STA 711: Probability \& Measure Theory<br>Robert L. Wolpert

## 12 Martingale Methods: Application to SPRT

## Random Walks and Martingales

Let $\left\{\xi_{j}\right\}$ be independent, identically distributed random variables, all with the same mean $\mu=\mathrm{E}\left[\xi_{j}\right]$, variance $\sigma^{2}=\mathrm{V}\left[\xi_{j}\right]$, and moment generating function $M(\lambda)=\mathrm{E}\left[e^{\lambda \xi_{j}}\right]$ finite in an interval around $\lambda=0$. Under suitable regularity conditions the $\log \operatorname{arithm} m(\lambda):=\log M(\lambda)$ has Taylor expansion $m(\lambda)=\mu \lambda+\sigma^{2} \lambda^{2} / 2+o\left(\lambda^{2}\right)$ near zero. Let $\mathcal{F}_{n}:=\sigma\left\{\xi_{j}: j \leq n\right\}$ be the filtration generated by $\left\{\xi_{j}\right\}$. For any $x \in \mathbb{R}$ consider the sequence $X_{n}=x+\sum_{j \leq n} \xi_{j}$ of partial sums, starting at $x ; X_{n}$ is a random walk starting at $x$. Fix real numbers $a<b$ and define a $\left\{\mathcal{F}_{n}\right\}$-stopping time $\tau=\tau_{a, b}$ by

$$
\tau:=\inf \left\{n: X_{n} \notin(a, b)\right\}
$$

and the "right exit probability" by $\alpha:=\mathrm{P}\left[\tau<\infty\right.$ and $\left.X_{\tau} \geq b\right]$. Our object is to compute $\alpha$ and $\mathrm{E}[\tau]$, the probability of exiting on the right and the expected exit time, as functions of $x \in(a, b)$. This classical probability problem is called the "gambler's ruin problem", for the example in which $x$ is a gambler's initial fortune and $\xi_{j}$ the net return on the $j$ th independent play of some game of chance; in that case $\alpha$ is the probability that the gambler attains a fortune of $X_{n} \geq b$ before going broke (for $a=0$ ). Below we will present an approach to sequential statistical significance testing based on the same probability model.

## The Symmetric Case

First suppose $\mu=0$. Then $X_{n}$ is a martingale, and so (by the optional sampling theorem) is $X_{\tau \wedge n}$, which moreover is bounded and hence uniformly integrable. It follows that

$$
\begin{aligned}
x & =\mathrm{E}\left[X_{\tau \wedge 0}\right] \\
& =\lim _{n \rightarrow \infty} \mathrm{E}\left[X_{\tau \wedge n}\right] \\
& =\mathrm{E}\left[X_{\tau}\right] \\
& \approx a(1-\alpha)+b(\alpha) ; \quad \text { solving, we find } \\
\alpha & \approx \frac{x-a}{b-a}
\end{aligned}
$$

(the estimates are exact if $\mathrm{P}\left[X_{\tau} \in\{a, b\}\right]=1$, but only approximate if there is a chance of "overshooting" the boundary). Also $\left(X_{n}\right)^{2}-n \sigma^{2}$ is a martingale, so

$$
\begin{aligned}
x^{2} & =\mathrm{E}\left[\left(X_{\tau \wedge 0}\right)^{2}-(\tau \wedge 0) \sigma^{2}\right] \\
& =\mathrm{E}\left[\left(X_{\tau}\right)^{2}-\tau \sigma^{2}\right] \\
& \approx a^{2}(1-\alpha)+b^{2}(\alpha)-\mathrm{E}[\tau] \sigma^{2}, \quad \text { so } \\
\mathrm{E}[\tau] & \approx \frac{a^{2}+\alpha\left(b^{2}-a^{2}\right)-x^{2}}{\sigma^{2}} \\
& =(b-x)(x-a) / \sigma^{2} .
\end{aligned}
$$

For example, for the standard symmetric random walk on the integers with $\xi= \pm 1$ with probability $1 / 2$ each, then $\mu=0, \sigma^{2}=1$, and for $a<x<b, \alpha=\mathrm{P}\left[X_{n} \geq b\right.$ before $\left.X_{n} \leq a \mid X_{0}=x\right]=(x-$ $a) /(b-a)$ and the exit time $\tau:=\min \left[n: X_{n} \notin(a, b)\right]$ has expectation $\mathrm{E}\left[\tau \mid X_{0}=x\right]=(b-x)(x-a)$.

## The Asymmetric Case

Now suppose $\mu \neq 0$, that $\mathrm{P}\left[\xi_{j}<0\right]>0$ and $\mathrm{P}\left[\xi_{j}>0\right]>0$, and that $M(t)$ is smooth enough for the logarithm $m(\lambda):=\log M(\lambda)$ to have Taylor expansion $m(\lambda)=\mu \lambda+\sigma^{2} \lambda^{2} / 2+o\left(\lambda^{2}\right)$ near zero. Then $m(\lambda) \rightarrow \infty$ as $\lambda \rightarrow \pm \infty$ while $m(0)=0$ and $m^{\prime}(0)=\mu \neq 0$, so there exists some $\lambda^{*} \neq 0$ (approximately $\lambda^{*} \approx-2 \mu / \sigma^{2}$ ) for which $m\left(\lambda^{*}\right)=0$. For any $\lambda \in \mathbb{R}, Y_{n}:=e^{\lambda X_{n}-n m(\lambda)}$ is a martingale (well, any $\lambda$ for which $M(\lambda)<\infty$ ). In particular, $e^{\lambda^{*} X_{n}}$ is a martingale, so again the optional sampling theorem and UI limit theorem give

$$
\begin{aligned}
e^{\lambda^{*} x} & =\mathrm{E}\left[e^{\lambda^{*} X_{\tau \wedge n}}\right] \\
& =\mathrm{E}\left[e^{\lambda^{*} X_{\tau}}\right] \\
& \approx e^{\lambda^{*} a}(1-\alpha)+e^{\lambda^{*} b}(\alpha), \text { so } \\
\alpha & \approx \frac{e^{\lambda^{*} x}-e^{\lambda^{*} a}}{e^{\lambda^{*} b}-e^{\lambda^{*} a}} .
\end{aligned}
$$

To find $P=\mathrm{P}\left[X_{t}\right.$ ever exceeds $\left.b\right]$, take $a \rightarrow-\infty$ to find $P \approx e^{-\lambda^{*}(b-x)}<1$, if $\lambda^{*}>0$, or $P=1$, if $\lambda^{*} \leq 0$.
Since $\left(X_{n}-n \mu\right)$ is also a martingale,

$$
\begin{aligned}
x & =\mathrm{E}\left[X_{\tau \wedge n}-(\tau \wedge n) \mu\right] \\
& =\mathrm{E}\left[X_{\tau}-\tau \mu\right] \\
& \approx a(1-\alpha)+b(\alpha)-\mathrm{E}[\tau] \mu, \text { so } \\
\mathrm{E}[\tau] & \approx \frac{a-x+\alpha(b-a)}{\mu} .
\end{aligned}
$$

For example, for the standard asymmetric random walk on the integers with $\xi= \pm 1$ with probabilities $p, q=1-p$, respectively, then $\mu=p-q, \sigma^{2}=4 p q$, and $M(\lambda)=p e^{\lambda}+q e^{-\lambda}=1+\left(p e^{\lambda}-q\right)\left(1-e^{-\lambda}\right)$ so $m\left(\lambda^{*}\right)=0$ for $\lambda^{*}=\log q / p$. Thus for $a<x<b, \alpha=\mathrm{P}\left[X_{n} \geq b\right.$ before $\left.X_{n} \leq a \mid X_{0}=x\right]=$ $\left((q / p)^{x}-(q / p)^{a}\right) /\left((q / p)^{b}-(q / p)^{a}\right)$ and the exit time $\tau:=\min \left[n: X_{n} \notin(a, b)\right]$ has expectation $\mathrm{E}\left[\tau \mid X_{0}=x\right]=(a-x+\alpha(b-a)) / \mu$. If $\alpha \approx 0$ (resp., $\alpha \approx 1$ ) this is close to $\mathrm{E}[\tau] \approx(x-a) / \mu$ (resp., $\mathrm{E}[\tau] \approx(b-x) / \mu)$, just what you would expect for a heavily biased random walk.

## Sequential Probability Ratio Test

Let $\left\{Y_{j}\right\}$ be independent, identically-distributed random variables with absolutely-continuous distributions and density function $f(y)$, and consider the statistical problem of trying to tell from observed values $y_{1}, \ldots, y_{n}$ which of two possible density functions $\left\{f_{0}, f_{1}\right\}$ governs the distribution of the $\left\{Y_{j}\right\}$. All of the standard statistical tests of the hypothesis $H_{0}: f=f_{0}$ against its alternative $H_{1}: f=f_{1}$ make use of the Likelihood Ratio (against the null)

$$
\Lambda_{n}:=\frac{f_{1}\left(y_{1}\right) \cdots f_{1}\left(y_{n}\right)}{f_{0}\left(y_{1}\right) \cdots f_{0}\left(y_{n}\right)}=\prod_{j \leq n} \frac{f_{1}\left(y_{j}\right)}{f_{0}\left(y_{j}\right)}
$$

or, equivalently, its logarithm $\ell_{n}=\sum_{j \leq n} \log \left(f_{1}\left(y_{j}\right) / f_{0}\left(y_{j}\right)\right)$. For example, the Bayesian posterior probability of $H_{0}$ (starting with prior $\pi_{0}=\mathrm{P}\left[H_{0}\right], \pi_{1}=1-\pi_{0}=\mathrm{P}\left[H_{1}\right]$ ) is given by

$$
\begin{aligned}
& \mathrm{P}\left[H_{0} \mid Y_{1} \ldots Y_{n}\right]=\frac{\pi_{0} f_{0}\left(y_{1}\right) \cdots f_{0}\left(y_{n}\right)}{\pi_{0} f_{0}\left(y_{1}\right) \cdots f_{0}\left(y_{n}\right)+\pi_{1} f_{1}\left(y_{1}\right) \cdots f_{1}\left(y_{n}\right)}=\frac{\left(\pi_{0} / \pi_{1}\right)}{\left(\pi_{0} / \pi_{1}\right)+\Lambda_{n}} \\
& \frac{\mathrm{P}\left[H_{1} \mid Y_{1} \ldots Y_{n}\right]}{\mathrm{P}\left[H_{0} \mid Y_{1} \ldots Y_{n}\right]}=\frac{\pi_{1}}{\pi_{0}} \Lambda_{n}
\end{aligned}
$$

so the posterior odds against $H_{0}$ are the prior odds multiplied by $\Lambda_{n}$ (which in this context is called the "Bayes factor against $H_{0}$ "), while the Neyman-Pearson Lemma says that the most powerful (frequentist) test of level $\alpha$ is to reject $H_{0}$ whenever $\Lambda_{n} \geq r$, where $r$ is chosen to ensure that the probability of a "Type-I error" (rejecting a true null hypothesis) is no more than some specified value $\mathrm{P}_{0}\left[\Lambda_{n} \geq r\right] \leq \alpha$ if $H_{0}: f=f_{0}$ is true (the subscript zero on $\mathrm{P}_{0}$ indicates that this probability should be computed assuming $H_{0}$ ). The "power" of the test is then $\mathrm{P}_{1}\left[\Lambda_{n} \geq r\right]$, the probability of rejecting $H_{0}$ when in fact $H_{1}$ is true, or one minus the probability $\beta=\mathrm{P}_{1}\left[\Lambda_{n}<r\right]$ of a "Type-II error", failing to reject a false null hypothesis.
Only a large sample-size $n$ will ensure that both of these error probabilities will be small, but how large $n$ must be will depend on how different $f_{0}$ and $f_{1}$ are, something that may be difficult to anticipate. One possibility, initially proposed by Wald and Wolfowitz (1948), is to design a sequential test in which data are drawn successively until the evidence becomes compelling either that $H_{0}$ is false and must be rejected (large values of $\Lambda_{n}$ ) or that $H_{0}$ is true and must not be rejected (small values of $\Lambda_{n}$ ). A simple process is to select numbers $0<A<1<B$ and continue sampling until either $\Lambda_{n} \geq B$, in which case we stop and reject $H_{0}$, or $\Lambda_{n} \leq A$, in which case we stop sampling and accept $H_{0}$.
But this is exactly equivalent to drawing samples until the random walk $\ell_{n}:=\log \Lambda_{n}$, which starts at $x=\log 1=0$, reaches either a lower boundary $a:=\log A<0$ or an upper boundary $b:=\log B>0$, a problem we have just solved. For this random walk and any $\lambda \in \mathbb{R}$, under the hypothesis $H_{0}$, the means of $\xi_{j}:=\log \left(f_{1}\left(Y_{j}\right) / f_{0}\left(Y_{j}\right)\right)$ and of $\exp \left(\lambda \xi_{j}\right)$ are given by:

$$
\begin{aligned}
\mu_{0} & :=\mathrm{E}_{0}\left[\xi_{j}\right]=\int \log \frac{f_{1}(y)}{f_{0}(y)} f_{0}(y) d y \quad=-K\left(f_{0}: f_{1}\right) \\
M_{0}(\lambda) & :=\mathrm{E}_{0}\left[e^{\lambda \xi_{j}}\right]=\int f_{1}(y)^{\lambda} f_{0}(y)^{1-\lambda} d y
\end{aligned}
$$

and under hypothesis $H_{1}$ they are

$$
\begin{aligned}
\mu_{1} & :=\mathrm{E}_{1}\left[\xi_{j}\right]=\int \log \frac{f_{1}(y)}{f_{0}(y)} f_{1}(y) d y=K\left(f_{1}: f_{0}\right) \\
M_{1}(\lambda) & :=\mathrm{E}_{1}\left[e^{\lambda \xi_{j}}\right]=\int f_{1}(y)^{1+\lambda} f_{0}(y)^{-\lambda} d y
\end{aligned}
$$

so $\mu_{0}<0<\mu_{1}$. The quantity $K(f: g) \geq 0$ is the Kullback-Leibler divergence from $f$ to $g$, a measure of how different $f$ and $g$ are; for example, the K-L divergence from a standard normal distribution to the $\operatorname{No}(\mu, 1)$ distribution is $\mu^{2} / 2$. Also $M_{0}\left(\lambda^{*}:=1\right)=1=M_{1}\left(\lambda^{*}:=-1\right)$, so the exit time $\tau:=\min \left[n \geq 0: \ell_{n} \notin(a, b)\right]$ leads to right-exit probability (and Type-I error probability)
$\alpha=\mathrm{P}_{0}\left[\ell_{\tau} \geq b\right]$ and to left-exit probability (and Type-II error probability) $\beta=\mathrm{P}_{1}\left[\ell_{\tau} \leq a\right]$ of

$$
\begin{aligned}
& \alpha=\mathrm{P}_{0}\left[\ell_{\tau} \geq b\right] \approx \frac{e^{0}-e^{a}}{e^{b}-e^{a}}=\frac{1-A}{B-A} \\
& \beta=\mathrm{P}_{1}\left[\ell_{\tau} \leq a\right] \approx 1-\frac{e^{-0}-e^{-a}}{e^{-b}-e^{-a}}=\frac{A(B-1)}{B-A}
\end{aligned}
$$

with expected sample-size

$$
\begin{aligned}
& \mathrm{E}_{0}[\tau] \approx-\frac{(B-1) \log A+(1-A) \log B}{(B-A) K\left(f_{0}: f_{1}\right)} \\
& \mathrm{E}_{1}[\tau] \approx \frac{A(B-1) \log A+B(1-A) \log B}{(B-A) K\left(f_{1}: f_{0}\right)}
\end{aligned}
$$

In the symmetric case $A B=1, \alpha=\beta=1 /(1+B)=A /(1+A)$, and

$$
\mathrm{E}_{0}[\tau]=\frac{(B-1) \log B}{(B+1) K\left(f_{0}: f_{1}\right)} \quad \mathrm{E}_{1}[\tau]=\frac{(B-1) \log B}{(B+1) K\left(f_{1}: f_{0}\right)} .
$$

Evidently $\alpha$ and $\beta$ may be made as small as desired by taking $A^{-1}=B=\log \frac{1-\alpha}{\alpha}$ sufficiently large, but doing so will increase the expected sample size to approximately $\mathrm{E}[\tau] \approx \log \frac{1}{\alpha} / K\left(f_{0}: f_{1}\right)$.
Exercise 1a: Starting with $X_{0}=\$ 80$ and betting $\$ 1$ each turn at even odds, what is the chance of reaching $b=\$ 100$ before going broke (i.e., reaching $a=\$ 0$ )? On average, how long will it take to reach one of these?
Exercise 1b: Same question, but now playing US Roulette with probability $p=9 / 19$ of winning and $q=10 / 19$ of losing each turn.

Exercise 2: (Berk, 1966). Suppose that both hypotheses are wrong, and that $\xi_{j} \sim f(x) d x$ but $f \neq\left\{f_{0}, f_{1}\right\}$. Show that $\ell_{n}$ is again a random walk, now with drift $\mu=K\left(f: f_{1}\right)-K\left(f: f_{0}\right)$ and conclude that, almost surely, $\Lambda_{n} \rightarrow 0$ as $n \rightarrow \infty$ if $K\left(f: f_{0}\right)<K\left(f: f_{1}\right)$ and $\Lambda_{n} \rightarrow \infty$ as $n \rightarrow \infty$ if $K\left(f: f_{0}\right)>K\left(f: f_{1}\right)$. What do you think would happen if $K\left(f: f_{0}\right)=K\left(f: f_{1}\right)$ ?

Exercise 3a: Find $K\left(f_{0}: f_{1}\right)$ if each $f_{i}$ is $\operatorname{No}\left(\mu_{i}, \sigma^{2}\right)$ (different means, same variance).
Exercise 3b: Find $K\left(f_{0}: f_{1}\right)$ if each $f_{i}$ is $\operatorname{No}\left(0, \sigma_{i}^{2}\right)$ (same mean, different variances).
Exercise 3c: Find $K\left(f_{0}: f_{1}\right)$ if each $f_{i}$ is $\operatorname{Ex}\left(\lambda_{i}\right)$, exponentially distributed with rate $\lambda_{i}$.
Exercise 3d: Find $K\left(f_{0}: f_{1}\right)$ if each $f_{i}$ is $\operatorname{Bi}\left(N, p_{i}\right)$, binomial with the same $N$ but possibly different probabilities $p_{i}$.

## References

Berk, R. H. (1966), "Limiting Behavior of Posterior Distributions when the Model is Incorrect," Annals of Mathematical Statistics, 37, 51-58, doi:10.1214/aoms/1177699597.

Wald, A. and Wolfowitz, J. (1948), "Optimal Character of the Sequential Probability Ratio Test," Annals of Mathematical Statistics, 19, 326-339, doi:10.1214/aoms/1177730197.

