = q(1 + bf)$, i.e. $pb - q = bf$, whence $f^* = (pb - q)/b = e/b$. If $e \leq 0$ then $g'(0) \leq 0$ and g is decreasing on $[0, 1)$, so $f^* = 0$. $\square$

Proposition 2 answers Question 1 under its stated assumptions. Breiman [2] proved that full Kelly staking maximizes the asymptotic growth rate of capital and minimizes the expected time to reach distant wealth goals. The same principle underlies the recently proposed program of ergodicity economics, which recasts a range of economic decisions as the maximization of the time-average growth rate of wealth in the vein of Kelly [15]. Although the Kelly criterion plays a core role in profitable systems in blackjack, sports betting, and investing [14, 20, 21], it has been forcefully contested, most famously by Samuelson [17].

1.2 Thorp’s case for fractional Kelly

Definition 3 (Fractional Kelly). A strategy that stakes αf^* on each round, for $\alpha \in (0, 1)$ and f^* the Kelly fraction in (5), is called α -fractional Kelly.

From a purely mathematical point of view, any α -fractional Kelly strategy, with $0 < \alpha < 1$, is strictly dominated (in the sense of maximizing expected log growth rate (3)) by full Kelly staking. Under the same conditions, a full Kelly bettor’s wealth grows exponentially faster and larger in the long run with probability 1. Yet, in practice, professional gamblers overwhelmingly prefer to stake less than full Kelly.

In discussing applications of the Kelly criterion, Thorp makes “the case for fractional Kelly” [21, Section 7.3] by laying out mathematical arguments in favor of the widely used heuristic of betting a fraction $0 < \alpha < 1$ of the prescribed Kelly stake. Quoting from Thorp [21]:

“During a large downward oscillation experience shows that bettors will generally either quit or be eliminated by a minimum bet size requirement.

Some lessons here are:

- (1) To the extent m_e is an uncertain estimate of m_t , it is wise to assume $m_t < m_e$ and to choose $f < f^*$ by enough to prevent $g \leq 0$.
- (2) Subject to (1), choosing f in the range $0.5f^* \leq f < f^*$ offers protection against $g \leq 0$ with a reduction of g that is likely to be no more than 25%.”

(In Thorp’s notation, m_e is the estimated edge, m_t the true edge, and g the growth rate.) Thorp’s explanation highlights two practical issues with full Kelly staking:

1. The variance is high. For any $0 < d < 1$, a full Kelly bettor at even money has probability $1 - d$ of experiencing a drawdown of at least $d \times B_0$: that is, a 50% probability of losing 50% of the initial bankroll, a 20% probability of losing 80% of the initial bankroll, and so on. Such oscillations, per Thorp’s observation above, cause many bettors to quit before their edge is realized.
2. The edge is uncertain. In Kelly’s theory, the edge is assumed to be known, in which case the sizing is mathematically proven to be optimal. In practice, the edge used in calculating the Kelly criterion is a best-guess estimate of the true edge. In the best case, the estimated edge correctly approximates the true edge. But because of adverse selection, the estimated edge often overestimates the true edge, in which case the full Kelly fraction leads to overbetting, which can lead to negative bankroll growth even if the edge is positive.

From these observations, Thorp concludes that it is prudent to assume the true edge is less than the estimated edge, and to choose $f < f^*$ in order to prevent bankroll growth from becoming negative due to overbetting. He goes on to suggest fractions in the range $0.5 \leq \alpha < 1$ while stating,

“My gambling and investment experience, as well as reports from numerous blackjack players and teams, suggests that most people strongly prefer the increased safety and psychological comfort of “half Kelly” (or some nearby value), in exchange for giving up 1/4 of their growth rate.”

Although the rationale underlying fractional Kelly sizing aligns with common sense, the choice of fractional amount still follows heuristic guidelines, and these heuristics often disagree. Whereas Thorp’s analysis supports fractional corrections in the range $0.5 \leq \alpha < 1$, the fractional values used in practice are typically in the range $\alpha \leq 0.5$. The framework developed below reconciles the two by making the choice of α precise based on the bettor’s preference for optimizing bankroll growth subject to controlling the errors of type (I) and (II) above.

1.3 Overview of approach and main results

Thorp’s observation that bettors who experience large or extended losing streaks tend to quit reflects the common sense principle that the bettor’s edge is merely an estimate. This principle is not relevant in the theory underlying the Kelly criterion, which tacitly assumes that the edge is known. But when the edge is estimated rather than known, a losing streak that produces a significant drawdown cannot simply be dismissed as natural variance, as such events can alternatively be explained as the estimated edge being wrong. The practicing bettor, therefore, must weigh two competing explanations in real time: whether the drawdown is a result of natural variation or of an overestimated edge. The bettor’s objective is twofold: (i) to realize the edge if it is real and (ii) to find out cheaply and quickly if it is not.

The framework proposed here formalizes this objective. At the outset, the bettor fixes

- a **stop loss** E : the fraction of bankroll he is willing to lose before abandoning the strategy;
- a **premature stopping probability** π : the probability of hitting the stop loss if the estimated edge is correct; and
- an **objective**, such as expected wealth or the growth rate of wealth, to be maximized subject to the constraint that (E, π) imposes.

The stop loss reflects that some amount of risk must be taken in order to realize the potential edge. The premature stopping probability controls the possibility that the bettor loses and quits on a profitable opportunity. The bettor then stakes an amount that optimizes the stated objective subject to the constraint imposed by (E, π) .

As in Kelly’s original analysis, we assume the bettor stakes a fixed fraction f of bankroll on each round. Write B_0 for the initial bankroll, B_n for the bankroll after round $n \geq 1$, $S_n = \log(B_n/B_0)$ for the log change in wealth after round $n \geq 1$, and let $\theta^* \equiv \theta^*(f)$ be the *Cramér root* of the per-round log-return Y . In particular, $\theta^*(f)$ is the unique $\theta > 0$ solving $\mathbb{E}[\exp(-\theta Y)] = 1$. Then $\exp(-\theta^* S_n)$ is a martingale, and optional stopping shows that, when the estimated edge is correct, the probability of ever hitting the stop loss is $(1 - E)^{\theta^*}$, up to an explicit overshoot factor that is bounded uniformly over stakes (Theorem 7, Proposition 11). A bettor with no edge, by contrast, hits every stop loss with probability 1 (Lemma 5). Hitting the stop loss therefore carries a likelihood ratio of $(1 - E)^{-\theta^*} \geq 1/\pi$ against the estimated edge. The premature stopping probability is the significance level of a sequential test of the hypothesis that the estimated edge is accurate (Section 4.5).

The main results are as follows.

- *The optimal staking formula* (Theorem 18). The constraint $(1 - E)^{\theta^*(f)} \leq \pi$ caps the stake at the unique $\bar{f}$ solving $\theta^*(\bar{f}) = \log \pi / \log(1 - E) =: \theta_{\min}$. Expressed as a fraction $\bar{\alpha} = \bar{f}/f^*$ of the Kelly stake, the cap has the small-edge closed form

$$\bar{\alpha} = \frac{2 \log(1 - E)}{\log(1 - E) + \log \pi}.$$

The expected-wealth maximizer stakes $\bar{f} = \bar{\alpha}f^*$, which can exceed full Kelly under some conditions. Under the objective of maximum growth rate (i.e., Kelly’s objective), the growth-rate maximizer stakes $\min(\bar{f}, f^*)$. The two prescriptions coincide precisely when $\pi \leq 1 - E$ (Corollary 20), a condition met by most practically sensible choices of (E, π) . For example, under the mild requirement that $E \geq 0.05$ (i.e., risk at least 5% of total bankroll on a strategy) and $\pi \leq 0.05$ (i.e., accept at most 5% probability of hitting stop loss when edge is positive), we have $\pi \leq 1 - E$ and the above conditions agree.

- *Exact benchmarks* (Corollary 8, Theorem 26). At full Kelly the Cramér root equals $\theta^* = 1$ for every payoff structure and every edge. Therefore, for every E , full Kelly staking hits the stop loss with probability $1 - E$. At even money and shorter odds the identity extends to a one-parameter family of supermartingales: for every fraction $\alpha \leq 1$ of f^* , $(B_0/B_n)^{2/\alpha-1}$ is a supermartingale, so that the diffusion power law bounds the drawdown probability at every edge. For example, at half Kelly ($\alpha = 1/2$) we arrive at the *cube law*, which bounds the drawdown probability by $(1 - E)^3$ (Theorem 26).
- *Target probabilities and the sequential test* (Theorem 31, Corollary 35, Section 4.5). The probability of multiplying the bankroll by a target factor M before hitting the stop loss follows from the same martingale in closed form, and at the optimal stake it is at least $1 - \pi$ for every target M . Hitting the stop loss is a level- π , power-1 sequential test of the hypothesis that the estimated edge is correct.
- *Robustness* (Theorem 38). If the true edge is a fraction $\kappa < 1$ of the estimate, the better staking αf^* is effectively an (α/κ) -fractional Kelly better. Growth stays positive if and only if $\kappa > \alpha/2$, and scaling the prescribed stake by a chosen factor $\kappa_{\min} \in (0, 1]$ restores the (E, π) guarantee whenever the true edge is at least $\kappa_{\min}$ times the estimate.
- *General returns* (Remark 39). The arguments do not use the binary structure of the wager, the binary wager is the worst case among returns with the same mean and maximum payoff, and for Gaussian log-returns and in continuous time the closed forms hold exactly.

Monte Carlo simulations confirm the closed form results and develop intuition throughout.

An important practical consequence is that the framework determines the fractional Kelly value α from (E, π) , rather than choosing α arbitrarily. The *exponent rule* $\bar{\alpha} = 2/(1 + \theta_{\min})$ of Corollary 21 provides an explicit recipe for a given choice of (E, π) . For many reasonable specifications, the prescribed $\bar{\alpha}$ lands in the range typically adopted in practice. For example, a bettor who accepts a 10% chance of prematurely abandoning a winning strategy at a 50% stop loss is prescribed 46% Kelly, while the stricter $\pi = 0.02$ at the same stop loss prescribes 30% Kelly (Table 3). Guarding additionally against an overestimated edge, with $1/2 < \kappa_{\min} < 3/4$, moves the prescription into the 15–23% range, in agreement with the fifth- to quarter-Kelly values often reported in practice. As our analysis shows, no single fraction is universally optimal: $\bar{\alpha}$ depends on (E, π) , the odds, and the bettor’s skepticism about his own estimate. Section 6 relates the framework to risk aversion, adverse selection, and existing drawdown-control approaches.

2 The bettor’s problem

The framework guards against the two errors identified in the Introduction: (I) betting without an edge and (II) not betting with an edge. In the notation of Section 1.1, error (I) is betting

when $e \leq 0$ and error (II) is declining to bet or quitting when $e > 0$. For a disciplined bettor, variance is the dominant driver of both errors: variance convinces losing bettors that they are winners (error I) and winning bettors that they are losers (error II). Therefore, in the proposed framework, the practical problem with variance is not volatility or risk per se, but how it corrupts decision making.

Without loss of generality, we assume that the bettor's stake will be scaled proportionally to bankroll (as in the Kelly formulation), so we express $f_n := s_n/B_{n-1}$ as the fraction of bankroll risked on round n . The two error conditions suggest that the bettor specify the following parameters at the outset.

- **Stop loss (E):** the maximum amount, as a fraction of initial bankroll, that the bettor is willing to lose before stopping. The bettor commits to abandoning the strategy as soon as bankroll reaches a level at or below $(1 - E)B_0$.
- **Premature stopping probability (π):** the probability of hitting the stop loss when the true edge equals the estimated edge; that is, the probability of incorrectly quitting a winning strategy.
- **Objective:** the bettor's goal subject to the above two conditions. For example, the bettor may seek to maximize absolute wealth, maximize growth rate of wealth, or maximize the probability of doubling bankroll.

The stop loss E mitigates the impact of error (I). The premature stopping probability π bounds the chance of error (II).

Remark 4 (Observation without betting). One might recommend that the bettor should initially stake nothing while recording outcomes over enough rounds to pin down the edge before risking money. This option is rarely available in practice. First, the circumstances of the bet are often unstable. In sports betting and poker the counterparty's tendencies drift, so past observation need not transfer to future rounds, and an adversarial opponent responds to the bettor's own play, so the strategy cannot be assessed from the sideline. Second, even when the game itself is stable, observation establishes only a theoretical edge rather than an executable one. A card counter can tally hypothetical profits from recorded play, but executing the count under live conditions, without error and without detection, is necessary to realize the edge. The relevance of executability is even more pronounced in more advanced advantage situations [9, 11].

The above parameters E and π are based on the bettor's preferences. The stop loss represents the bettor's willingness to risk capital on a strategy that may not work out. For example, a bettor with a bankroll of \$1,000,000 might allocate \$100,000 to testing and executing a new strategy. In doing so, the bettor caps the potential downside in case the edge is not as high as estimated. Within the framework, setting a stop loss of \$100,000 on a bankroll of \$1,000,000 corresponds to $E = 0.10$, as \$100,000 is 10% of \$1,000,000.

Objectives likewise differ across bettors. We focus here on maximizing the expected growth rate of B_n subject to the constraints of E and π , which allows direct comparison with Kelly's treatment and the fractional-Kelly literature. Section 4 treats maximizing expected wealth alongside the growth rate, and Section 4.4 works out a further alternative, the probability of multiplying the bankroll by a target factor before hitting the stop loss, which generalizes questions addressed by Dubins and Savage [6].




Figure 1: Six simulated bankroll trajectories under full Kelly betting with a 10% edge at even money, starting from \$100,000 (dotted line; the dashed line marks \$20,000), on a logarithmic scale. *Left*: the first 2,000 rounds, exhibiting deep and prolonged drawdowns. *Right*: the first 5,000 rounds, by which point the positive drift dominates.

2.1 Variance and information

The pair (E, π) constrains the bet size, and the constraint has an informational interpretation. Smaller stake induces less variance in the bankroll trajectory. With less variance, natural fluctuation is less likely to carry the bankroll below the stop loss, and so hitting the stop loss becomes stronger evidence that the estimated edge is wrong. The pair (E, π) thus sets the balance between bankroll growth and information gain based on the bettor's preferences.

Figure 1 shows simulated bankroll trajectories under full Kelly betting with a 10% edge at even money. Over the first 2,000 rounds the paths are widely dispersed. Several experience drawdowns of 50% or more, and one drops by more than 90%. Over a longer horizon the positive drift pulls every trajectory into positive territory. In Kelly's idealized world, where the edge is known with certainty, this volatility is harmless, and the bettor simply keeps betting and waits for the law of large numbers. In practice, few bettors could withstand a loss of half the bankroll within the first 1,000 rounds, which happens about half the time under full Kelly, without quitting or reassessing. The framework proposed here accounts for this.

3 Premature stopping probability and evidence in a drawdown

3.1 The log-wealth random walk

Fix a stake $f \in (0, 1)$ and let

$$S_n := \log \frac{B_n}{B_0} = \sum_{i=1}^n Y_i, \quad Y_i := \begin{cases} u := \log(1 + bf), & X_i = 1, \\ d := \log(1 - f), & X_i = 0, \end{cases} \quad (6)$$

so that $(S_n)_{n \geq 0}$ is a random walk with i.i.d. increments and mean step $\mathbb{E}[Y_i] = g(f)$ as in (3). Write $\mathbb{P}_e$ for the law when the edge is e ($p = (e + 1)/(b + 1)$), so that $\mathbb{P}_0$ is the law when the edge is 0 ($p = 1/(b + 1)$), but the bettor stakes the same amount f in either case. For $E \in (0, 1)$, let

$$\tau_E := \inf\{n \geq 1 : B_n \leq (1 - E)B_0\} = \inf\{n \geq 1 : S_n \leq \log(1 - E)\} \quad (7)$$

be the first time the bankroll has drawn down by the fraction E ; that is, the first time the stop loss is hit. The probability of hitting the stop loss when the true edge is e is precisely $\mathbb{P}_e(\tau_E < \infty)$. The drawdown probabilities below are first-passage probabilities for the random walk (6).

Lemma 5 (Under no edge, the stop loss is reached with probability 1). *If $e \leq 0$ and $f > 0$, then $g(f) < 0$. Consequently, $S_n \rightarrow -\infty$ almost surely (a.s.) and $\mathbb{P}_e(\tau_E < \infty) = 1$ for every $E \in (0, 1)$; in particular, $\mathbb{P}_0(\tau_E < \infty) = 1$.*

Proof. Let $R = bf$ on a win and $R = -f$ on a loss, so $1 + R$ is the per-round bankroll multiplier and $\mathbb{E}[R] = pbf - qf = ef \leq 0$. By strict concavity of $\log(\cdot)$ and Jensen's inequality applied to the nondegenerate R , $g(f) = \mathbb{E} \log(1 + R) < \log(1 + \mathbb{E}[R]) = \log(1 + ef) \leq 0$. By the strong law of large numbers, $S_n/n \rightarrow g(f) < 0$ a.s., so that $S_n \rightarrow -\infty$ a.s. and the random walk crosses every level $\log(1 - E) < 0$ with probability 1. $\square$

Lemma 5 says that a bettor with no edge is certain to hit any stop loss if he plays the game long enough. The key question is how this compares to the same event when the edge is positive. If the two probabilities are close, hitting the stop loss tells the bettor almost nothing about his estimated edge. If the probabilities are far apart, hitting the stop loss provides meaningful information to the bettor.

3.2 The evidence in a drawdown

Definition 6 (Likelihood ratio of a drawdown). The *likelihood ratio of an E -drawdown under edge e* is the Bayes factor for the null of no edge against the hypothesis of an edge e ,

$$\text{LR}(E) := \frac{\mathbb{P}_0(\tau_E < \infty)}{\mathbb{P}_e(\tau_E < \infty)}. \quad (8)$$

A value $\text{LR}(E) = k$ means the drawdown is k times more likely when the bettor has no edge than when he has the edge he assumed. Large k , therefore, implies that hitting the stop loss provides strong evidence in support of quitting, whereas k near 1 provides weak or no evidence at all.

By Lemma 5 the numerator in (8) equals 1, so

$$\text{LR}(E) = \frac{1}{\mathbb{P}_e(\tau_E < \infty)} \geq \frac{1}{\pi} \quad (9)$$

whenever the stake is calibrated so that the premature stopping probability is no greater than π . Therefore, a bettor who calibrates bet sizing so that $\pi = 0.02$ is, equivalently, demanding that hitting the stop loss constitute 50-to-1 evidence against his edge before quitting.

It remains to compute $\mathbb{P}_e(\tau_E < \infty)$. Let θ^* be the Cramér root of the log-return, that is, the unique strictly positive solution of

$$\psi(\theta) := \mathbb{E}_e[\exp(-\theta Y)] = p \exp(-\theta u) + q \exp(-\theta d) = 1, \quad (10)$$

for u and d as defined in (6). In the language of classical risk theory, θ^* is the Lundberg adjustment coefficient of the log-wealth walk (6), and the closed form below is Lundberg's inequality together with its Cramér–Lundberg overshoot correction [1, 5].

Theorem 7 (Premature stopping probability and likelihood ratio). *Suppose $e > 0$ and $f > 0$ is such that $g(f) > 0$. Then (10) has a unique root $\theta^* \equiv \theta^*(f) > 0$ and*

$$\mathbb{P}_e(\tau_E < \infty) = C_E^{-1} (1 - E)^{\theta^*} \leq (1 - E)^{\theta^*}, \quad (11)$$

where C_E satisfies $1 \leq C_E \leq (1 - f)^{-\theta^*}$. Consequently

$$\text{LR}(E) = C_E (1 - E)^{-\theta^*} \geq (1 - E)^{-\theta^*}. \quad (12)$$

Proof. The function ψ is strictly convex with $\psi(0) = 1$ and $\psi'(0) = -\mathbb{E}_e[Y] = -g(f) < 0$, while $\psi(\theta) \rightarrow \infty$ as $\theta \rightarrow \infty$ (since $d < 0$ makes $\exp(-\theta d) \rightarrow \infty$). Therefore, ψ has a unique root $\theta^* > 0$.

At this root θ^* , $M_n := \exp(-\theta^* S_n)$ is a nonnegative martingale:

$$\begin{aligned}\mathbb{E}_e[M_n \mid \mathcal{F}_{n-1}] &= \mathbb{E}_e[\exp(-\theta^* S_n) \mid \mathcal{F}_{n-1}] \\ &= \exp(-\theta^* S_{n-1}) \mathbb{E}_e[\exp(-\theta^* Y_n) \mid \mathcal{F}_{n-1}] \\ &= M_{n-1} \psi(\theta^*) \\ &= M_{n-1}.\end{aligned}$$

Note that τ_E , and therefore also $\tau_E \wedge n$ for any $n = 1, 2, \dots$, is a stopping time for (S_n) . Optional stopping at $\tau_E \wedge n$ gives $\mathbb{E}_e[M_{\tau_E \wedge n}] = \mathbb{E}_e[M_0] = 1$.

On the event $\{\tau_E = \infty\}$, we have $S_n \rightarrow +\infty$ with probability 1 (as $g(f) > 0$), so that $M_n \rightarrow 0$ with probability 1. On the event $\{\tau_E < \infty\}$, we have $M_{\tau_E \wedge n} \rightarrow M_{\tau_E}$. The stopped martingale is uniformly bounded: for $k < \tau_E$ we have $S_k > \log(1 - E)$ and therefore $M_k < (1 - E)^{-\theta^*}$, while at τ_E the single overshooting step satisfies $S_{\tau_E} > \log(1 - E) + d$, giving $M_{\tau_E} < (1 - E)^{-\theta^*} \exp(-\theta^* d)$. Hence by bounded convergence,

$$1 = \mathbb{E}_e[M_{\tau_E} \mathbf{1}\{\tau_E < \infty\}] = \mathbb{E}_e[\exp(-\theta^* S_{\tau_E}) \mathbf{1}\{\tau_E < \infty\}]. \quad (13)$$

At τ_E the walk is at or below the barrier, $S_{\tau_E} \leq \log(1 - E)$, so $\exp(-\theta^* S_{\tau_E}) \geq (1 - E)^{-\theta^*}$. Writing $C_E := \mathbb{E}_e[\exp(-\theta^*(S_{\tau_E} - \log(1 - E))) \mid \tau_E < \infty] \geq 1$ for the expected exponential overshoot, $1 = (1 - E)^{-\theta^*} C_E \mathbb{P}_e(\tau_E < \infty)$, which is (11); (12) follows from $\text{LR}(E) = 1/\mathbb{P}_e(\tau_E < \infty)$. Since the overshooting step gives $S_{\tau_E} > \log(1 - E) + d$, the overshoot exponent is less than $-\theta^* d$, so $C_E \leq \exp(-\theta^* d) = (1 - f)^{-\theta^*}$. $\square$

Corollary 8 (Diffusion limit and the full-Kelly identity). *Let $f = \alpha f^*$ with $\alpha \in (0, 1]$ fixed, and define $\theta^*(f)$ as in Theorem 7.*

(i) *As $e \rightarrow 0$, $\theta^*(\alpha f^*) = 2/\alpha - 1 + O(e)$ for every payoff b , and $C_E = 1 + O(e)$, so the bounds of Theorem 7 are tight in the small-edge limit.*

(ii) *At full Kelly the diffusion value is exact. For every payoff b and every edge $e > 0$, $\theta^*(f^*) = 1$. Equivalently, the reciprocal bankroll $1/B_n$ is a martingale under full Kelly staking.*

Proof. Expand (10) to second order, so that $1 = 1 - \theta g(f) + \frac{1}{2}\theta^2 \mathbb{E}_e[Y^2] + o(\theta^2)$ and hence $\theta^* \approx 2g(f)/\mathbb{E}_e[Y^2] \approx 2g(f)/\text{Var}_e(Y)$. Take $f = \alpha f^* = \alpha e/b$ and expand in the edge e . To first order, $g(f) \approx ef - \frac{1}{2}f^2(pb^2 + q)$ and $\text{Var}_e(Y) \approx f^2(pb^2 + q)$. Because $p \rightarrow 1/(b + 1)$ as $e \rightarrow 0$, and $pb^2 + q \rightarrow b$, it follows that $g(f) \approx \frac{e^2}{b}(\alpha - \frac{\alpha^2}{2})$ and $\text{Var}_e(Y) \approx \frac{\alpha^2 e^2}{b}$, so $\theta^* \approx \frac{2(\alpha - \alpha^2/2)}{\alpha^2} = \frac{2}{\alpha} - 1$ regardless of b . The same expansion gives $-\theta^* \log(1 - f) = (2 - \alpha)e/b + O(e^2)$, so $C_E \leq (1 - f)^{-\theta^*} = 1 + O(e)$ and the overshoot correction vanishes as the edge becomes small. It does not vanish for small stakes at a fixed edge, as Remark 14 quantifies.

For (ii), at $\alpha = 1$ we have $f = f^* = e/b$, so $1 + bf^* = 1 + e$, and substituting $\theta = 1$ into (10) with $p = (e + 1)/(b + 1)$ and $q = (b - e)/(b + 1)$ gives

$$\psi(1) = \frac{p}{1 + e} + \frac{q}{1 - e/b} = \frac{1}{b + 1} + \frac{b - e}{b + 1} \cdot \frac{b}{b - e} = \frac{1}{b + 1} + \frac{b}{b + 1} = 1,$$

so $\theta^* = 1$ exactly, for every b and every $e > 0$. Since $\psi(1) = \mathbb{E}_e[(1 + R)^{-1}] = 1$ at full Kelly, the reciprocal bankroll $1/B_n$ is a martingale under full Kelly staking, whatever the odds. $\square$

Remark 9 (On θ^* notation). The Cramér root depends on the staking factor f through (10). For convenience, we often suppress this dependence and write $\theta^* \equiv \theta^*(f)$ when f is understood from context and no confusion will result. As fractional Kelly is an important special case, we may also write

$$\theta_\alpha^* := \theta^*(\alpha f^*), \quad (14)$$

where f^* is the optimal fraction given by the Kelly criterion, and the edge e and payoff b are fixed by context.

Let $f_c := \sup\{f > 0 : g(f) > 0\}$ denote the *zero-growth stake*, so that $g(f) > 0$ only for $f \in (0, f_c)$ and $g(f_c) = 0$. Note that $f^* < f_c < 1$ when $e > 0$; when $e \leq 0$, $g(f) < 0$ for every $f \in (0, 1)$ and no such stake exists.

Lemma 10 (The Cramér root decreases in the stake). *Suppose $e > 0$. The map $f \mapsto \theta^*(f)$ is continuous and strictly decreasing on $(0, f_c)$, with $\theta^*(f) \rightarrow \infty$ as $f \downarrow 0$ and $\theta^*(f) \rightarrow 0$ as $f \uparrow f_c$.*

Proof. Write

$$\psi_f(\theta) := p(1 + bf)^{-\theta} + q(1 - f)^{-\theta}$$

for the function of (10), with its dependence on the stake f now explicit. Recall from the proof of Theorem 7 that for each fixed $f \in (0, f_c)$ the map $\theta \mapsto \psi_f(\theta)$ is strictly convex, equals 1 at $\theta = 0$ with strictly negative slope $-g(f)$ there, and has unique positive root $\theta^*(f)$; consequently

$$\psi_f(\theta) < 1 \text{ for } 0 < \theta < \theta^*(f), \quad \psi_f(\theta) > 1 \text{ for } \theta > \theta^*(f). \quad (15)$$

The argument uses one further fact: the shape of $\psi_f(\theta)$ as a function of the stake f for fixed $\theta > 0$. We have $\psi_0(\theta) = p + q = 1$,

$$\begin{aligned} \frac{\partial}{\partial f} \psi_f(\theta) \Big|_{f=0} &= \theta(q - pb) = -\theta e < 0 \\ \frac{\partial^2}{\partial f^2} \psi_f(\theta) &= \theta(\theta + 1)[pb^2(1 + bf)^{-\theta-2} + q(1 - f)^{-\theta-2}] > 0, \end{aligned}$$

so $f \mapsto \psi_f(\theta)$ is strictly convex on $[0, 1)$, starts at 1, and dips below 1 immediately to the right of $f = 0$.

Strict decrease. Fix $0 < f_1 < f_2 < f_c$ and put $\theta_2 := \theta^*(f_2)$, so that $\psi_{f_2}(\theta_2) = 1$. As a function of the stake, $\psi_f(\theta_2)$ is strictly convex and takes the value 1 at the two stakes $f = 0$ and $f = f_2$; a strictly convex function lies strictly below its chord, so $\psi_{f_1}(\theta_2) < 1$. By (15) applied at the stake f_1 , this forces $\theta_2 < \theta^*(f_1)$, that is, $\theta^*(f_2) < \theta^*(f_1)$.

Continuity. A strictly convex function is strictly increasing where it crosses a level from below, so $\frac{\partial}{\partial \theta} \psi_f(\theta) > 0$ at $\theta = \theta^*(f)$. The implicit function theorem applied to $(f, \theta) \mapsto \psi_f(\theta) - 1$ then gives continuity of $f \mapsto \theta^*(f)$ on $(0, f_c)$.

Limits. Fix any $\theta > 0$. Since the slope of $f \mapsto \psi_f(\theta)$ at $f = 0$ is $-\theta e < 0$, we have $\psi_f(\theta) < 1$ for all sufficiently small $f > 0$, and (15) then gives $\theta^*(f) > \theta$. As $\theta > 0$ was arbitrary, $\theta^*(f) \rightarrow \infty$ as $f \downarrow 0$. At the zero-growth stake the drift vanishes, $g(f_c) = 0$, so $\theta \mapsto \psi_{f_c}(\theta)$ has zero slope at $\theta = 0$ and, being strictly convex with $\psi_{f_c}(0) = 1$, satisfies $\psi_{f_c}(\theta) > 1$ for every $\theta > 0$. By continuity of $\psi_f(\theta)$ in f , for each fixed $\theta > 0$ we have $\psi_f(\theta) > 1$ for all f close enough to f_c , and (15) then gives $\theta^*(f) < \theta$. Hence $\theta^*(f) \rightarrow 0$ as $f \uparrow f_c$. $\square$

By Theorem 7, the closed form $(1 - E)^{\theta^*}$ is governed entirely by the exponent θ^* , which decreases as the stake grows (Lemma 10). Therefore, the more aggressively one bets, the more probable any given drawdown, and therefore the less evidence against the perceived edge such a drawdown provides.

Proposition 11 (Overshoot bounds). *Under the hypotheses of Theorem 7, $1 \leq C_E \leq (1 - f)^{-\theta^*}$, and consequently*

$$\begin{aligned} (1 - f)^{\theta^*} (1 - E)^{\theta^*} &\leq \mathbb{P}_e(\tau_E < \infty) \leq (1 - E)^{\theta^*} \quad \text{and} \\ (1 - E)^{-\theta^*} &\leq \text{LR}(E) \leq (1 - f)^{-\theta^*} (1 - E)^{-\theta^*}. \end{aligned} \quad (16)$$

Proof. At time $\tau_E - 1$ the walk is still above the barrier, $S_{\tau_E - 1} > \log(1 - E)$, and the crossing step is at least $d = \log(1 - f)$, so the overshoot $O := \log(1 - E) - S_{\tau_E}$ satisfies $0 \leq O < -d$. Hence $C_E = \mathbb{E}_e[\exp(\theta^* O) \mid \tau_E < \infty] \in [1, (1 - f)^{-\theta^*}]$. Substituting into (11) and (12) gives (16). $\square$

In the diffusion regime $\theta_\alpha^* |d| \approx (2/\alpha - 1) \alpha f^* = (2 - \alpha) f^* \leq 2f^*$, so the two sides of (16) differ by at most a factor $\exp(2f^*)$ regardless of the Kelly fraction. For the running 10% even-money example the factor $(1 - f)^{\theta^*}$ has a uniform lower bound $\exp(-2f^*) = 0.82$ and ranges only from 0.90 at full Kelly to 0.83 at tenth Kelly.

Remark 12 (Drawdown from a new level). By the strong Markov property and scale invariance of fractional staking, the framework renews at any stopping time: at the moment the bankroll first reaches a level L , the probability of ever falling back to cL is c^{θ^*} for any $c \in (0, 1)$.

Remark 13 (Computing the Cramér root). Equation (10) admits closed form solutions only in special cases, but its numerical solution can be computed easily. Because ψ is strictly convex with $\psi(0) = 1$ and $\psi'(0) = -g(f) < 0$, the function dips below 1 immediately to the right of 0 and recrosses it exactly once; therefore, θ^* is the unique sign change of $\psi(\theta) - 1$ on $(0, \infty)$. It can be found numerically by scanning a coarse grid of θ values to bound $\theta^* \in (\theta_{\text{lo}}, \theta_{\text{hi}})$, with $\psi(\theta_{\text{lo}}) < 1 < \psi(\theta_{\text{hi}})$. For the worked example of Section 4 (i.e., 10% edge at even money, $\alpha = 0.30$, $f = 0.03$), the integer grid gives $\psi(5) = 0.9985 < 1 < 1.0009 = \psi(6)$, and further bisection steps yield $\theta^* \approx 5.685$.

3.3 Accuracy of the diffusion approximation

The diffusion approximation $2/\alpha - 1$ is exact in the small-edge limit, for every fixed α . Table 1 compares the exact root to the diffusion approximation for a range of fractional Kelly values for the two scenarios: (1) a 10% edge at even money ($b = 1$, $e = 0.10$, $f^* = 10\%$) and (2) a 50% edge at 10-to-1 ($b = 10$, $e = 0.50$, $f^* = 5\%$).

The even-money column is well-approximated by $2/\alpha - 1$, showing a deviation less than 1% in all cases. The diffusion approximation is less accurate for the second case, a 50% edge at 10-to-1 odds. In the typical range $\alpha \leq 0.50$, the approximation $2/\alpha - 1$ overestimates θ^* by at least 5% (at $\alpha = 0.50$) and as much as 10.5% (at $\alpha = 0.10$). Smaller θ^* implies that every drawdown is more probable and less informative than the diffusion limit estimates, so the odds-free formula will lead to overbetting relative to the desired parameters.

Table 1: The exact Cramér root θ_α^* of (10) against the diffusion approximation $2/\alpha - 1$, for the two running examples. Both structures agree exactly with the diffusion value at full Kelly ($\theta^* = 1$); the even-money structure tracks it to a fraction of a percent throughout, while the 10-to-1 structure runs systematically below it.

Kelly fraction α	$2/\alpha - 1$	exact root θ_α^*	
		10% edge, $b = 1$	50% edge, $b = 10$
10%	19.00	19.06	17.03
25%	7.00	7.02	6.35
$33\frac{1}{3}\%$	5.00	5.02	4.57
50%	3.00	3.01	2.79
$66\frac{2}{3}\%$	2.00	2.00	1.90
75%	1.67	1.67	1.60
100%	1.00	1.00	1.00
125%	0.60	0.60	0.64
150%	0.33	0.33	0.40

Table 2: Premature stopping probability at a 50% stop loss and its likelihood ratio, under fractional Kelly betting with a 10% edge at even money. Simulated values are Monte Carlo estimates over the first $N = 1,000$ rounds (4×10^5 paths), with the simulated LR computed as the ratio of the simulated hitting frequencies under the null and under the edge over the same horizon; closed form values are the infinite-horizon expressions $\mathbb{P}_e = (1 - E)^{\theta^*}$ and $\text{LR} = (1 - E)^{-\theta^*}$ of Theorem 7, with $E = 0.5$.

Kelly fraction α	θ^*	simulated ($N = 1,000$)		closed form ($N = \infty$)	
		$\mathbb{P}_e(\tau_E < \infty)$	LR	$\mathbb{P}_e(\tau_E < \infty)$	LR
100%	1.00	47.0%	2.1	50%	2.0
50%	3.01	11.5%	7.4	12.4%	8.0
25%	7.02	0.70%	73	0.77%	130
10%	19.1	$< 0.1\%$	$> 10^3$	$1.8 \times 10^{-4}\%$	5.5×10^5

3.4 Finite horizons and numerical confirmation

The closed form above applies to the infinite time horizon, that is, the probability of a drawdown of at least E at any time. In practice the bettor may only care about a finite number of rounds. Table 2 compares the infinite-horizon formula of Theorem 7 with Monte Carlo estimates over the first $N = 1,000$ rounds for the even-money 10%-edge example at a 50% stop loss. The simulated probabilities track the closed form closely, with a small systematic deficit quantified in Remark 14 below. The closed form in turn recovers the numbers long quoted in the Kelly literature [21, Section 7.3]. In the diffusion limit the probability of ever halving the bankroll is $(1/2)^{2/\alpha-1}$, i.e., $1/2$ at full Kelly, $1/8$ at half Kelly, $1/128 \approx 0.8\%$ at quarter Kelly, with the simulated values for $N = 1,000$ shown in the table (47.0%, 11.5%, 0.70%) confirming that most of the probability of the adverse event is realized within the first 1,000 rounds.

Remark 14 (The size of the overshoot deficit). The simulated probabilities in Table 2 sit a few percent below the closed form, and the overshoot factor C_E of Theorem 7 accounts for

the gap. The overshoot at the crossing is at most one losing step. When the two step sizes are incommensurable and the barrier lies many losing steps deep, renewal theory implies the position of the crossing within that step is approximately uniformly distributed between 0 and the maximum overshoot size [1]. Uniformity gives the approximation

$$C_E \approx \frac{\exp(\delta) - 1}{\delta}, \quad \delta := -\theta^* \log(1 - f) = -\theta^* d. \quad (17)$$

At half Kelly, where $\delta = 0.154$, it predicts a simulated-to-closed form ratio of $1/1.081 = 0.925$ against the observed $11.5/12.4 = 0.927$. At quarter Kelly it predicts 0.914 against the observed 0.909. Note that (17) is a heuristic rather than a theorem. Rigorous statements are given by the bounds of Proposition 11.

Two consequences are worth noting. First, the deficit does not vanish for small stakes at a fixed edge. As $f \downarrow 0$, the root $\theta^*(f)$ grows like $1/f$ while the steps shrink like f , and the product $y = \theta^* f$ converges. In this limit the Cramér equation (10) becomes $p \exp(-by) + q \exp(y) = 1$, which is the Lundberg equation of a flat-stake walk with the exponent measured in units of one losing step. The root expands as $y = 2e/(pb^2 + q) + O(e^2)$, and therefore $\delta \rightarrow 2e/(pb^2 + q) + O(e^2)$, which is a nonzero constant by assumption $e > 0$. By (17) the overshoot factor then tends to $C_E \approx 1 + \delta/2$, so that the closed form carries a stable percentage bias of about $\delta/2 \approx e/(pb^2 + q)$, which is 10% for a 10% edge at even money. Second, the deficit does vanish as the edge becomes small. For $f = \alpha f^*$ the losing step is $-\log(1 - f) = \alpha e/b + O(e^2)$ while $\theta^* = 2/\alpha - 1 + O(e)$ by Corollary 8, so $\delta = (2 - \alpha)e/b + O(e^2)$, and this is the sense in which the diffusion limit is exact.

Proposition 15 (Residual risk decays geometrically). *Under the hypotheses of Theorem 7, for every $N \geq 0$ and every $\theta \in (0, \theta^*)$,*

$$\mathbb{P}_e(N < \tau_E < \infty) \leq (1 - E)^{-\theta} \psi(\theta)^N, \quad (18)$$

and $\psi(\theta) < 1$ for $\theta \in (0, \theta^*)$.

Proof. Since ψ is strictly convex with $\psi(0) = \psi(\theta^*) = 1$, we have $\psi(\theta) < 1$ on $(0, \theta^*)$, so $\exp(-\theta S_n)$ is a nonnegative supermartingale. On the event $\{\tau_E > N\}$, hitting the stop loss after time N requires the post- N walk, again a random walk with the same increments, to fall by $S_N - \log(1 - E) > 0$, which by Ville's inequality has conditional probability at most $\exp(-\theta(S_N - \log(1 - E)))$. Hence

$$\mathbb{P}_e(N < \tau_E < \infty) \leq (1 - E)^{-\theta} \mathbb{E}_e[\exp(-\theta S_N)] = (1 - E)^{-\theta} \psi(\theta)^N.$$

□

The best decay rate in (18) comes from minimizing ψ . Expanding $\psi(\theta) = 1 - \theta g(f) + \frac{1}{2} \theta^2 \mathbb{E}_e[Y^2] + O(\theta^3 \mathbb{E}_e|Y|^3)$ places the minimum near $\theta = g(f)/\text{Var}_e(Y)$, which is $\frac{1}{2} \theta^*$ to first order, where $\psi \approx 1 - g(f)^2/(2 \text{Var}_e(Y))$. The residual risk therefore decays at a rate of about $g^2/(2 \text{Var})$ per round. For the worked example of Section 4, with a 10% edge at even money, $\alpha = 0.30$, and $E = 50\%$, this rate is about 0.0037 per round, so the residual bound halves roughly every 190 rounds, and optimizing (18) over θ certifies that at most 0.002, a tenth of the tolerance $\pi = 0.02$, remains outstanding after $N \approx 2,300$ rounds. The Monte Carlo estimates of Table 2, which put most of the infinite-horizon probability inside the first 1,000 rounds, validate this estimate.

Figure 2 displays the trade-off between growth and the evidence carried by a drawdown. As the Kelly fraction α increases, the growth rate $g(\alpha f^*)$ rises to its maximum at full Kelly, while the likelihood ratio of a 50% drawdown falls monotonically. The bettor's problem is to choose a point on this trade-off, which is the subject of the next section.




Figure 2: The growth–evidence trade-off for a 10% edge at even money. The expected log-growth $g(\alpha f^*)$ (dashed) is maximized at full Kelly ($\alpha = 1$, dotted vertical line). The evidence provided by a drawdown decreases monotonically with bet size (solid blue line).

4 The optimal staking formula

We now address the bettor’s problem posed in Section 2. Given a stop loss E , a premature stopping probability π , and an objective, find the bettor’s optimal stake.

We consider three objectives in particular:

- Maximize expected wealth: maximize $\mathbb{E}B_n$ for every $n \geq 1$.
- Maximize expected growth rate: maximize $g(f)$, defined in (3).
- Maximize the probability of reaching a level before E : maximize $\mathbb{P}(B_n \geq MB_0 \text{ before } B_n \leq (1 - E)B_0)$.

Definition 16 (Admissible stake). Fix a stop loss $E \in (0, 1)$ and a premature stopping probability $\pi \in (0, 1)$. A fixed-fraction stake f is *admissible* for (E, π) if

$$(1 - E)^{\theta^*(f)} \leq \pi. \quad (19)$$

Remark 17. By Theorem 7, the closed form criterion (19) is conservative for the exact premature stopping probability (since $\mathbb{P}_e(\tau_E < \infty) \leq (1 - E)^{\theta^*}$) and tight in the small-edge limit.

Taking logarithms and rearranging (19), we see that admissibility requires

$$\theta^*(f) \geq \theta_{\min} := \frac{\log \pi}{\log(1 - E)}. \quad (20)$$

By Lemma 10, θ^* is strictly decreasing in f , so the admissible stakes form an interval $(0, \bar{f}]$, where $\bar{f}$ is the unique stake with $\theta^*(\bar{f}) = \theta_{\min}$. Geometrically, a choice of (π, E) draws a horizontal line at height $\theta_{\min}$ across the decreasing curve $f \mapsto \theta^*(f)$. The admissible stakes are those to the left of the crossing, and the optimal stake is the crossing itself. Figure 3 displays this geometry for the two running examples, with the stake expressed as a Kelly fraction α . The first two objectives are increasing in the stake up to their respective optima. The third is treated in Section 4.4.




Figure 3: A plot of $\theta^*(\alpha f^*) \equiv \theta_\alpha^*$ for two running examples. A specification (π, E) determines a horizontal line at height $\theta_{\min}$. Stakes to the left of the crossing are admissible and the crossing itself occurs at $\bar{\alpha}$. Shown are the worked example $\pi = 0.02$, $E = 50\%$, with $\theta_{\min} = 5.64$ and $\bar{\alpha} \approx 0.30$, and the permissive specification $\pi = 0.10$, $E = 75\%$, with $\theta_{\min} = 1.66$ and $\bar{\alpha} \approx 0.75$, both for the even-money structure.

Theorem 18 (Optimal staking formula). *Fix $E \in (0, 1)$ and $\pi \in (0, 1)$, suppose $e > 0$, and let $\bar{f} \in (0, f_c)$ be the unique stake with $\theta^*(\bar{f}) = \theta_{\min}$, expressed as a fraction of Kelly by $\bar{\alpha} := \bar{f}/f^*$. The optimum among admissible fixed-fraction stakes is as follows.*

- (i) *When the objective is to maximize expected wealth, $\mathbb{E}[B_n] = B_0(1 + ef)^n$ is strictly increasing in f for every n , and the optimal stake is*

$$f_{\text{opt}} = \bar{f} = \bar{\alpha} f^*.$$

- (ii) *When the objective is to maximize the growth rate, $g(f)$ is strictly increasing on $(0, f^*]$ and strictly decreasing on the region $f > f^*$, and the optimal stake is*

$$f_{\text{opt}} = \min(\bar{f}, f^*) = \min(\bar{\alpha}, 1) f^*.$$

As $e \rightarrow 0$ with (E, π) fixed, $\theta^*(\alpha f^*) = 2/\alpha - 1 + O(e)$ by Corollary 8, and the optimal fraction satisfies

$$\bar{\alpha} = \frac{2}{1 + \theta_{\min}} + O(e) = \frac{2 \log(1 - E)}{\log(1 - E) + \log \pi} + O(e). \quad (21)$$

Remark 19. Note that $\bar{\alpha}$ may exceed 1 in case (i) above. When the constraints (E, π) are loose (i.e., π is large relative to E), the expected-wealth maximizer bets more than full Kelly.

Proof. Existence and uniqueness of $\bar{f}$ follow from Lemma 10, since θ^* is continuous and strictly decreasing on $(0, f_c)$ with range $(0, \infty)$ and $\theta_{\min} \in (0, \infty)$. Monotonicity also gives that (20) holds exactly for $f \in (0, \bar{f}]$. For (i), the per-round wealth multiplier has mean $p(1 + bf) + q(1 - f) = 1 + ef$, and the rounds are independent, so $\mathbb{E}[B_n] = B_0(1 + ef)^n$, which is strictly increasing in f

since $e > 0$. The maximum over $(0, \bar{f}]$ is therefore attained at $\bar{f}$. For (ii), g is strictly concave with interior maximizer f^* (Proposition 2), hence strictly increasing on $(0, f^*]$ and strictly decreasing on $[f^*, f_c)$. The maximum over $(0, \bar{f}]$ is attained at $\min(\bar{f}, f^*)$. Finally, substituting $\theta^*(\bar{\alpha}f^*) = 2/\bar{\alpha} - 1 + O(e)$ from Corollary 8 into $\theta^*(\bar{f}) = \theta_{\min}$ and solving for $\bar{\alpha}$ gives $\bar{\alpha} = 2/(1 + \theta_{\min} + O(e)) = 2/(1 + \theta_{\min}) + O(e)$, which is (21). $\square$

Note that the constraint (E, π) determines the stake in both cases above, with the objective only determining whether the optimal stake is capped at full Kelly. Given the bettor's objective, the optimal fractional Kelly value is the most aggressive stake consistent with the bettor's stated willingness to lose (E) and his stated tolerance for quitting a winning strategy (π) . The next corollary makes the agreement between the two objectives precise, and shows that the condition for it is exact.

Corollary 20 (When the two objectives coincide). *The optimal stakes in Theorem 18(i) and (ii) are identical if and only if*

$$\bar{f} \leq f^* \iff \theta_{\min} \geq 1 \iff \pi \leq 1 - E.$$

The equivalence holds for every payoff b and every edge $e > 0$.

Proof. The two prescriptions are $\bar{f}$ and $\min(\bar{f}, f^*)$, which agree if and only if $\bar{f} \leq f^*$. By Lemma 10, θ^* is strictly decreasing in the stake, so $\bar{f} \leq f^*$ if and only if $\theta_{\min} = \theta^*(\bar{f}) \geq \theta^*(f^*) = 1$, the last equality being the exact full-Kelly identity of Corollary 8, valid at every b and $e > 0$. Finally, by (20), $\theta_{\min} \geq 1$ if and only if $-\log \pi \geq -\log(1 - E)$, i.e. $\pi \leq 1 - E$. $\square$

Corollary 21 (The exponent rule). *Fix the edge $e > 0$ and the payoff b , and write $\theta_{\min}(E, \pi) = \log \pi / \log(1 - E)$. The admissible set $(0, \bar{f}]$ and the optimal stakes of Theorem 18(i) and (ii) depend on (E, π) only through $\theta_{\min}(E, \pi)$. Consequently, two pairs (E_1, π_1) and (E_2, π_2) with $\theta_{\min}(E_1, \pi_1) = \theta_{\min}(E_2, \pi_2)$ prescribe the same stake. In particular, fix $j > 0$. Every pair on the curve $\pi = (1 - E)^j$, $0 < E < 1$, has $\theta_{\min} = j$ and prescribes the same stake $\bar{f}$ with $\theta^*(\bar{f}) = j$, which corresponds to the Kelly fraction $\bar{\alpha} = \bar{f}/f^* = 2/(1 + j) + O(e)$ as $e \rightarrow 0$.*

Proof. Since $\log(1 - E) < 0$, taking logarithms in (19) shows that f is admissible if and only if $\theta^*(f) \geq \theta_{\min}(E, \pi)$, so the admissible set is determined by $\theta_{\min}$ alone, and by Lemma 10 it is the interval $(0, \bar{f}]$ with $\theta^*(\bar{f}) = \theta_{\min}$. The objectives $\mathbb{E}B_n = B_0(1 + ef)^n$ and $g(f)$ do not involve E or π , so their maximizers over this interval, namely $\bar{f}$ and $\min(\bar{f}, f^*)$ by Theorem 18, are functions of $\theta_{\min}$ alone. The remaining claims follow by substituting $\theta_{\min} = j$ into $\theta^*(\bar{f}) = \theta_{\min}$ and into (21). $\square$

The named fractions arise at odd integer exponents. The values $j = 1, 3, 5$, and 7 correspond to full, half, third, and quarter Kelly, and the corresponding curves $\pi = (1 - E)^j$ are developed as frontiers in Section 4.2.

Example 22. Consider a bettor who allocates half his bankroll to testing a strategy ($E = 0.50$) and will accept a 2% chance of abandoning a winning strategy ($\pi = 0.02$). Then $\theta_{\min} = \log 50 / \log 2 \approx 5.64$ and, by (21),

$$\bar{\alpha} = \frac{2 \log 2}{\log 2 + \log 50} \approx 0.30.$$

Because $\pi \leq 1 - E$, the optimal fractional Kelly is $\alpha = 0.30$ under either objective covered in Theorem 18. With a 10% edge at even money ($f^* = 0.10$), the optimal stake is $\bar{f} = \bar{\alpha}f^* = 3.0\%$ of bankroll per round. The exact Cramér-root computation gives $\bar{\alpha} = 30.2\%$.

Table 3: Optimal Kelly fraction $\bar{\alpha}(\pi, E)$, 10% edge at even money.

π	stop loss E		
	25%	50%	75%
0.10	22%	46%	75%
0.05	18%	38%	63%
0.02	14%	30%	52%

Table 4: Optimal Kelly fraction $\bar{\alpha}(\pi, E)$, 50% edge at 10-to-1.

π	stop loss E		
	25%	50%	75%
0.10	20%	43%	73%
0.05	16%	35%	61%
0.02	12%	28%	50%

Remark 23 (Betting more than Kelly). Under objective (i), loose constraints prescribe more than full Kelly. For instance, $\pi = 0.50$ and $E = 0.75$ give $\theta_{\min} = \log 2 / \log 4 = \frac{1}{2}$ and $\bar{\alpha} = 4/3$. So a bettor maximizing expected wealth, who tolerates a coin-flip chance of losing 75% of his bankroll, should stake 33% more than the Kelly prescription. As $\pi \uparrow 1$ the admissible stake $\bar{f}$ increases to the zero-growth stake f_c , for which $\bar{\alpha} \rightarrow 2$ in the diffusion limit, and only at $\pi = 1$ exactly does the constraint disappear, recovering the bet-it-all strategy for objective (i) and full Kelly for objective (ii).

Remark 24 (Full Kelly as a special case). Full Kelly has $\theta^* = 1$ at any odds (Corollary 8), so its premature stopping probability at stop loss E is exactly $1 - E$ up to the overshoot factor C_E . The bettor who sets $\pi = 0.50$ at a 50% stop loss is therefore prescribed exactly full Kelly. In particular, full Kelly staking is an acceptable strategy for a bettor who is indifferent to having a 50% probability of losing half of his bankroll and quitting a winning system; or, in general, a bettor who is indifferent to having a $1 - E$ probability of losing a fraction E of his bankroll and quitting a winning system. This makes the implicit risk posture of full Kelly staking explicit, and illustrates why practicing bettors tend to prefer smaller stakes.

Tables 3 and 4 report the optimal fraction, computed from the exact root of (10) via (20), for two running examples: a 10% edge at even money and a 50% edge at 10-to-1. We vary $\pi \in \{0.10, 0.05, 0.02\}$ and $E \in \{25\%, 50\%, 75\%\}$. Note that $\bar{\alpha} < 1$ within these ranges, so that the two objectives coincide.

First, the familiar heuristics emerge as special cases rather than universal constants. A moderate posture that accepts a 5% chance of quitting a winning strategy while refusing to risk more than half the bankroll gives approximately 38% Kelly, between the quarter- and half-Kelly rules of thumb. A stricter demand of $\pi = 0.02$ at the same stop loss suggests 30% Kelly.

Second, the optimal fraction depends on the odds. Long-shot odds (Table 4) call for slightly smaller fractions than the even-money case (Table 3) at the same (π, E) . The diffusion closed form (21) is independent of odds and matches Table 3 reasonably closely, while the large edge and skewed payoff of Table 4 pull the exact values 1.5–2.5 percentage points of Kelly fraction lower (e.g., 27.8% against the diffusion estimate 30.1% at $\pi = 0.02$, $E = 50\%$). Because the exact root sits below $2/\alpha - 1$ at high odds, the stake prescribed by (21) is slightly above the admissibility threshold, and the true premature stopping probability exceeds π . A practical rule of thumb is to use (21) at near-even odds and modest edges, and to solve (10) exactly (Remark 13) for long odds or large edges. At odds-on ($b \leq 1$) the first half of this rule is a theorem: the closed form fraction is admissible at every edge (Corollary 27).

Third, and most importantly, the inputs are not preferences over volatility. The bettor who chooses a larger E is willing to risk more capital, and the bettor who chooses a larger π is willing

Table 5: Implied premature stopping probability $\pi(\alpha, E) = (1 - E)^{\theta_\alpha^*}$, 10% edge at even money. Its reciprocal is the evidence carried by the stop loss.

α	stop loss E			
	25%	50%	75%	90%
100%	75%	50%	25%	10%
75%	62%	31%	9.9%	2.1%
50%	42%	12%	1.5%	0.098%
$33\frac{1}{3}\%$	24%	3.1%	0.096%	$9.7 \times 10^{-4}\%$
25%	13%	0.77%	$5.9 \times 10^{-3}\%$	$9.5 \times 10^{-6}\%$
10%	0.42%	$1.8 \times 10^{-4}\%$	$3.3 \times 10^{-10}\%$	$8.6 \times 10^{-18}\%$

to act on weaker evidence. Neither is a statement about risk aversion in the utility-theoretic sense.

Finally, note that no single Kelly fraction is optimal across all situations. The same (E, π) prescribes a different $\bar{\alpha}$ at each b . Therefore, a bettor who faces wagers at different odds cannot optimally apply one fraction of Kelly to all of them.

4.1 The trade-off in all three directions

Tables 3 and 4 fix (π, E) and solve for the optimal stake. But the constraint is a single identity linking three quantities, and can be inverted to answer other practically relevant questions. A bettor who already has a staking habit and a stop loss can ask what premature stopping probability he is implicitly adopting in doing so. A bettor who has settled on a stake and an evidence standard can ask how much capital must be at risk before the arrangement is self-consistent. We record these below.

Corollary 25 (The admissibility frontier). *Fix $e > 0$ and recall from (14) that θ_α^* denotes the Cramér root at stake αf^* . When equality holds in (19), i.e., (α, E, π) satisfies $(1 - E)^{\theta_\alpha^*} = \pi$, each coordinate is determined by the other two:*

$$\pi(\alpha, E) = (1 - E)^{\theta_\alpha^*}, \quad E(\alpha, \pi) = 1 - \pi^{1/\theta_\alpha^*}, \quad \theta_{\bar{\alpha}}^* = \theta_{\min}(\pi, E). \quad (22)$$

As $e \rightarrow 0$ with the remaining coordinates fixed, $\theta_\alpha^* = 2/\alpha - 1 + O(e)$ by Corollary 8, and each parameter has the closed form

$$\pi = (1 - E)^{\frac{2}{\alpha} - 1} (1 + O(e)), \quad 1 - E = \pi^{\frac{\alpha}{2 - \alpha}} (1 + O(e)), \quad \bar{\alpha} = \frac{2}{1 + \theta_{\min}} + O(e). \quad (23)$$

Proof. The first two expressions in (22) rearrange $(1 - E)^{\theta_\alpha^*} = \pi$ directly. The third is (20) with equality, solvable uniquely by Lemma 10. Substituting $\theta_\alpha^* = 2/\alpha - 1 + O(e)$ from Corollary 8 gives (23), since with (α, E, π) fixed the $O(e)$ exponent perturbations enter as $(1 - E)^{O(e)} = 1 + O(e)$ and $\pi^{O(e)} = 1 + O(e)$. $\square$

Figure 4 illustrates (22) for the even-money 10%-edge example, and Table 5 tabulates $\pi(\alpha, E)$ at the exact roots.

Table 5 shows the premature stopping probability implied by common staking habits. The half-Kelly bettor with a 50% stop loss is implicitly running $\pi = 12\%$. He will abandon a winning




Figure 4: The admissibility frontier $(1 - E)^{\theta_\alpha^*} = \pi$ of Corollary 25 read in each direction, for a 10% edge at even money (exact Cramér root throughout). (a) The optimal fraction $\bar{\alpha}$ as a function of the stop loss, one curve per tolerance π . The dotted line marks full Kelly, which is exceeded only for deep stop losses at loose tolerances. (b) The premature stopping probability implied by a stake and a stop loss (log scale). Steeper-than-exponential in α , so small reductions in stake buy orders of magnitude of evidence. (c) The stop loss required to hold a stake at a given tolerance. Near full Kelly the required E approaches total loss of the bankroll.

strategy about one time in eight, and when the stop loss is reached it carries only 8-to-1 evidence. Quarter-Kelly at the same stop loss runs $\pi = 0.77\%$ (130-to-1 evidence); tenth-Kelly makes premature stopping a practical impossibility at any stop loss of 50% or more. The full-Kelly row makes the perils of full Kelly clear, as even a bettor prepared to lose 90% of the bankroll retains a 10% chance of hitting that stop loss when the estimated edge is accurate. Reading down any column, the implied π falls faster than exponentially in the stake, demonstrating that small reductions in stake increase the amount of evidence at the stop loss by orders of magnitude.

The inversion $E(\alpha, \pi) = 1 - \pi^{1/\theta_\alpha^*}$ of (22) gives the stop loss required to support a given stake at a given tolerance, displayed in panel (c) of Figure 4. Fixing the evidence standard at 20-to-1, so $\pi = 0.05$, a bettor who insists on betting half Kelly must be prepared to lose 63% of the bankroll before quitting. At full Kelly the requirement is 95%, and at tenth Kelly it is 15%.

4.2 The exponent rule and its frontiers

By the exponent rule (Corollary 21), the specifications that prescribe a given fraction are the level sets $\pi = (1 - E)^{\theta_\alpha^*}$ of Corollary 25. For half Kelly, the diffusion limit ($\theta_{1/2}^* = 3$) gives the *cube-law frontier*

$$\pi = (1 - E)^3. \quad (24)$$

The half-Kelly row of Table 5 follows this curve: the exact entries 42% at $E = 25\%$, 12% at $E = 50\%$, and 1.5% at $E = 75\%$ lie closely to the cube-law values 42.2%, 12.5%, and 1.56%. More generally, the conventional fractions trace a power-law frontier in the (E, π) square with exponent $2/\alpha - 1$: for example, full Kelly ($\alpha = 1$) corresponds to the straight line $\pi = 1 - E$, exact at every payoff and edge because $\theta_1^* = 1$; quarter Kelly ($\alpha = 0.25$) corresponds to the seventh-power curve $\pi = (1 - E)^7$; and so on.

Theorem 26 (Power-law drawdown bounds at even money and shorter odds). *Let $0 < b \leq 1$ and $e \in (0, b)$, and for $\alpha \in (0, 2)$ write $f = \alpha f^* = \alpha e/b$ for the stake and $\theta_0 := 2/\alpha - 1$ for its*

diffusion exponent.

- (i) If $\alpha \in (0, 1)$, then $\psi_f(\theta_0) < 1$. Consequently $(B_0/B_n)^{\theta_0}$ is a nonnegative supermartingale, $\theta_\alpha^* > \theta_0$, and for every stop loss $E \in (0, 1)$,

$$\mathbb{P}_e(\tau_E < \infty) \leq (1 - E)^{2/\alpha - 1} \quad \text{and} \quad \text{LR}(E) \geq (1 - E)^{-(2/\alpha - 1)}. \quad (25)$$

If $\alpha = 1$, then $\psi_f(1) = 1$, the full-Kelly identity of Corollary 8. If $\alpha \in (1, 2)$ and $e < b/\alpha$, the inequality reverses, $\psi_f(\theta_0) > 1$, so $\theta_\alpha^* < \theta_0$ whenever $g(f) > 0$.

- (ii) Half Kelly ($\alpha = 1/2$) satisfies the cube law $\mathbb{P}_e(\tau_E < \infty) \leq (1 - E)^3$ at every edge and every payoff $b \leq 1$. At $b = 1$ the difference $1 - \psi_f(3)$ has the exact value

$$1 - \psi_{e/2}(3) = \frac{x^4(5 - x^2)}{(1 - x^2)^3}, \quad x := e/2. \quad (26)$$

- (iii) At $b = 1$, for $\alpha \in (0, 1)$,

$$\theta_\alpha^* = \theta_0 + \frac{(\theta_0 - 1)(\theta_0 + 2)}{3(\theta_0 + 1)} e^2 + O(e^4), \quad (27)$$

in particular $\theta_{1/2}^* = 3 + \frac{5}{6}e^2 + O(e^4)$ and $\theta_{1/4}^* = 7 + \frac{9}{4}e^2 + O(e^4)$. At an arbitrary payoff $b > 0$, for $\alpha \in (0, 1)$,

$$\theta_\alpha^* = \theta_0 - \frac{(\theta_0 - 1)(b - 1)}{3b} e + O(e^2). \quad (28)$$

Proof. (i) At payoff b we have $p = (1 + e)/(b + 1)$ and $e = bcf$ with $c := 1/\alpha = (\theta_0 + 1)/2$, so

$$(b + 1)\psi_f(\theta_0) = (1 + bcf)(1 + bf)^{-\theta_0} + b(1 - cf)(1 - f)^{-\theta_0} =: G_b(f),$$

with $G_b(0) = 1 + b$, and the claim is that $G_b < 1 + b$ when $\theta_0 > 1$. Differentiating and substituting $c - \theta_0 = \frac{1 - \theta_0}{2}$ and $c(1 - \theta_0) = \frac{(1 - \theta_0)(\theta_0 + 1)}{2}$,

$$\begin{aligned} G'_b(f) &= b \left\{ (1 + bf)^{-\theta_0 - 1} [c(1 + bf) - \theta_0(1 + bcf)] + (1 - f)^{-\theta_0 - 1} [\theta_0(1 - cf) - c(1 - f)] \right\} \\ &= \frac{b(1 - \theta_0)}{2} [\varphi(bf) - \varphi(-f)], \quad \text{where } \varphi(t) := (1 + t)^{-\theta_0 - 1} (1 + (\theta_0 + 1)t). \end{aligned}$$

Since $\varphi'(t) = -\theta_0(\theta_0 + 1)t(1 + t)^{-\theta_0 - 2}$, the function φ is strictly decreasing on $(0, \infty)$, and

$$\varphi(f) - \varphi(-f) = \int_0^f [\varphi'(t) + \varphi'(-t)] dt = \theta_0(\theta_0 + 1) \int_0^f t [(1 - t)^{-\theta_0 - 2} - (1 + t)^{-\theta_0 - 2}] dt > 0$$

for every $f \in (0, 1)$ and every $\theta_0 > 0$, the integrand being strictly positive on $(0, 1)$. For $b \leq 1$ we have $bf \leq f$, so monotonicity of φ gives

$$\varphi(bf) \geq \varphi(f) > \varphi(-f),$$

and therefore $G'_b < 0$ throughout when $\theta_0 > 1$, $G'_b \equiv 0$ when $\theta_0 = 1$ (recovering the full-Kelly identity), and $G'_b > 0$ when $\theta_0 \in (0, 1)$, which proves the three cases of (i) at the level of $\psi_f(\theta_0)$. The consequences follow as in Section 3: $M_n := (B_0/B_n)^{\theta_0} = \exp(-\theta_0 S_n)$ satisfies $\mathbb{E}_e[M_n | \mathcal{F}_{n-1}] = M_{n-1} \psi_f(\theta_0) \leq M_{n-1}$, a nonnegative supermartingale with $M_0 = 1$, so Ville's

maximal inequality [22] gives $\mathbb{P}_e(\sup_n M_n \geq \lambda) \leq 1/\lambda$; taking $\lambda = (1 - E)^{-\theta_0}$ and noting that $\tau_E < \infty$ implies $M_{\tau_E} \geq (1 - E)^{-\theta_0}$ yields the first bound in (25), and the second follows from Lemma 5 via (9). Finally, since ψ_f is strictly convex in θ with $\psi_f(0) = 1$ and strictly negative slope at $\theta = 0$ (proof of Theorem 7), $\psi_f(\theta) < 1$ holds exactly for $0 < \theta < \theta^*(f)$; thus $\psi_f(\theta_0) < 1$ forces $\theta_\alpha^* > \theta_0$, and $\psi_f(\theta_0) > 1$ forces $\theta_\alpha^* < \theta_0$ whenever the root exists.

(ii) The cube law is (i) with $\alpha = \frac{1}{2}$. For the identity, take $b = 1$, so $\theta_0 = 3$, $f = x = e/2$, and $c = 2$. Expanding the two products in G_1 ,

$$(1 + 2x)(1 - x)^3 = 1 - x - 3x^2 + 5x^3 - 2x^4, \quad (1 - 2x)(1 + x)^3 = 1 + x - 3x^2 - 5x^3 - 2x^4,$$

so

$$\psi_f(3) = \frac{(1 + 2x)(1 - x)^3 + (1 - 2x)(1 + x)^3}{2(1 - x^2)^3} = \frac{2 - 6x^2 - 4x^4}{2(1 - x^2)^3},$$

and subtracting from 1, using $2(1 - x^2)^3 = 2 - 6x^2 + 6x^4 - 2x^6$, gives

$$1 - \psi_f(3) = \frac{10x^4 - 2x^6}{2(1 - x^2)^3} = \frac{x^4(5 - x^2)}{(1 - x^2)^3},$$

which is (26).

(iii) At $b = 1$, writing $\psi_f(\theta)$ as $\frac{1}{2}[(1 + f)^{-\theta} + (1 - f)^{-\theta}] - \frac{e}{2}[(1 - f)^{-\theta} - (1 + f)^{-\theta}]$ with $e = f/\alpha$, the binomial series gives

$$\psi_f(\theta) = 1 + \left[\frac{\theta(\theta+1)}{2} - \frac{\theta}{\alpha} \right] f^2 + \left[\frac{\theta(\theta+1)(\theta+2)(\theta+3)}{24} - \frac{\theta(\theta+1)(\theta+2)}{6\alpha} \right] f^4 + O(f^6),$$

uniformly for θ in a neighborhood of θ_0 . At $\theta = \theta_0$ the f^2 coefficient vanishes. Using $1/\alpha = (\theta_0 + 1)/2$,

$$1 - \psi_f(\theta_0) = \frac{\theta_0(\theta_0 + 1)(\theta_0 + 2)(\theta_0 - 1)}{24} f^4 + O(f^6), \quad \frac{\partial \psi_f}{\partial \theta}(\theta_0) = \frac{\theta_0}{2} f^2 + O(f^4).$$

Since ψ_f is smooth and strictly increasing at its root,

$$\theta_\alpha^* - \theta_0 = \frac{1 - \psi_f(\theta_0)}{\frac{\partial \psi_f}{\partial \theta}(\theta_0)} (1 + o(1)) = \frac{(\theta_0 + 1)(\theta_0 + 2)(\theta_0 - 1)}{12} f^2 + o(f^2),$$

and substituting $f = 2e/(\theta_0 + 1)$ gives (27); $\theta_0 = 3$ and $\theta_0 = 7$ give the half- and quarter-Kelly values. At a general payoff, the same program with the exact moments $pb^2 + q = b + e(b - 1)$ and $pb^3 - q = b(b - 1) + e(b^2 - b + 1)$ gives

$$\psi_f(\theta_0) - 1 = \frac{2\theta_0(\theta_0 - 1)(b - 1)}{3b^2(\theta_0 + 1)^2} e^3 + O(e^4), \quad \frac{\partial \psi_f}{\partial \theta}(\theta_0) = \frac{2\theta_0}{b(\theta_0 + 1)^2} e^2 + O(e^3),$$

whence $\theta_\alpha^* - \theta_0 = -\frac{(\theta_0 - 1)(b - 1)}{3b} e + O(e^2)$, which is (28). $\square$

Table 6 illustrates the theorem across stakes, payoffs, and edges. For $b \leq 1$ every exact exponent exceeds its diffusion value θ_0 , as part (i) guarantees. At $b = 1$ the excess follows the second-order formula (27): for instance $\theta_{1/2}^* = 3.034$ at $e = 0.20$, compared to the approximation $3 + \frac{5}{6}(0.20)^2 = 3.033$. At $b = \frac{1}{2}$ the excess is first order in the edge, per (28), and larger throughout ($\theta_{1/2}^* = 3.252$ at $e = 0.20$). For $b > 1$ the exact exponent falls below θ_0 , with the deficit tracking

Table 6: Exact Cramér root θ_α^* against the power-law exponent $\theta_0 = 2/\alpha - 1$ of Theorem 26, at half Kelly ($\theta_0 = 3$) and quarter Kelly ($\theta_0 = 7$), across payoffs b and edges e . Entries exceed θ_0 exactly when $b \leq 1$ (Theorem 26(i)); for $b > 1$ the deficit is governed to first order by (28).

payoff b	half Kelly ($\theta_0 = 3$)				quarter Kelly ($\theta_0 = 7$)			
	$e = 0.05$	$e = 0.10$	$e = 0.20$	$e = 0.40$	$e = 0.05$	$e = 0.10$	$e = 0.20$	$e = 0.40$
0.5	3.039	3.091	3.252	4.182	7.116	7.267	7.719	10.180
1	3.002	3.008	3.034	3.149	7.006	7.023	7.092	7.400
2	2.985	2.972	2.953	2.939	6.954	6.914	6.854	6.800
10	2.971	2.945	2.897	2.820	6.913	6.833	6.688	6.449

(28): at $b = 10$, $e = 0.20$, the first-order formula gives $3 - \frac{2.9 \cdot (0.20)}{30} = 2.88$ against the exact 2.897. In probability terms, at a 50% stop loss and $e = 0.20$, the half-Kelly cube law $(1 - E)^3 = 12.5\%$ compares with exact premature stopping probabilities (neglecting overshoot) of 10.5% at $b = \frac{1}{2}$, 12.2% at $b = 1$, and 13.4% at $b = 10$. The bound (25) holds at $b \leq 1$ with a margin that widens as b falls, while for $b > 1$ it understates the true probability by an amount governed by (28).

Theorem 26(i) extends the full-Kelly reciprocal bankroll martingale to a one-parameter family. For $b \leq 1$, betting the fraction α of Kelly makes $(B_0/B_n)^{2/\alpha-1}$ a supermartingale, with the martingale case at full Kelly. In particular, the closed form fraction of Theorem 18 then satisfies the (E, π) constraint at every edge.

Corollary 27 (Admissibility of the closed form fraction when $b \leq 1$). *Fix $E, \pi \in (0, 1)$ with $\pi \leq 1 - E$ and let $\bar{\alpha} = 2/(1 + \theta_{\min})$ be the closed form fraction (21). At any payoff $b \leq 1$, the stake $\bar{\alpha}f^*$ satisfies the admissibility constraint (19) at every edge:*

$$\mathbb{P}_e(\tau_E < \infty) \leq (1 - E)^{\theta_{\min}} = \pi \quad \text{for every } e \in (0, b).$$

Proof. $\pi \leq 1 - E$ gives $\theta_{\min} \geq 1$, hence $\bar{\alpha} \leq 1$ and $2/\bar{\alpha} - 1 = \theta_{\min}$. By Theorem 26(i), with the boundary case $\bar{\alpha} = 1$ covered by the full-Kelly identity, the process $(B_0/B_n)^{\theta_{\min}}$ is a nonnegative supermartingale, and Ville's inequality gives $\mathbb{P}_e(\tau_E < \infty) \leq (1 - E)^{\theta_{\min}} = \pi$. $\square$

Remark 28 (Integer exponents and long odds). (i) At integer exponents the difference $1 - \psi_f(\theta_0)$ has a closed polynomial form. For example, at quarter Kelly, with $b = 1$, $\theta_0 = 7$, and $x := e/4$,

$$2(1 - x^2)^7(1 - \psi_f(7)) = 252x^4 + 84x^6 + 2x^8(39 - 21x^2 + 7x^4 - x^6).$$

(ii) By (28), for $b > 1$ the exponent falls short of θ_0 at first order in the edge, so a bettor who sizes by the diffusion power law overbets. Solving $\theta_0 - \frac{(\theta_0-1)(b-1)}{3b}e = \theta_{\min}$ for the stake yields the skew-corrected fraction

$$\bar{\alpha}_b = \bar{\alpha} \left[1 - \frac{(\theta_{\min} - 1)(b - 1)}{3b(1 + \theta_{\min})} e \right] + O(e^2),$$

a practical substitute for solving (10) exactly. In the worked example pushed to long odds ($\pi = 0.02$, $E = 50\%$, $b = 10$, $e = 0.5$) it prescribes 26.9% of Kelly against the exact 27.8% and the uncorrected diffusion value 30.1%.

4.3 Rules of thumb

This subsection collects some practical consequences of the preceding results.

Remark 29 (The staking formula in bits). The closed form (21) reduces to mental arithmetic in special cases. Write $m(E) \equiv m := -\log_2(1 - E)$ for the number of halvings of the bankroll that the stop loss represents ($E = 1 - 2^{-m}$) and $k := -\log_2 \pi$ for the weight of evidence, in bits, that the bettor demands of the stop loss. Then $\theta_{\min} = k/m$ and

$$\bar{\alpha} = \frac{2m}{m + k}.$$

A bettor prepared to lose half the bankroll ($m = 1$) who demands 100-to-1 evidence ($k = \log_2 100 \approx 6.6$) is prescribed $2/7.6 \approx 26\%$ of Kelly; one prepared to lose $3/4$ ($m = 2$) at the same standard is prescribed $4/8.6 \approx 46\%$.

By Theorem 7, a drawdown of j halvings carries likelihood ratio $2^{j\theta^*}$, so a stake with Cramér root θ^* earns θ^* bits of evidence per halving of bankroll, and the admissibility constraint (20) requires the stake to deliver at least k/m bits per halving. Full Kelly delivers exactly one bit per halving, at every payoff and every return distribution (Corollary 8 and Remark 39). At even money or shorter odds, half Kelly delivers at least three and quarter Kelly at least seven bits of information per halving (Theorem 26). In this framing, the coincidence condition $\pi \leq 1 - E$ of Corollary 20 becomes $k \geq m$.

Remark 30 (A significance level interpretation). The premature stopping probability can be interpreted as a significance level for the test of the null hypothesis H_0 that the edge is $e > 0$ against the alternative H_1 that there is no edge ($e \leq 0$). In statistical parlance, erroneously rejecting the null hypothesis is called a Type I error, and the significance level of the test is the probability of committing a Type I error. By construction, any sequence of bets sized based on (E, π) has a probability of at most π of reaching the stop loss when the estimated edge is correct (i.e., committing a Type I error). Thus, π defines the significance level of the corresponding hypothesis test.

Conventional choices of significance level in many scientific fields are $\pi = 0.05$ or $\pi = 0.01$. A bettor who adopts one of these conventions $\pi = 0.05$ or $\pi = 0.01$, together with a stop loss $E = \frac{1}{2}$, is prescribed

$$\bar{\alpha}(0.05, \tfrac{1}{2}) = \frac{2 \log 2}{\log 2 + \log 20} \approx 0.38, \quad \bar{\alpha}(0.01, \tfrac{1}{2}) = \frac{2 \log 2}{\log 2 + \log 100} \approx 0.26.$$

The 5% convention lands below half Kelly, and the 1% convention agrees closely with quarter-Kelly. Table 7 tabulates $\bar{\alpha} = 2m/(m + k)$ over the dyadic grid $E \in \{\frac{1}{8}, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, \frac{7}{8}\}$, $\pi \in \{\frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \frac{1}{32}, \frac{1}{128}\}$. The named fractions appear as exact entries (e.g., half Kelly at $(E, \pi) = (\frac{1}{2}, \frac{1}{8})$, third Kelly at $(\frac{1}{2}, \frac{1}{32})$, quarter Kelly at $(\frac{1}{2}, \frac{1}{128})$) and full Kelly runs down the diagonal where $k = m$. Entries above 100% occur where $\pi > 1 - E$: there the expected-wealth objective prescribes more than full Kelly and the growth objective caps $\bar{\alpha}$ at 100% (Theorem 18).

4.4 Reaching a target before the stop loss

The Kelly criterion targets the long run growth rate, but bettors sometimes have shorter term objectives. A famous result of Dubins and Savage [6] concerns the gambler who seeks to maximize

Table 7: The staking formula over dyadic conventions: $\bar{\alpha} = 2m/(m+k)$ (diffusion limit) for stop losses $E = 1 - 2^{-m}$ and premature stopping probabilities $\pi = 2^{-k}$. Each row lists the evidence $1/\pi$ that the stop loss carries. Entries above 100% (where $\pi > 1 - E$) apply to the expected-wealth objective and are capped at full Kelly under the growth objective.

π ($1/\pi$)	stop loss E (halvings m)				
	$\frac{1}{8}$ (0.19)	$\frac{1}{4}$ (0.42)	$\frac{1}{2}$ (1)	$\frac{3}{4}$ (2)	$\frac{7}{8}$ (3)
1/4 (4)	18%	34%	67%	100%	120%
1/8 (8)	12%	24%	50%	80%	100%
1/16 (16)	9.2%	19%	40%	67%	86%
1/32 (32)	7.4%	15%	33%	57%	75%
1/128 (128)	5.4%	11%	25%	44%	60%

the probability of doubling an initial fortune at an unfavorable game, for which the optimal strategy is bold play, staking everything available or just enough to reach the target. Here we generalize this consideration to the setting of this paper. The bettor specifies (E, π) together with the objective of multiplying the initial bankroll by a factor $M > 1$ before reaching the stop loss.

For the target multiple $M > 1$, let

$$\tau_M := \inf\{n \geq 1 : B_n \geq MB_0\} = \inf\{n \geq 1 : S_n \geq \log M\} \quad (29)$$

be the first time the bankroll reaches M times its starting value. The theorem below gives the probability of reaching the target before the stop loss at an arbitrary stake f , and does not involve π . The tolerance π enters in Corollary 35, where the stake is set to the prescription of Theorem 18.

Theorem 31 (Probability of reaching a target before the stop loss). *Suppose $e > 0$ and $f \in (0, f_c)$, let $\theta^* = \theta^*(f)$ be the Cramér root of (10), and fix $E \in (0, 1)$ and $M > 1$. Then $\tau_E \wedge \tau_M < \infty$ almost surely, and*

$$\mathbb{P}_e(\tau_M < \tau_E) = \frac{1 - D_E^{-1}(1 - E)^{\theta^*}}{1 - \frac{D_M}{D_E}\left(\frac{1-E}{M}\right)^{\theta^*}}, \quad (30)$$

where

$$D_E := \mathbb{E}_e[\exp(-\theta^*(S_{\tau_E} - \log(1-E))) \mid \tau_E < \tau_M], \quad D_M := \mathbb{E}_e[\exp(-\theta^*(S_{\tau_M} - \log M)) \mid \tau_M < \tau_E] \quad (31)$$

are the conditional expectations of the exponential overshoots given $\tau_E < \tau_M$ and $\tau_M < \tau_E$, respectively, and satisfying

$$1 \leq D_E \leq (1 - f)^{-\theta^*}, \quad (1 + bf)^{-\theta^*} \leq D_M \leq 1. \quad (32)$$

Proof. Since $g(f) > 0$, $S_n \rightarrow +\infty$ almost surely, so $\tau_M < \infty$ a.s. and the race resolves. Stop the martingale $M_n = \exp(-\theta^* S_n)$ of Theorem 7 at $\tau := \tau_E \wedge \tau_M$. Before τ the walk is confined to the strip $\log(1 - E) < S_n < \log M$, and the single overshooting step is bounded as in the proof of Theorem 7, so the stopped martingale is uniformly bounded and bounded convergence gives

$$1 = \mathbb{E}_e[\exp(-\theta^* S_{\tau_E}) \mathbf{1}\{\tau_E < \tau_M\}] + \mathbb{E}_e[\exp(-\theta^* S_{\tau_M}) \mathbf{1}\{\tau_M < \tau_E\}].$$

Writing $\bar{q} := \mathbb{P}_e(\tau_E < \tau_M)$ and conditioning each term on the event it carries,

$$1 = \bar{q} D_E (1 - E)^{-\theta^*} + (1 - \bar{q}) D_M M^{-\theta^*},$$

with D_E and D_M as in (31). Solving this linear equation for $1 - \bar{q} = \mathbb{P}_e(\tau_M < \tau_E)$ and multiplying numerator and denominator by $(1 - E)^{\theta^*}/D_E$ gives (30). The bounds (32) hold because each barrier is missed by less than one step: given $\tau_E < \tau_M$, the crossing step is a loss, so $\log(1 - E) + d < S_{\tau_E} \leq \log(1 - E)$ and the overshoot exponent $-\theta^*(S_{\tau_E} - \log(1 - E))$ lies in $[0, -\theta^*d]$; given $\tau_M < \tau_E$, the crossing step is a win, so $\log M \leq S_{\tau_M} < \log M + u$ and the exponent lies in $(-\theta^*u, 0]$. Letting $M \rightarrow \infty$ in (30) recovers Theorem 7: the event $\{\tau_E < \tau_M\}$ increases to $\{\tau_E < \infty\}$, so $D_E \rightarrow C_E$, while the denominator of (30) tends to 1. $\square$

Corollary 32 (Closed form in the small-edge limit). *Let $f = \alpha f^*$. As $e \rightarrow 0$ with E and M fixed,*

$$\mathbb{P}_e(\tau_M < \tau_E) = \frac{1 - (1 - E)^{\theta^*}}{1 - \left(\frac{1 - E}{M}\right)^{\theta^*}} (1 + O(e)). \quad (33)$$

Proof. For $f = \alpha f^*$ both θ^*u and $-\theta^*d$ are $O(e)$ (proof of Corollary 8), so the bounds (32) give $D_E = 1 + O(e)$ and $D_M = 1 + O(e)$, and substituting $D_E = D_M = 1$ into (30) gives (33). For the monotonicity, set $A := -\log(1 - E) > 0$ and $B := \log M > 0$, so that the ratio in (33) is $\varphi(\theta) := (1 - \exp(-\theta A))/(1 - \exp(-\theta(A + B)))$. Then

$$\frac{d}{d\theta} \log \varphi(\theta) = \frac{A}{\exp(\theta A) - 1} - \frac{A + B}{\exp(\theta(A + B)) - 1} > 0,$$

because $x \mapsto x/(\exp(x) - 1)$ is strictly decreasing on $(0, \infty)$ and $\theta A < \theta(A + B)$. First-order expansion of the exponentials gives $\varphi(\theta) \rightarrow A/(A + B)$ as $\theta \downarrow 0$, and $\varphi(\theta) \rightarrow 1$ as $\theta \rightarrow \infty$. Lemma 10 maps these limits to $f \uparrow f_c$ and $f \downarrow 0$ respectively; the $\theta \downarrow 0$ limit is the classical gambler's-ruin probability of a driftless walk, as befits the zero-growth stake. Substituting $\theta^* = 1$ into the ratio gives $E/(1 - \frac{1-E}{M}) = EM/(M - 1 + E)$, which equals $\frac{2}{3}$ at $E = \frac{1}{2}$, $M = 2$. $\square$

Remark 33. The ratio in (33) is strictly increasing in θ^* and therefore, by Lemma 10, strictly decreasing in the stake f . It increases to 1 as $f \downarrow 0$ and decreases to $(-\log(1 - E))/(\log M - \log(1 - E))$ as $f \uparrow f_c$. At full Kelly, the ratio equals $EM/(M - 1 + E)$. In particular, the probability of doubling the bankroll before losing half is $\frac{2}{3} + O(e)$, where the value $\frac{2}{3}$ does not depend on b .

Table 8 confirms (33) by Monte Carlo for the running even-money example. A full Kelly better doubles his bankroll before halving it approximately 2/3 of the time; the quarter-Kelly better does so 99.2% of the time; and the probability only improves as the stake shrinks further.

Because the target probability is decreasing in the stake while both objectives of Theorem 18 are increasing in it (up to full Kelly), maximizing (33) over the admissible stakes $(0, \bar{f}]$ is a degenerate problem: the supremum is approached as $f \downarrow 0$, where the probability tends to 1. Taken alone, then, “maximize the probability of reaching M before the stop loss” prescribes no bet at all. This reflects that a better with an edge who bets timidly enough reaches every target eventually, and the constraint (E, π) , which caps the stake from above, is never reached. Betting less lengthens the time needed to reach the target.

Table 8: Probability of doubling the bankroll before hitting a 50% stop loss ($M = 2$, $E = 0.50$), under fractional Kelly betting with a 10% edge at even money. Simulated values are Monte Carlo estimates (2×10^5 paths, run to $N = 6,000$ rounds, by which point every path has resolved the race); closed form values are (33). The probability *falls* as the stake rises.

Kelly fraction α	θ^*	simulated	closed form (33)
100%	1.00	67.9%	66.7%
50%	3.01	89.6%	88.9%
30%	5.68	98.2%	98.1%
25%	7.02	99.3%	99.2%
10%	19.1	100.0%	100.0%

Remark 34 (Probability versus speed of reaching the target). To reach the target as surely as possible, the bettor should stake as little as patience allows. By Theorem 31 the probability increases to 1 as $f > 0$ decreases. To reach it as quickly as possible, maximize the growth rate: by Wald’s identity $\mathbb{E}_e[S_{\tau_M}] = g(f) \mathbb{E}_e[\tau_M]$, and $S_{\tau_M} \approx \log M$, so $\mathbb{E}_e[\tau_M] \approx \log M / g(f)$, which is minimized over admissible stakes at the optimum $\min(\bar{\alpha}, 1)f^*$ of Theorem 18(ii). This recovers, inside the constrained framework, Breiman’s classical characterization of Kelly staking as the strategy that minimizes the expected time to reach a distant wealth goal [2].

Corollary 35 (Target guarantee at the optimal stake). *Let $\bar{f}$ be the unique stake with $\theta^*(\bar{f}) = \theta_{\min}$, as in Theorem 18, so that $(1 - E)^{\theta^*(\bar{f})} = \pi$. Then, setting $D_E = D_M = 1$ as in (33), for every target $M > 1$,*

$$\mathbb{P}_e(\tau_M < \tau_E) = \frac{1 - \pi}{1 - \pi M^{-\theta_{\min}}} \geq 1 - \pi, \quad (34)$$

which decreases in M with limit $1 - \pi$ as $M \rightarrow \infty$. In particular, for the target $M = 1/(1 - E)$, we have $M^{-\theta_{\min}} = \pi$ and

$$\mathbb{P}_e(\tau_M < \tau_E) = \frac{1}{1 + \pi}.$$

Proof. Substitute $(1 - E)^{\theta_{\min}} = \pi$ into (33), noting $(\frac{1-E}{M})^{\theta_{\min}} = \pi M^{-\theta_{\min}}$. The denominator of (34) is at most 1 and increases to 1 as $M \rightarrow \infty$, giving the bound and the limit. For $M = 1/(1 - E)$, $M^{-\theta_{\min}} = (1 - E)^{\theta_{\min}} = \pi$, and $(1 - \pi)/(1 - \pi^2) = 1/(1 + \pi)$. $\square$

Corollary 35 gives the premature stopping probability yet another interpretation. By Theorem 7, π is (up to overshoot) the probability of ever hitting the stop loss at the optimal stake; therefore, $1 - \pi$ is the probability that the bettor is never stopped at all, in which case the bankroll eventually exceeds every finite target. The guarantee (34) interpolates the two extremes: the race to any finite M is won with probability at least $1 - \pi$. For the worked example ($\pi = 0.02$, $E = 0.50$, $\bar{\alpha} \approx 0.30$), the bettor doubles his bankroll before reaching the stop loss with probability $1/1.02 \approx 98.0\%$, and reaches any multiple with probability at least 98%.

4.5 The stop loss as a sequential test

Under the edge hypothesis, $\exp(-\theta^* S_n)$ is a mean-1 nonnegative martingale by Theorem 7, and the stop loss is the Ville event $\{\sup_n \exp(-\theta^* S_n) \geq (1 - E)^{-\theta^*} = 1/\pi\}$, an anytime-valid rejection of the edge hypothesis at level π (exactly π neglecting overshoot), with power 1 by Lemma 5

[16, 18, 19, 22]. In the language of sequential analysis, the (E, π) framework sizes bets so that the strategy's own failure is a level- π , power-1 sequential test of the hypothesis that motivated it. Section 2 bounds the capital that the test consumes. The next proposition bounds its duration under the null, namely the expected number of rounds until the stop loss is hit when there is no edge.

Proposition 36 (Expected time to the stop loss under no edge). *Fix $f \in (0, 1)$, write $A := -\log(1 - E)$ and $g_0 := p_0 \log(1 + bf) + q_0 \log(1 - f) < 0$ for the log-drift under $\mathbb{P}_0$, where $p_0 = 1/(b + 1)$ and $q_0 := 1 - p_0$. Then $\mathbb{E}_0[\tau_E] < \infty$ and*

$$\frac{A}{|g_0|} \leq \mathbb{E}_0[\tau_E] \leq \frac{A - \log(1 - f)}{|g_0|}. \quad (35)$$

For $f = \alpha e/b$, $|g_0| = \alpha^2 e^2/(2b) + O(e^3)$, so $\mathbb{E}_0[\tau_E] = 2bA/(\alpha^2 e^2)(1 + O(e))$.

Proof. That $g_0 < 0$ is Lemma 5 at $e = 0$. The process $S_n - ng_0$ is a $\mathbb{P}_0$ -martingale with bounded increments, and on $\{n \leq \tau_E\}$ we have $S_n > \log(1 - E) + d$; optional stopping at $\tau_E \wedge n$ gives $\mathbb{E}_0[\tau_E \wedge n] |g_0| = -\mathbb{E}_0[S_{\tau_E \wedge n}] \leq A - d$, and monotone convergence yields $\mathbb{E}_0[\tau_E] \leq (A - d)/|g_0| < \infty$. Wald's identity then applies exactly: $\mathbb{E}_0[S_{\tau_E}] = g_0 \mathbb{E}_0[\tau_E]$, and the crossing step satisfies $\log(1 - E) + d < S_{\tau_E} \leq \log(1 - E)$, which is (35). The expansion of g_0 uses $p_0 b - q_0 = 0$ and $p_0 b^2 + q_0 = b$: $g_0 = -\frac{1}{2}bf^2 + O(f^3)$, and $f = \alpha e/b$ gives $|g_0| = \alpha^2 e^2/(2b) + O(e^3)$. $\square$

A bettor with no edge loses at most the fraction E of the bankroll, and the expected number of rounds until the stop loss is hit is about $2bA/(\alpha^2 e^2)$. In the worked example ($\alpha = 0.30$, $e = 0.10$, $b = 1$, $E = 50\%$), the bounds (35) put this expectation between 1,540 and 1,610 rounds (Monte Carlo estimate 1,546), roughly six times the $\log 2/g \approx 270$ rounds the same stake needs to double the bankroll when the edge is real (Remark 34). Halving the stake quadruples the expected time to the stop loss, since it grows as $1/\alpha^2$, while the premature stopping probability improves faster than exponentially (Table 5).

5 Betting with an overestimated edge

In theory, the Kelly criterion bets most when the edge is highest. In practice, it bets most when adverse selection is highest. Estimates large enough to prescribe aggressive staking are the most likely to be inflated, because the price is based on the collective assessment of all counterparties in the market.

Suppose that the bettor sizes bets to an estimated edge e , staking $f = \alpha f^*(e) = \alpha e/b$ for some $0 < \alpha < 1$, while the true edge is κe for $\kappa \in (0, 1]$, so that the true law is $\mathbb{P}_{\kappa e}$ in the notation of Section 3. The entire preceding analysis transfers through one elementary observation.

Lemma 37 (Effective Kelly fraction). *Staking $\alpha f^*(e)$ when the true edge is κe is identical to staking $(\alpha/\kappa) f^*(\kappa e)$. In particular, under $\mathbb{P}_{\kappa e}$ the bettor is an (α/κ) -fractional Kelly bettor, and every result of Sections 3–4 applies with α replaced by the effective fraction $\alpha_{\text{eff}} = \alpha/\kappa$.*

Proof. The Kelly criterion prescribes a fraction $f^*(e) = e/b$ for edge $e > 0$ and odds $b > 0$. Assuming an edge of κe , therefore, gives a Kelly fraction of $f^*(\kappa e) = \kappa e/b$. Betting $\alpha f^*(e) = \alpha e/b$ under the assumption of edge e is, therefore, equivalent to

$$\alpha e/b = (\alpha/\kappa)(\kappa e/b) = (\alpha/\kappa)f^*(\kappa e).$$

$\square$

Overestimating the edge by a factor $1/\kappa$ therefore inflates the bettor's effective Kelly fraction by the same factor. A half-Kelly bettor whose true edge is half his estimate is unwittingly betting full Kelly, and the full-Kelly bettor in the same position is betting twice Kelly.

Theorem 38 (Bet sizing with an overestimated edge). *Let $f = \alpha f^*(e)$ with true edge κe for $\kappa > 0$, and let $e \rightarrow 0$ with b, α, κ, E , and π fixed.*

- (i) *There is a threshold $\kappa_c(\alpha) = \alpha/2 + O(e)$ such that the growth rate is negative for $\kappa < \kappa_c(\alpha)$ and positive for $\kappa > \kappa_c(\alpha)$.*
- (ii) *For $\kappa > \alpha/2$, the Cramér root is $\theta^* = \frac{2\kappa}{\alpha} - 1 + O(e)$, and the premature stopping probability at stop loss E is $(1 - E)^{2\kappa/\alpha - 1} (1 + O(e))$.*
- (iii) *At the stake $\alpha = \bar{\alpha}(E, \pi)$ prescribed by Theorem 18, the premature stopping probability under the true edge is*

$$\pi_\kappa = \pi^\kappa (1 - E)^{\kappa - 1} (1 + O(e)), \quad \kappa(1 + \theta_{\min}) > 1. \quad (36)$$

- (iv) (Robust staking rule.) *Fix $\kappa_{\min} \in (0, 1]$. The exponent $2\kappa/\alpha - 1$ in (ii) satisfies $2\kappa/\alpha - 1 \geq \theta_{\min}$ for every $\kappa \geq \kappa_{\min}$ if and only if $\alpha \leq 2\kappa_{\min}/(1 + \theta_{\min})$, and $2\kappa_{\min}/(1 + \theta_{\min}) = \kappa_{\min} \bar{\alpha}(1 + O(e))$ by (21).*

Proof. By Lemma 37 the bettor is an (α/κ) -fractional Kelly bettor with respect to the true edge κe , so Corollary 8 applies at the effective fraction $\alpha_{\text{eff}} = \alpha/\kappa$ and gives $\theta^* = 2\kappa/\alpha - 1 + O(e)$ together with $C_E = 1 + O(e)$, which is (ii). For (i), at fixed stake f the growth rate $g = p \log(1 + bf) + q \log(1 - f)$ is strictly increasing in the true edge, because p is increasing in it and $\log(1 + bf) > \log(1 - f)$. It is negative at $\kappa = 0$ by Lemma 5 and positive for κ large enough, so a unique root $\kappa_c(\alpha)$ exists. Expanding to second order as in the proof of Corollary 8 gives $g = (\kappa - \alpha/2) \alpha e^2/b + O(e^3)$, and setting this to zero gives $\kappa_c(\alpha) = \alpha/2 + O(e)$. For (iii), substituting $\bar{\alpha} = 2/(1 + \theta_{\min}) + O(e)$ from (21) into (ii) gives $\theta^* = \kappa(1 + \theta_{\min}) - 1 + O(e)$, so

$$(1 - E)^{\theta^*} = [(1 - E)^{\theta_{\min}}]^\kappa (1 - E)^{\kappa - 1} (1 - E)^{O(e)} = \pi^\kappa (1 - E)^{\kappa - 1} (1 + O(e)),$$

which requires the limiting exponent $\kappa(1 + \theta_{\min}) - 1$ to be positive. For (iv), the exponent $2\kappa/\alpha - 1$ is increasing in κ , so the inequality holds for every $\kappa \geq \kappa_{\min}$ exactly when it holds at $\kappa_{\min}$, and $2\kappa_{\min}/\alpha - 1 \geq \theta_{\min}$ rearranges to $\alpha \leq 2\kappa_{\min}/(1 + \theta_{\min})$. The identification with $\kappa_{\min} \bar{\alpha}(1 + O(e))$ is (21). $\square$

Part (i) is Thorp's first lesson as an exact threshold: α -fractional Kelly keeps the growth rate positive against any overestimate of the edge up to a factor $2/\alpha$. Full Kelly tolerates an overestimate of at most 2; half Kelly, an overestimate of 4; and quarter Kelly, of 8; and so on.

Part (iii) quantifies the effect of overestimation on the information guarantee. Assume $\pi = 0.02$, $E = 50\%$, and $\bar{\alpha} \approx 0.30$. An edge overestimated by one-third ($\kappa = 0.75$) gives $\pi_{0.75} = 6.4\%$, and an edge overestimated by half ($\kappa = 0.50$) gives $\pi_{0.50} = 20.1\%$. In the latter case the stop loss carries 5-to-1 evidence rather than the 50-to-1 evidence desired. For $\kappa = 1/2$, (36) reduces to $\pi_{1/2} = \sqrt{\pi/(1 - E)} (1 + O(e))$.

Part (iv) gives the practical response to a possibly overestimated edge. The bettor should decide the largest factor by which the edge may plausibly be overestimated, and scale the prescription by its reciprocal. A bettor who wants the (E, π) guarantee to hold in the event

that the true edge is half its estimate should stake $\kappa_{\min}\bar{\alpha} = 0.5 \times 30.1\% \approx 15\%$ of the estimated Kelly fraction. Guarding against a one-quarter edge costs another halving (7.5% of Kelly). The drift of practitioner folklore from Thorp’s half-Kelly toward the quarter- and fifth-Kelly fractions reported in practice is what the robust rule predicts for a bettor who assumes his edge is overestimated. Moderate parameters ($\pi = 0.02$, $E = 50\%$) combined with moderate skepticism ($\kappa_{\min}$ between $\frac{1}{2}$ and $\frac{3}{4}$) suggest a Kelly fraction in the range 15–23% is prudent.

Remark 39 (General returns). The martingale arguments of Sections 3–4 do not depend on the binary structure of a wager. In general, let the per-unit return R of each round satisfy $R \geq -1$, $\mathbb{P}(R < 0) > 0$, and $\mathbb{E}R = e > 0$, and replace (10) by $\psi_f(\theta) := \mathbb{E}[(1 + fR)^{-\theta}]$. Then the Cramér root exists and is unique whenever $g(f) := \mathbb{E}\log(1 + fR) > 0$, and Theorems 7, 18, and 31, Proposition 11, and Lemma 10 generalize verbatim. At an interior growth optimum f^* the first-order condition $\mathbb{E}[R/(1 + f^*R)] = 0$ gives $\mathbb{E}[(1 + f^*R)^{-1}] = 1$, so $\theta^*(f^*) = 1$ and the reciprocal bankroll is a martingale under full Kelly staking for every such return distribution. This is the Kuhn–Tucker condition for the log-optimal portfolio [4, Theorem 16.2.2]. In the small-edge limit $f^* = e/m_2 + O(e^2)$ with $m_2 := \mathbb{E}R^2$ and $\theta^*(\alpha f^*) = 2/\alpha - 1 + O(e)$, so the closed form (21) is unchanged. Among all returns with a given mean and a given maximum payoff b , the binary wager is the worst case for both the growth rate and the drawdown probability, by the chord inequality applied to the convex map $x \mapsto (1 + fx)^{-\theta}$ and the concave map $x \mapsto \log(1 + fx)$ on $[-1, b]$. For Gaussian log-returns, and in the continuous-time model $dB_t = B_t f(\mu dt + \sigma dZ_t)$, there is no overshoot and $\theta^*(\alpha f^*) = 2/\alpha - 1$ exactly, so the diffusion formulas hold with equality.

6 Discussion

The exponent rule. The paper prescribes betting $2/(1 + \theta_{\min})$ of the Kelly fraction, where $\theta_{\min} = \log \pi / \log(1 - E)$ is the number of bits of evidence demanded per halving of committed capital. The rule is exact in the small-edge limit at every payoff (Theorem 18, Corollary 21) and exact without qualification for Gaussian log-returns and in continuous time (Remark 39). At even money and shorter odds it is a rigorous bound at every edge (Theorem 26, Corollary 27), and remains so for every return distribution with the same mean and maximum payoff (Remark 39). At longer odds it systematically overbets. This overbetting can be repaired via the skew-corrected fraction of Remark 28.

Relation to risk aversion. The standard defense of fractional Kelly appeals, implicitly or explicitly, to a utility function more concave than the logarithm [14]. The present account requires no such appeal, and therefore avoids objections that growth-rate maximization assumes a particular choice of utility function [17]. A bettor who is risk-neutral over terminal wealth but uncertain whether his edge is real will still bet a fraction of Kelly. The associated variance reduction is incidental to the primary motivation, which is to detect a mistaken edge before it ruins him. Indeed, under the expected-wealth objective of Theorem 18(i) the bettor is risk-neutral by construction, and the fractional sizing comes entirely from the stop-loss discipline. Corollary 20 makes this precise. Whenever $\pi \leq 1 - E$, which holds for every configuration in Tables 3 and 4, the risk-neutral bettor and the log-utility bettor are prescribed the same stake. Within the constrained framework, the choice of objective at issue in the Kelly–Samuelson debate does not affect the prescription. The stop loss is a budget and the premature stopping probability is a tolerance for being wrong. Neither expresses a preference for a utility function.

Relation to Thorp. Both of Thorp’s lessons [21, Section 7.3] follow from the results above. His first lesson is to assume the true edge is less than the estimate and to bet small enough so that growth remains positive despite this estimation error. His second lesson is the prescription of the half-Kelly heuristic. These lessons correspond to the above as follows: any premature stopping requirement $\pi < 1$ confines the stake to the positive-growth region (Remark 23); and half-Kelly is the prescription when the stop loss and premature stopping probability lie on the cube-law frontier $\pi = (1 - E)^3$ of Section 4.2, for example $\pi = \frac{1}{8}$ at $E = \frac{1}{2}$.

Relation to drawdown control. Constraining the probability of a drawdown has a long history in growth-optimal investment. MacLean et al. [13] trade off growth against the probability of falling to a fixed wealth level, and the secure-growth literature it spawned obtains fractional-Kelly strategies as solutions; Grossman and Zhou [10] solve the continuous-time growth problem under an almost-sure constraint on drawdown from the running maximum. Closest to the present setting, Busseti et al. [3] impose the drawdown-probability bound $\mathbb{P}(\inf_n B_n/B_0 < 1 - E) \leq \pi$, derived from the same exponential supermartingale, as a convex constraint and optimize the stake numerically over portfolios. The present paper differs in two respects. First, the constraint is solved in closed form. The frontier, the exponent rule, the supermartingale families, and the robustness theorem are all explicit. Second, the constraint is given an evidential interpretation. The quantity of interest is the likelihood ratio that a drawdown induces, and (9) shows that this likelihood ratio is the reciprocal of the drawdown probability because ruin is certain under the null. For this reason the premature stopping probability is both a Type I error rate and a weight of evidence.

Adverse selection and the suspicious bettor. Two empirical regularities motivate the framework. First, apparently advantageous bets are more often traps than opportunities, so the prior probability that the maintained hypothesis $\mathbb{P}_e$ is correct should be modest, and the bettor should want cheap evidence against it. Second, a real edge is usually smaller than its estimate, because the estimate is a best-case account of the situation and adverse selection ensures the omissions run one way. Section 5 quantifies this asymmetry and supplies the correction. Fractional Kelly is the response of a suspicious bettor. He stakes enough to profit if the edge is real and little enough that he can afford to learn that it is not.

Limitations and extensions. We have assumed i.i.d. wagers with a single fixed edge. Remark 39 removes the binary-outcome restriction and Section 5 relaxes the assumption that the magnitude of the edge is correctly estimated, but realistic settings also involve heterogeneous bets, correlated outcomes, and simultaneous positions. Each merits separate treatment and is not considered here. The framework extends to these settings in principle. Whatever the dynamics, one can compute the probability that a given stop loss is hit under the maintained hypothesis and size bets so that hitting it is informative. We also treated the null as either the single fair alternative $e = 0$ or the composite alternative $e \leq 0$. Theorem 38 handles a composite alternative over $\kappa \geq \kappa_{\min}$ by least favorability. We have also restricted attention to fixed-fraction strategies. In practice, an adaptive bettor could stake next to nothing until his bankroll reaches a high level L and switch to full Kelly thereafter. Since the premature stopping probability from L is $((1 - E)B_0/L)^{\theta_1^*}$, which tends to 0 as $L \rightarrow \infty$, such two-phase play satisfies any (E, π) constraint while attaining the full-Kelly asymptotic growth rate. The constraint governs the early bankroll path, which asymptotic growth ignores. The fixed-fraction class is nevertheless the

right one here, because the fixed fraction is the strategy under test and the framework concerns the evidential meaning of its early failure. A natural extension is to couple the framework to sequential estimation of the edge, so that the stake adapts as the outcomes of the bets refine the estimate.

References

- [1] Søren Asmussen and Hansjörg Albrecher. *Ruin Probabilities*. World Scientific, Singapore, 2nd edition, 2010.
- [2] Leo Breiman. Optimal gambling systems for favorable games. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 65–78. University of California Press, 1961.
- [3] Enzo Busseti, Ernest K. Ryu, and Stephen Boyd. Risk-constrained Kelly gambling. *Journal of Investing*, 25(3):118–134, 2016.
- [4] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, Hoboken, NJ, 2 edition, 2006.
- [5] Harald Cramér. On the mathematical theory of risk. In *Försäkringsaktiebolaget Skandia 1855–1930, Jubilee Volume*. Skandia, Stockholm, 1930.
- [6] Lester E. Dubins and Leonard J. Savage. *How to Gamble If You Must: Inequalities for Stochastic Processes*. McGraw-Hill, New York, 1965.
- [7] Richard A. Epstein. *The Theory of Gambling and Statistical Logic*. Academic Press, 2 edition, 2009.
- [8] Stewart N. Ethier. *The Doctrine of Chances: Probabilistic Aspects of Gambling*. Springer, 2010.
- [9] James Grosjean. *Exhibit CAA: Beyond Counting. Exploiting Casino Games from Blackjack to Video Poker*. South Side Advantage Press, Las Vegas, 2009.
- [10] Sanford J. Grossman and Zhongquan Zhou. Optimal investment strategies for controlling drawdowns. *Mathematical Finance*, 3(3):241–276, 1993.
- [11] Eliot Jacobson. *Advanced Advantage Play: Beating and Safeguarding Modern Casino Table Games, Side Bets and Promotions*. Blue Point Books, Santa Barbara, CA, 2015.
- [12] J. L. Kelly. A new interpretation of information rate. *Bell System Technical Journal*, 35(4): 917–926, 1956.
- [13] Leonard C. MacLean, William T. Ziemba, and George Blazenko. Growth versus security in dynamic investment analysis. *Management Science*, 38(11):1562–1585, 1992.
- [14] Leonard C. MacLean, Edward O. Thorp, and William T. Ziemba, editors. *The Kelly Capital Growth Investment Criterion: Theory and Practice*. World Scientific, 2011.
- [15] Ole Peters. The ergodicity problem in economics. *Nature Physics*, 15:1216–1221, 2019.

- [16] Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 38(4):576–601, 2023.
- [17] Paul A. Samuelson. The fallacy of maximizing the geometric mean in long sequences of investing or gambling. *Proceedings of the National Academy of Sciences*, 68(10):2493–2496, 1971.
- [18] Glenn Shafer. Testing by betting: A strategy for statistical and scientific communication. *Journal of the Royal Statistical Society, Series A*, 184(2):407–431, 2021.
- [19] Glenn Shafer and Vladimir Vovk. *Game-Theoretic Foundations for Probability and Finance*. Wiley, Hoboken, 2019.
- [20] Edward O. Thorp. *Beat the Dealer: A Winning Strategy for the Game of Twenty-One*. Vintage, 1966.
- [21] Edward O. Thorp. The Kelly criterion in blackjack, sports betting, and the stock market. In *Handbook of Asset and Liability Management*, volume 1, pages 385–428. Elsevier, 2006.
- [22] Jean Ville. *Étude Critique de la Notion de Collectif*. Gauthier-Villars, Paris, 1939.