332:541 Stochastic Signals and Systems Fall 1999

Final Examination

Statement and Solutions

You have 3 hours to answer the following two questions. You are allowed four sides of 11x8.5 inch pieces of paper for notes. GOOD LUCK!

1. *(50)* **The Big Pay Day:** Fly By Night investment Corporation approaches you about an investment scheme. On any given day, if you give them C dollars, at the end of the day you will have doubled your money with probability 1/2, lost half of your money with probability 1/4 and lost 3/4 of your money with probability 1/4.

   (a) *(10)* Let X_1 be the random variable which denotes your investment portfolio at the end of day 1. Assuming you placed C dollars in the portfolio at the beginning of the day, what are $E[X_1]$ and the variance of X_1?

   Expected value:

$$m = E[X_1] = \frac{1}{2}(2C) + \frac{C}{2}\frac{1}{4} + \frac{C}{4}\frac{1}{4} = \frac{19}{16}C = 1.1875C$$

   Variance:

$$\begin{aligned} E[(X_1 - m)^2] &= E[X_1^2] - m^2 \\ &= \tfrac{1}{2}(4C^2) + \tfrac{C^2}{4}\tfrac{1}{4} + \tfrac{C^2}{16}\tfrac{1}{4} - (\tfrac{19}{16})^2 C^2 \\ &= C^2(2 + \tfrac{1}{16} + \tfrac{1}{64} - \tfrac{361}{256}) \\ &= \tfrac{171}{256}C^2 = 0.668C^2 \end{aligned}$$

   (b) *(10)* Suppose you take all the money in the portfolio on each day and re-invest it the following day. Let X_n be the size of your portfolio at the end of day n. Assuming that the result for each day is independent of all other days and that you invested C dollars at the beginning of day 1, what is the expected size of your portfolio at the end of day n. What is the variance?

   We've got independent trials so it's the product of expectations for the individual days. That is, the result at the end of day k is a random variable Z_k times X_{k-1} where Z is distributed as

$$p_{Z_k}(z_k) = \frac{1}{2}\delta(z_k - 1) + \frac{1}{4}\delta(z_k - \frac{1}{2}) + \frac{1}{4}\delta(z_k - \frac{1}{4})$$

   Since each day's factor Z_k is independent of all the others days we have $X_n = C\prod_{i=1}^n Z_i$ and

$$E[X_n] = \left(\frac{19}{16}\right)^n C$$

   The variances multiply as well and we have

$$\sigma_{X_n}^2 = C^2 \left(\frac{171}{256}\right)^n$$

Notice that the variance goes toward infinity with increasing n and that it seem to go there much faster than the square of the mean does. That should be your first clue that you'll need to be careful in making quantitative judgements about what's going on.

(c) *(10)* Define $Y_n = \log_2 X_n$. Calculate the expected value and variance of Y_n.

Same thing, but now we have a sum

$$Y_n = \sum_{k=1}^{n} \log Z_k + \log C$$

so

$$E[Y_n] = \sum_{k=1}^{n} E[\log Z_k] + \log C = n(1/2 - 1/4 - 1/2) + \log C = \log C - n/4$$

The variance is

$$\sigma^2_{Y_n} = (1/2 + 1/4 + 1 - 1/16)n = \frac{27}{16}n$$

What you want to notice is that the mean scales in n, but SO DOES THE VARIANCE!!!! That means that the standard deviation scales in $\sqrt{n}$.

(d) *(20)* Is this a good investment scheme? Obviously, yes or no will not suffice as an answer.

Any company which goes by the name fly-by-night is suspect. The the point is that in the previous parts, in the first two parts it looked like the investment scheme was good – the expected value of the portfolio tended to infinity. However, in the third part, we saw that the expected value of the log of the portfolio tended toward negative ∞. *This implies that the portfolio tends to zero. This is a contradiction.*

However, in the third part, we note that as $n \to \infty$, *the random variable* Y_n *is a gaussian with variance proportional to n. If we define a new random variable* $Q_n = \frac{1}{\sqrt{n}} Y_n$ *we see that* Q_n *also tends toward a Gaussian but with mean* $-1/4$ *and variance* $\frac{27}{16}\frac{1}{n}$ *— which tends to zero as* $n \to \infty$. *This means that as n goes to infinity,* Q_n *goes to* $-1/4$ *with probability 1!!!!!!!!. This in turn implies that* Y_n *goes to negative infinity with probability 1! And finally, the implication is that* $\lim_{n\to\infty} Prob(X_n = 0) = 1$ *so all the money is lost with probability 1.*

So the investment scheme sucks. The moral is that you must always look at the distribution of a random variable (or at least it's variance relative its mean) before you make any decisions based on second order statistics).

2. *(100 points)*

Rutgera and Fisher the Bird: Rutgera Universa, the world famous Rutgers University Graduate student owns a parakeet (small bird) named Fisher. Fisher is in either of two states: happy or sad. When Fisher is happy, he chirps at rate λ_H. When sad, he chirps at rate $\lambda_S < \lambda_H$. Whether Fisher is happy or sad depends on a variety of things, but the end result is that Fisher's happy and sad state transitions can be modeled as a continuous time markov process with transition rate from sad to happy, α, and from happy to sad β.

Every animal must eat and Fisher is no exception. However, when Fisher is sad, if he's fed he may get indigestion and die. As you can see parakeets are 1) very delicate and 2) not that intelligent.

So, Rutgera's problem is to listen to Fisher's chirping, figure out which of two states he's in, and either feed him or not.

(a) *(20 points)* Given α and β, please provide the differential equations for Fisher's two-state markov process. Derive a solution for $t \gg t_0$ assuming that Fisher is happy at $t = t_0$. What are the average proportions of time Fisher spends being happy and sad in terms of α and β. What is the probability that at some randomly chosen time ($t \gg t_0$) we find Fisher in the happy state.

Let x_s be the probability Fisher is sad and $1 - x_s$ the probability he is happy. We then have

$$\dot{x_s} = -\alpha x_s + \beta(1 - x_s)$$
$$1 \dot{-} x_s = -\dot{x_s} = -\beta(1 - x_s) + \alpha x_s$$

which are obviously equivalent. Rewriting we have

$$\dot{x_s} = -(\alpha + \beta)x_s + \beta$$

The solution to this simple first order linear differential equation for $t \geq t_0$ is

$$x_s(t) = e^{-(\alpha+\beta)(t-t_0)}(x_s(t_0) - \frac{\beta}{\beta + \alpha}) + \frac{\beta}{\beta + \alpha}$$

So, as $t \to \infty$ we have the (steady state) probability Fisher is in the sad state as $\frac{\beta}{\beta+\alpha}$ and in the happy state as $\frac{\alpha}{\beta+\alpha}$. So at a randomly chosen point in time the probability we find Fisher happy is $\frac{\alpha}{\beta+\alpha}$.

It turns out that the wording left you an out of just providing the steady state probabilities since I insisted upon large t. If you wrote down the steady state probability you got (begrudgingly) full credit.

(b) *(20 points)* Given an observation interval of duration T for a homogeneous poisson process, please provide an estimator for the rate. Is your estimator biased or unbiased. Is it efficient or inefficient? Is it asymptotically unbiased and/or efficient?

One could try to derive a maximum likelihood estimator or some other such, but the simplest way is to choose something simple and try it. So, let $\hat{\lambda} = n/T$ where n is the number of arrivals observed in the interval $[0, T]$. Since $E[n] = \lambda T$, $E[\hat{\lambda}] = \lambda$ so the estimator is unbiased. But is it efficient? Well we now need to calculate the variance of the estimator. That is

$$\sigma^2_{\hat{\lambda}} = E[(n/T - \lambda)^2] = E[(n/T)^2] - \lambda^2$$

Since we already know that the variance of n is λT we must have $E[n^2] = \lambda T + \lambda^2 T^2$. Thus

$$\sigma^2_{\hat{\lambda}} = \lambda / T$$

which is rather nice since it means our estimator variance goes to zero as $T \to \infty$. But now we need to calculate the Cramer-Rao bound and compare the variance of our estimator to it.

The likelihood function is

$$L(\lambda, n) = \log f_{N|\lambda}(n|\lambda) = n \log(\lambda T) - \log n! - (\lambda T) = n \log \lambda + n \log T - \log n! - \lambda T$$

It's derivative with respect to the parameter is

$$\frac{\partial L(\lambda, n)}{\partial \lambda} = n/\lambda - T$$

The expected value of the square of this quantity is

$$I = E[(n/\lambda - T)^2] = \frac{1}{\lambda^2} E[(n - T\lambda)^2] = T/\lambda$$

and $1/I$ is identically equal to the variance of our estimator. So our estimator is efficient.

(c) *(20 points)* Assume $0 \gg t_0$. Rutgera observes Fisher on $[0, T]$. What is the probability that Fisher experiences at least one transition (either happy $\to$ sad or sad $\to$ happy) on this interval? How big or small must T be such that the probability of no transitions on $[0, T]$ is p. Call this value T^*.

The probability Fisher is happy at $t = 0$ is $\frac{\alpha}{\alpha+\beta}$. The probability Fisher is sad at $t = 0$ is $\frac{\beta}{\alpha+\beta}$. Given Fisher is happy at $t = 0$, the probability he stays happy for T seconds is $e^{-\beta T}$ since the transition process is exponentially distributed. Likewise, the probability he stays sad given he's sad at $t = 0$ is $e^{-\alpha T}$. The probability there are no transitions on $[0, T]$ is then

$$p = \frac{\alpha}{\alpha + \beta} e^{-\beta T} + \frac{\beta}{\alpha + \beta} e^{-\alpha T}$$

This needs to be solved numerically for T^*.

(d) *(20 points)* Given T^* and $t_0 \ll 0$ and assuming there are no transitions on the observation interval, please devise a decision rule to determine whether Fisher was happy or sad on the observation interval based on your rate estimate which has minimum probability of error.

This is a decision problem. So there are two hypotheses with a priori *probabilities*

$$p_S = \frac{\beta}{\alpha + \beta}$$

and

$$p_H = \frac{\alpha}{\alpha+\beta}$$

at time $t = 0$. So we form the likelihood ratio

$$\frac{p_H f_{N|\lambda_H}(n|\lambda_H)}{p_S f_{N|\lambda_S}(n|\lambda_S)} \begin{array}{c} \text{say HAPPY} \\ > \\ < \\ \text{say SAD} \end{array} 1$$

Rewriting we have

$$\left(\frac{\lambda_H}{\lambda_S}\right)^n \begin{array}{c} \text{say HAPPY} \\ > \\ < \\ \text{say SAD} \end{array} \frac{p_S}{p_H} e^{(\lambda_H - \lambda_S)T^*}$$

and this is the decision rule.

(e) *(20 points)* Assume that Fisher takes S seconds to eat, and that if during feeding he is sad at any point, he dies. What is the probability that for a randomly chosen feeding time, he survives. Can Rutgera's observations and estimators improve Fisher's chances of surviving a given feeding?

The probability he survives is the probability Rutgera does not start feeding him when he's sad and that he doesn't ever get sad while eating. So for random incidence, the probability he's happy is $p_S = \frac{\alpha}{\alpha+\beta}$. The probability there are no transitions during a feeding interval of duration S can be found from part (c) of this problem (with T replaced by S). So the probability he survives with a randomly chosen feeding time is

$$Prob(survive|random\ incidence) = \frac{\alpha}{\alpha+\beta} e^{-\beta S}$$

The short answer to the second part is yes. The reason is that given a fixed feeding interval, his survival depends only on Rutgera's ability to figure out whether Fisher is happy or sad at $t = 0$ (p_H). So, since observations can only increase (or at least not decrease) the probability that Rutgera begins a feeding when Fisher is happy, Rutgera's machinations must help (assuming she uses her info wisely). Of course, that's pretty unsatisfying for my wonderfully curious Rutgers ECE graduate students. However, if you said something like the above, then you got full credit.

To delve deeper, the issue is how well Rutgera can determine whether Fisher is happy or sad based on measurements before $t = 0$. I don't think we really have the machinery to deal with this – the mutual information between all chirp arrival times for $t < 0$ and the state of Fisher at $t = 0$. This is why I did not ask you for the maximum possible survival probability though I had originally constructed the *problem to ask just that question.*

Of course, there might be a simpler way which I don't see. This problem is inspired by a topic called hidden markov models. They are used in speech recognition and also in network congestion control. In the networking literature, the seminal paper is by San Qi Li at the University of Texas, Austin. The problem is slightly different, but the basic idea is that you'd like to estimate the state of a network (congested or not) and apply some control action (feeding or not feeding :). There is also a problem in neurophysiology (and chemistry) where you don't even know the transition rates (α *and* β*) and must estimate them from your observations of the arrival process. I did my doctoral thesis in auditory physiology on an aspect of this problem.)*