Betting with Prior on Horizon-Aware Anytime-Valid Testing

Ojas Kalra

1 Literature Review

This literature review goes over use of coin betting and connections with online learning, concentration inequalities, and statistical hypothesis testing. The first connection is with online learning where Orabona and Pál showed regret minimization connection with wealth maximization game and showed how to get parameter free OCO [Orabona and Pál, 2016]. The second connection is with concentration inequalities, where regret bounds from a betting algorithm translate into tail bounds on the sample mean, yielding tight confidence sequences via Ville's inequality applied to the wealth process [Jun and Orabona, 2019, Orabona and Jun, 2023, Ville, 1939, Waudby-Smith and Ramdas, 2024]. Finally, more recent work has been doing hypothesis testing and connection with coin betting under a known deadline N where we have N flips to try and reject null hypothesis [Taga et al., 2026, Voráček and Orabona, 2025].

We organize the review into three parts: (i) the coin betting framework and parameter free online learning (ii) concentration inequalities (iii) horizon aware statistical inference.

1.1 Coin-Betting and Parameter-Free Online Learning

Starting with coin betting and OCO, Orabona and Pál were able to show an equivalent coin betting game [Orabona and Pál, 2016]. In this game, they describe it as a bettor sees a sequence of coins $c_t \in [-1, 1]$ and they will wager fraction $\beta_t \in [-1, 1]$ of their current wealth on the next coin toss. The connection they found was that a lower bound on the bettor's wealth is a regret upper bound on the original OCO problem where we get $O(\|u\|_2\sqrt{T \log(\|u\|_2 T)})$ regret but we don't need to tune any parameters. This is significant because prior bounds required variables like the learning rate (which depends on the unknown competitor norm $\|u\|_2$) to be "tuned".

1.2 Betting-Based Concentration and Confidence Sequences

Alongside the work of coin betting and parameter free online learning, Orabona and Jun also showed a bridge between coin betting and statistical inference [Jun and Orabona, 2019, Orabona and Jun, 2023]. This comes from the typical wealth process defined as

$$W_n(m) = \prod_{t=1}^{n} \big(1 + \lambda_t (X_t - m)\big),$$

where $\{\lambda_t\}$ is a predictable sequence in the safe range $[-1/(1-m),\ 1/m]$, under the null $\mathbb{E}[X_t] = m$ which makes $W_n(m)$ a nonnegative martingale starting at 1. Additionally, Ville's inequality [Ville, 1939] states that $\mathbb{P}(\sup_{n\geq 1} W_n \geq 1/\alpha) \leq \alpha$ and that we can reject if wealth crosses $1/\alpha$. This gives us a framework for doing statistical inference.

1.3 Horizon-Aware Methods

The most recent work related to coin betting and inference is related to betting on finite horizon settings. Voráček and Orabona [2025] proposed STaR-Bets which dynamically updates bets at each step based on the log distance to the rejection threshold. In other words, they proposed an algorithm to dynamically adjust betting sizes based on how far we are away from rejecting the null hypothesis. They showed that the resulting confidence sequence width at the deadline is within $(1+o(1))$ of the optimal fixed-sample confidence interval. Taga et al. [2026] did something similar but looked at the finite horizon optimal control setting where we consider the state $(t, \log W_t)$. They showed that depending on whether we are betting ahead, on schedule, or behind schedule, the optimal betting strategy sometimes deviates away from Kelly. If we are ahead of schedule, we can relax, but if we are behind, we would have to bet more aggressively. They trained a Deep Q-Network to estimate the optimal policy over unknown distribution P_X since it is intractable in closed form.

1.4 History and Motivation

In 1945, Wald set up a test to determine which of two distributions data comes from [Wald, 1945]. We will say that under H_0 the data comes from P_0 and under H_1 the data comes from P_1. The likelihood function is defined as

$$\Lambda_n = \frac{\prod_{i=1}^n p_1(X_i)}{\prod_{i=1}^n p_0(X_i)},$$

where p_0, p_1 are the densities of P_0, P_1, and this test will be called Sequential Probability Ratio Test (SPRT). We use two thresholds $A > B$: if $\Lambda_n \geq A$ reject H_0, if $\Lambda_n \leq B$ we accept H_0, otherwise keep sampling. In this scenario, there is no deadline but it terminates almost surely.

In 1960, Anderson took this problem and asked what happens if data comes from neither P_0 nor P_1 [Anderson, 1960]. In this scenario, the likelihood ratio would drift around and would possibly not terminate. Anderson decided to pick a deadline N in advance where we would have to force a decision to be made. Anderson was able to show that this test's Type I and Type II error rates were similar to the original SPRT. This was one of the earliest works in horizon-aware testing.

Then in 1968, Darling and Robbins did a nonparametric setup [Darling and Robbins, 1968]. We no longer know where the data comes from and try to test a functional of the distribution (i.e., mean, variance, or median). This test also had power one, i.e., if the alternative is true you will eventually reject with probability 1. By the law of the iterated logarithm, we know that the sample mean fluctuates around the true mean by an order of $\sqrt{(\log\log n)/n}$, and they showed a boundary like $\bar{X}_n \pm c\sqrt{(\log\log n)/n}$, where c is a function of α.

In more recent work, Waudby-Smith and Ramdas in 2024 operationalized the betting and wealth framework connection for testing means [Waudby-Smith and Ramdas, 2024]. They define

$$W_n(m) = \prod_{i=1}^{n} \left(1 + \lambda_i (X_i - m)\right),$$

with predictable bets $\lambda_i \in [-1/(1-m), 1/m]$, which is a nonnegative martingale under $H_0 : \mathbb{E}[X] = m$. This is equivalent to hypothesis testing where we reject W_n if it crosses $1/\alpha$.

	Parametric	**Nonparametric (bounded mean)**
Infinite horizon	Wald (1945)	Darling & Robbins (1968); WSR (2024)
Fixed deterministic N	Anderson (1960)	Voráček & Orabona (2025); Taga et al. (2026)
Random horizon $N \sim \pi$	—	**This report**

Table 1: Sequential testing landscape

1.5 The Taga et al. Setup in Detail

We will now give a detailed summary of Taga et al. [2026] as our work will build directly off this paper. Define $X_i \overset{\text{i.i.d.}}{\sim} P_X$ on $[0,1]$ with mean μ_X, and we test $H_0 : \mu_X = m$ versus $H_1 : \mu_X \neq m$ for known $m \in (0,1)$. The log-wealth increment is

$$h_m(\lambda, x) = \log\bigl(1 + \lambda(x - m)\bigr), \qquad \lambda \in \Lambda_m := \left[-\tfrac{1}{1-m}, \tfrac{1}{m}\right],$$

and the cumulative log-wealth is $Y_t = \sum_{i=1}^{t} h_m(\lambda_i, X_i)$. We denote Kelly's bet as λ_m^{Kelly}, which maximizes the expected log-wealth $\mathbb{E}_X[h_m(\lambda, X)]$. We also define $b = \log(1/\alpha)$, $B = \sup_x |h_m(\lambda_m^{\text{Kelly}}, x)|$, $L_{\max} = \mathbb{E}[h_m(\lambda_m^{\text{Kelly}}, X)]$, and the rejection time $\tau = \inf\{t \geq 1 : Y_t \geq b\}$. The main result we will build off of is their Theorem 3.1, which states:

> Consider any time t, and log-wealth $y = \log W_t$. Let $b = \log(1/\alpha)$ denote the threshold, and $T = N - t$ the remaining time. Fix a $\delta > 0$, and assume that there exists an $\epsilon \equiv \epsilon(\delta) > 0$, such that $|\lambda - \lambda_m^{\text{Kelly}}| \geq \delta$ implies $L(\lambda) \leq L_{\max} - \epsilon$. Let Δ denote the term $TL_{\max} - (b - y)$.
>
> - Suppose we follow the policy that sets $\lambda_i = \lambda_m^{\text{Kelly}}$ for all $i \in \{t+1, \ldots, N\}$. Then, we have
> $$\Delta \geq B\sqrt{8T \log T} \implies \mathbb{P}(\tau \leq N) \geq 1 - \tfrac{1}{T}.$$
> - Now, let us consider any other policy that plays $\{\lambda_{t+1}, \ldots, \lambda_N\}$ taking values in Λ such that for some $\rho > 0$, the condition
> $$\bigl|\{i \in \{t+1, \ldots, t+k\} : |\lambda_i - \lambda_m^{\text{Kelly}}| \geq \delta\}\bigr| \geq \rho k$$
> holds for all $k \in \{1, 2, \ldots, T\}$. Then, we have
> $$\Delta \leq \rho\epsilon T - B\sqrt{8T \log 2} \implies \mathbb{P}(\tau \leq N) \leq \tfrac{1}{2}.$$

We will then consider extensions to this theorem and see what happens if we have a prior over N instead of a hard deadline.

2 Extensions to Prior on Horizon-Aware Valid Testing

In Taga et al. [2026] work, they talk about three cases you can be in during the betting game. Firstly is when you are "on schedule" to reject null hypothesis in which case using Kelly is near optimal strategy (Theorem 3.1). If you are behind schedule, then you would need to be more aggressively (Proposition 3.4) and if you are ahead of schedule, you can be more passive (Proposition 3.6). We will look at each of these cases and see what happens if instead of using fixed N, we consider a prior π over N.

2.1 Random-Horizon Kelly Bound: General Prior Case

Setup

Let X_i be i.i.d. in $[0,1]$ drawn from distribution P_X with mean $\mu_X \neq m$ (i.e., alternative hypothesis is true) and let $\{\mathcal{F}_n\}_{n\geq 0}$ be the natural filtration with $\mathcal{F}_n = \sigma(X_1, X_2, \ldots, X_n)$. For fixed null value $m \in (0,1)$, we define log-wealth increment per step as

$$h_m(\lambda, x) = \log\big(1 + \lambda(x-m)\big), \qquad \lambda \in \Lambda_m := \left[-\tfrac{1}{1-m}, \tfrac{1}{m}\right].$$

Then we define $\lambda_m^{\text{Kelly}} \in \Lambda_m$ to denote the Kelly bet where

$$\lambda_m^{\text{Kelly}} = \arg\max_{\lambda\in\Lambda_m} \mathbb{E}_{X\sim P_X}\big[h_m(\lambda, X)\big].$$

Then $L_{\max} := \mathbb{E}[h_m(\lambda_m^{\text{Kelly}}, X)]$ and $B := \sup_{x\in[0,1]} |h_m(\lambda_m^{\text{Kelly}}, x)|$. We define $b = \log(1/\alpha)$, rejection time as $\tau := \inf\{t \geq 1 : Y_t \geq b\}$, log-wealth process $Y_t = \sum_{i=1}^t h_m(\lambda_m^{\text{Kelly}}, X_i)$ and $Y_0 = 0$. Then since $\mu_X \neq m$ and $L_{\max} > 0$, the strong law of large numbers gives $\tau < \infty$ almost surely. Wald's identity gives $\mathbb{E}[\tau] \leq (b+B)/L_{\max} < \infty$.

Let N be a random horizon taking values in $\mathbb{N}$ and be independent of $\{X_i\}$ distributed with a known prior π. We define $S(k)$ to be the survival function with $S(k) > 0$ for all $k \geq 1$ and $S(k) := \pi(N \geq k)$. Assume S is log-convex, then we get hazard rate $r(k) := 1 - S(k+1)/S(k)$ which is non-increasing in k.

Main Result

Theorem 1 (Prior on Horizon Kelly bound)**.** *Using the setup above, the rejection time τ of Kelly betting against an independent random horizon $N \sim \pi$ satisfies*

$$\mathbb{P}(\tau \leq N) \;\geq\; S\big(\mathbb{E}[\tau]\big),$$

where $S(k) = \mathbb{P}(N \geq k)$. Assume that the survival sequence $\{S(k)\}_{k\geq 0}$ is log convex and extend S to $\mathbb{R}_{\geq 0}$ by log linear interpolation: i.e., $x \in [k, k+1]$,

$$\log S(x) = (k+1-x)\log S(k) + (x-k)\log S(k+1).$$

Then, $\log S$ is convex so $S = \exp(\log S)$ is also convex.

In particular, using the Wald bound $\mathbb{E}[\tau] \leq (b+B)/L_{\max}$, we get the closed form lower bound

$$\mathbb{P}(\tau \leq N) \;\geq\; S\left(\frac{b+B}{L_{\max}}\right). \tag{1}$$

Proof. Since N is independent of the data stream $\{X_i\}$, and hence of τ, conditioning on τ gives

$$\mathbb{P}(\tau \leq N) = \mathbb{E}_\tau\big[\mathbb{P}(N \geq \tau \mid \tau)\big] = \mathbb{E}_\tau\big[S(\tau)\big],$$

where we use that $\tau < \infty$ almost surely. By the log convexity of S, we have that the function $S = \exp\circ\log S$ is the exponential of a convex function so it is also convex itself. Applying Jensen's inequality:

$$\mathbb{E}_\tau[S(\tau)] \;\geq\; S\big(\mathbb{E}[\tau]\big).$$

For the closed form bound (1), note that S is non-increasing, so any upper bound on $\mathbb{E}[\tau]$ yields a lower bound on $S(\mathbb{E}[\tau])$. Using Wald's identity applied to the martingale $Y_t - tL_{\max}$ at the stopping time τ gives

$$\mathbb{E}[Y_\tau] = L_{\max} \cdot \mathbb{E}[\tau].$$

Since $Y_\tau \in [b, b+B]$ almost surely, we have $\mathbb{E}[Y_\tau] \leq b+B$, giving $\mathbb{E}[\tau] \leq (b+B)/L_{\max}$. ☐

Example with Geometric Prior

When $N \sim \text{Geometric}(q)$ with $\mathbb{E}[N] = 1/q$, the survival function $S(k) = (1-q)^{k-1}$ is log linear and log convex. Theorem 1 then gives the closed form bound

$$\mathbb{P}(\tau \le N) \;\ge\; (1-q)^{(b+B)/L_{\max}-1} \;\approx\; \exp\!\left(-\frac{b+B}{\mathbb{E}[N]\cdot L_{\max}}\right),$$

The interpretation of this is that the rejection probability decreases exponentially in the ratio of the expected hitting time $(b+B)/L_{\max}$ to the expected horizon $\mathbb{E}[N]$.

2.2 Translations to Random Variable N

Theorem 1 talks about what happens when we are on schedule and that using Kelly's is good enough. We now look at what happens when we are no longer on schedule, specifically extensions of Propositions 3.4 and 3.6 of Taga et al. [2026]. Consider

$$r(t, y, b) = \frac{b-y}{N-t},$$

which represents the per step log wealth growth needed to cross the threshold by the deadline. The off schedule conditions of Propositions 3.4 and 3.6 compare this drift to $L_{\max}$.

Under a random horizon, $N - t$ is itself random, so r no longer makes sense as written. We can however have an analogous replacement conditions on survival to time t and uses the expected residual life

$$\mathrm{ERL}(t) := \mathbb{E}[N-t \mid N \ge t] = \sum_{k \ge t} \frac{S(k)}{S(t)} - 1.$$

We define the random horizon required drift as

$$r_\pi(t, y, b) := \frac{b-y}{\mathrm{ERL}(t)}.$$

In the fixed horizon proof from Taga, they use $T/2$ as a floor on hitting time inside concentration arguments. In our context, $T/2$ doesn't make sense because the deadline is random now. We will instead use

$$T_0(y) := \left\lceil \frac{b-y}{B_K} \right\rceil,$$

this represents the minimum number of steps required to cross the threshold from log-wealth y since each step is at most B_K.

Similarly, a list of the translations and difference from fixed N to $N \sim \pi$ looks as follows:

Quantity	**Fixed N**	**Random $N \sim \pi$**
Remaining time	$T = N - t$	$\mathrm{ERL}(t) = \mathbb{E}[N-t \mid N \ge t]$
Required drift	$r = (b-y)/T$	$r_\pi = (b-y)/\mathrm{ERL}(t)$
Off-schedule gap	$\Delta = TL_{\max} - (b-y)$	$\Delta_\pi = \mathrm{ERL}(t)\cdot L_{\max} - (b-y)$
Min. hitting floor	$T/2$	$T_0(y) = \lceil (b-y)/B_K \rceil$
Mirror drift	$r_- = 2(r - \frac{1}{2}B_K)$	$r_\pi^- = 2(r_\pi - \frac{1}{2}B_K)$
Schedule (behind)	$r > \max\{L_{\max}, \frac{1}{2}B_K\}$	$r_\pi > \max\{L_{\max}, \frac{1}{2}B_K\}$
Schedule (ahead)	$\frac{1}{2}B_K < r < L_{\max}$	$\frac{1}{2}B_K < r_\pi < L_{\max}$
Rate functions $I^\pm$	$I^\pm(\lambda, r)$	$I^\pm(\lambda, r_\pi)$

Table 2: Proposed translations from fixed N to prior over N

2.3 Extension of Proposition 3.4: Behind Schedule Regime

We now state the random horizon analog of Taga et al. [2026]'s behind schedule result.

Conjecture 1 (Random horizon analog of Proposition 3.4)**.** *Fix a state (t, y) with $N \geq t$, and suppose*

$$r_\pi(t, y, b) > \max\left\{L_{\max},\ \tfrac{1}{2}B_K\right\},$$

where $B_K = \max_x h_m(\lambda_m^{\text{Kelly}}, x)$. Let $\tau(\lambda)$ denote the stopping time for a constant policy λ played from (t, y) onward. Then for any $\lambda^{\text{agg}} > \lambda_m^{\text{Kelly}}$ satisfying

$$I^+(\lambda^{\text{agg}}, r_\pi) < \tfrac{1}{2}I^+(\lambda_m^{\text{Kelly}}, r_\pi) - c^+_{\text{ERL}(t)},$$

we have

$$\mathbb{P}\big(\tau(\lambda^{\text{agg}}) \leq N \mid N \geq t\big) \ \geq\ \mathbb{P}\big(\tau(\lambda_m^{\text{Kelly}}) \leq N \mid N \geq t\big).$$

Proof sketch. Conditioning on N and using independence of N and the data stream:

$$\mathbb{P}(\tau(\lambda) \leq N \mid N \geq t) = \sum_{k\geq 1} \frac{S(t+k-1) - S(t+k)}{S(t)} \cdot \mathbb{P}(\tau(\lambda) \leq t+k).$$

For each fixed remaining time $T = k$, Proposition 3.4 of Taga et al. [2026] provides a Sanov type ordering between $\mathbb{P}(\tau(\lambda^{\text{agg}}) \leq t+T)$ and $\mathbb{P}(\tau(\lambda_m^{\text{Kelly}}) \leq t+T)$, provided the drift condition holds at remaining time T. Leaving this for future work.

2.4 Extension of Proposition 3.6: Ahead of Schedule Regime

Conjecture 2 (Random horizon analog of Proposition 3.6)**.** *Fix a state (t, y) with $N \geq t$, and suppose*

$$\tfrac{1}{2}B_K < r_\pi(t, y, b) < L_{\max}.$$

Let $\tau(\lambda)$ denote the stopping time for a constant policy λ played from (t, y) onward. Then for any feasible $\lambda^{\text{def}} < \lambda_m^{\text{Kelly}}$ satisfying

$$I^-(\lambda^{\text{def}}, r_\pi) > \tfrac{1}{2}I^-(\lambda_m^{\text{Kelly}}, r_\pi^-) + c^-_{\text{ERL}(t)},$$

where $r_\pi^- := 2(r_\pi - \frac{1}{2}B_K)$, we have

$$\mathbb{P}\big(\tau(\lambda^{\text{def}}) \leq N \mid N \geq t\big) \ \geq\ \mathbb{P}\big(\tau(\lambda_m^{\text{Kelly}}) \leq N \mid N \geq t\big).$$

3 Experiments

We will validate the results over simulations. Specifically questions:

1. Across priors with the same mean horizon, what feature most strongly predicts problem difficulty? Is betting on certain distributions much harder and if so why?

2. How does Kelly betting compare to the DP optimal betting under random horizons?

3.1 Setup

Using similar Bernoulli setup of Taga et al. [2026, Example 3.3]: $X_i \overset{\text{i.i.d.}}{\sim}$ Bernoulli(0.7), null mean $m = 0.5$, $\alpha = 0.05$, threshold $b \approx 2.996$. The Kelly bet is $\lambda_m^{\text{Kelly}} = 0.8$, drift $L_{\max} \approx 0.0823$, maximum jump $B \approx 0.5108$, and Wald upper bound $\mathbb{E}[\tau] \leq (b + B)/L_{\max} \approx 42.6$.

We will make $\mathbb{E}[N] = 100$ across all priors and compare: Point mass at 100; Uniform $[70, 130]$ and Uniform $[10, 190]$; Bimodal $0.7\,\delta_{50}+0.3\,\delta_{217}$; Geometric(1/100); an equal mixture of Geometric(1/40) and Geometric(1/160); and a discrete Pareto with $\pi(N = k) \propto (k + k_0)^{-2.5}$.

Also note that Wald upper bound is $\mathbb{E}[\tau] \leq (b + B)/L_{\max} \approx 42.6$. However from simulation, $\mathbb{E}[\tau] \approx 36$ via Monte Carlo, slightly tighter than the Wald bound so we use this value in our experiments - specifically we care about $\pi(N \leq 36)$.

3.2 Methodology

For each prior, we compute the empirical Kelly rejection rate via Monte Carlo over 30,000 trials, the DP-optimal value $V^\star(\pi) = \max_{\text{policy}} \mathbb{P}(\tau \leq N)$ via backward induction

$$V_t(y) = \max_{\lambda \in \mathcal{A}} \mathbb{E}_X\Big[\mathbf{1}\{y + h_m(\lambda, X) \geq b\} + \mathbf{1}\{y + h_m(\lambda, X) < b\} \cdot p_t \cdot V_{t+1}\big(y + h_m(\lambda, X)\big)\Big],$$

with $p_t = \pi(N \geq t + 2)/\pi(N \geq t + 1)$ and action set $\mathcal{A} = \{\frac{1}{2}\lambda_m^{\text{Kelly}}, \lambda_m^{\text{Kelly}}, \lambda_{\max}\}$. This is similar set up to Taga et al. [2026].

3.3 Main Results

Table 3: Performance under seven priors with $\mathbb{E}[N] = 100$. Comparing using "$P(N \leq 36)$" which is the prior mass below the expected Kelly hitting time. "Bound" is $S(\mathbb{E}[\tau]_{\text{Wald}})$. Monte Carlo estimates use 30,000 trials.

Prior	std(N)	$P(N \leq 36)$	Kelly	$V^\star$
Point mass δ_{100}	0.00	0.000	0.959	0.9892
Uniform $[70, 130]$	17.60	0.000	0.951	0.9737
Uniform $[10, 190]$	52.26	0.149	0.843	0.8486
Bimodal $(50, 217;\ 0.7/0.3)$	76.54	0.000	0.847	0.8877
Geometric(1/100)	99.36	0.304	0.709	0.7172
MixGeom(40, 160)	130.84	0.400	0.635	0.6395
DiscPareto($\alpha = 2.5$)	186.10	0.476	0.564	0.5742

3.4 Findings

Lower tail mass is a near sufficient predictor Across all seven priors, $\pi(N \leq 36)$ predicts $1 - V^\star$ with Pearson $\rho = 0.98$. The hazard bound depends on π only through its survival at $\mathbb{E}[\tau]$, which for log-convex distributions is dominated by lower tail mass so simulation matches conjugectures. Essentially, two priors with same mean can have vastly different rejection rates: i.e., U(70,130) has ($P(N \leq 36) = 0$) with $V^\star \approx 0.99$, while DiscPareto ($P(N \leq 36) = 0.48$) has $V^\star \approx 0.57$.

Kelly is still near optimal under random N. The DP gap ranges from 0.4 pp (Geometric) to 4.1 pp (Bimodal). These results indicate that Kelly is very close to optimal. These results are aligned with the fixed N that showed that Kelly was near optimal.

3.5 Limitations

The main limitation of our experiment is that we fix $P_X = \text{Bernoulli}(0.7)$. However, the problem comes down do the behavior of both P_x and π, P_x affects our $\mathbb{E}[\tau]$. An experiment over combinations of (P_x, π) is left as future work.

References

T. M. Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, 1991.

K.-S. Jun and F. Orabona. Parameter-free online convex optimization with sub-exponential noise. In *Conference on Learning Theory*, pages 1802–1823. PMLR, 2019.

F. Orabona and K.-S. Jun. Tight concentrations and confidence sequences from the regret of universal portfolio. *IEEE Transactions on Information Theory*, 70(1):436–455, 2023.

F. Orabona and D. Pál. Coin betting and parameter-free online learning. In *Advances in Neural Information Processing Systems*, pages 577–585, 2016.

E. O. Taga, S. Oymak, and S. Shekhar. Learning to bet for horizon-aware anytime-valid testing. *arXiv preprint arXiv:2603.19551*, 2026.

J. Ville. *Étude critique de la notion de collectif.* Gauthier-Villars, Paris, 1939.

V. Voráček and F. Orabona. STaR-Bets: Sequential target-recalculating bets for tighter confidence intervals. *arXiv preprint arXiv:2505.22422*, 2025.

I. Waudby-Smith and A. Ramdas. Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society Series B*, 86(1):1–27, 2024.

T. W. Anderson. A modification of the sequential probability ratio test to reduce the sample size. *The Annals of Mathematical Statistics*, 31(1):165–197, 1960.

D. A. Darling and H. Robbins. Some nonparametric sequential tests with power one. *Proceedings of the National Academy of Sciences*, 61(3):804–809, 1968.

G. Shafer and V. Vovk. *Game-Theoretic Foundations for Probability and Finance.* John Wiley & Sons, 2019.

A. Wald. Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16(2):117–186, 1945.