1.6 Sequential decision theory. As previously, let the sample space be a measurable space $(X, \mathcal{B})$ and $\mathcal{P} = \{P_\theta,\ \theta \in \Theta\}$ a family of laws on it. Let $X_1, X_2, \ldots,$ be independent and identically distributed with law P_θ. Let A be the space of possible (specific) actions, with a σ-algebra $\mathcal{E}$. We have a σ-algebra $\mathcal{T}$ on Θ and a loss function L which is a measurable function L from $\Theta \times A$ to $[0, \infty]$. A prior π may or may not be given on Θ.

A *sequential decision rule* will consist of two functions N and δ as follows. Let X^∞ be the set of all sequences $\{x_n\}_{n\geq 1}$ with $x_n \in X$ for all n. For $n = 1, 2, \ldots,$ let $\mathcal{B}_n$ be the smallest σ-algebra of subsets of X^∞ for which $X_1, \ldots, X_n$ are measurable. Let $\mathcal{B}_0$ be the trivial σ-algebra $\{\phi, X^\infty\}$. Then N is a function from X^∞ into $\mathbb{N} := \{0, 1, 2, \ldots\}$ such that for each $k = 0, 1, 2, \ldots,$ $\{N \leq k\} \in \mathcal{B}_k$. Such a function is called a *stopping time* or *stopping rule*. The *terminal decision rule*, δ, is a sequence of functions $\{\delta_n\}_{0\leq n<\infty}$ where $\delta_0 \in A$ and for each $n \geq 1$, δ_n is a measurable function from X^n into A. The action actually taken will be $\delta_N := \delta_N(X_1, \ldots, X_N)$. Let $\phi := \{N(\cdot), \delta\}$.

If $c \geq 0$ is the cost of each observation, the total loss (including costs of observations) in a given case is $L(\theta, \delta_N) + Nc$. The *risk* is then

$$r(\theta, \phi) := c \cdot EN + EL(\theta, \delta_N)$$

where the expectations are with respect to P_θ^∞ on X^∞. Note that if $c = 0$, and N is required to take finite values as in the above definition, then in general, optimal rules do not exist.

Example. This is actually not a statistical decision problem as just formulated but it will illustrate some possible difficulties.

Suppose that a gambler can play a sequence of games as follows. In the nth game, the gambler wagers \$1 and wins \$$100{\cdot}2^{n+1}$ with probability $0.01/2^n$. Thus the expected gain in each game is \$2 - \$1 = \$1. So the "Bayes" or optimal strategy would seem to be to continue playing indefinitely. But the probability that the gambler ever wins is $\leq \sum_{n\geq 1} 0.01/2^n = 0.01$. If the gambler never wins, which occurs with probability ≥ 0.99, then the gambler wagers and loses infinitely many dollars. The expected or average gain from games won is also infinite, if the gambler continues to play, so that the overall average gain is $\infty - \infty$ (undefined). There is actually no Bayes (optimal) strategy.

Let f_n be the net winnings after n plays. Then $Ef_n = n \to +\infty$ while $f_n \to -\infty$ a.s. by the Borel-Cantelli lemma.

We can also consider sequential *randomized* decision rules defined as follows. For $n = 1, 2, \ldots,$ let $(A_n, \mathcal{E}_n)$ be a measurable space, where A_n is the space of specific actions which can be taken after n observations. Often, all $(A_n, \mathcal{E}_n)$ will be equal to one space $(A, \mathcal{E})$. Assume that for each n, $0 \in A_n$, where the action "0" will mean taking another observation X_{n+1}, while all other actions in A_n will imply taking no more observations. Each ϕ_n is a measurable function from $(X^n, \mathcal{B}_n)$ into the space $D_\mathcal{E}$ of all probability laws on $(A, \mathcal{E})$. So, given $X_1, \ldots, X_n$, if no decision to stop has been made earlier, we then take another observation with probability $\phi_n(X_1, \ldots, X_n)(\{0\})$ and otherwise stop and take an action chosen from $A_n \setminus \{0\}$ with distribution $\phi_n(X_1, \ldots, X_n)/(1 - \phi_n(X_1, \ldots, X_n)(\{0\}))$.

So for a sequential randomized test of P vs. Q, we can take $A_n = A = \{-1, 0, 1\}$ for all n, where (as in Sec. 1.5) -1 means choosing P and $+1$ means choosing Q.

PROBLEMS

1. In the example at the end of Sec. 1.5 and where $R_{Q/P}$ has only values t or $1/t$ (not 1) with $t = 2$, let ϕ be a randomized test that does SPRT$(1/8, 2)$ or SPRT$(1/2, 8)$ with probability 1/2 each. Compare the performance of this test to SPRT$(1/4, 4)$ in terms of error probabilities and average sample numbers.
2. For a sequential test of P vs. Q as in Problem 1, suppose that the nth observation costs $1/3^n$, while $L_{PQ} = L_{QP} = 3$ and $\pi(P) = \pi(Q) = 1/2$. A decision rule must reach a decision after a finite number N of observations. Is there an optimal (Bayes) sequential test in this case? Why, or why not?