STA-226, Spring Problem set 9. Due Thursday April ffth.

1. Suppose that $x_1, x_2, \ldots$ are conditionally independent Bernoulli experiments, all with success probability θ. You interested in estimating θ and your terminal decision loss function is

$$L(\theta, \boldsymbol{a}) = \frac{(\theta - \boldsymbol{a})^2}{\theta(1-\theta)}.$$

This is the traditional squared error loss, standardized using the variance of a single observation. With this loss function, errors of the same absolute magnitude are considered worse if the are made near the extremes of the interval $(0, 1)$. You have a uniform prior on θ. Each observation costs you c. What is the optimal fixed sample size? You may start using computers at any stage in the solution of this problem.

2. You are the mayor of Solvantville and you interested in estimating the difference in the concentration of a certain contaminant in the water supply of homes in two different sections of town, each with its own well. If you found that a difference you woould further investigate whether there may be contamination due to industrial waste. How big a sample of homes do you need for your purposes? Specify a terminal loss function, a cost of observation, a likelihood and a prior. Keep it very simple, but justify your assumptions and point out limitations if necessary.

3. Can you find the optimal sample size in example 18.0.1 when σ is not known? You can use a conjugate prior, and standard results such as those of section 3.3 in "Bayesian Data Analysis" by Gelman et al.

4. Sequential problems get very hard very quickly. This is one of the simplest you can try, but it will take some time.

Try to address the same question as worked exercise 19.4, except with Poisson observation and a Gamma$(1, 1)$ prior. Set $c = 1/12$.