STA 450/2102 — Assignment #3

Due at the beginning of class on March 24. Worth 18% of the final mark.

Part 1: Write an S/R function called `midpoint.rule` that approximates the integral of a function using the midpoint rule. This function should take as arguments the function to integrate, the low and high ends of the range over which to integrate it, and the number of points to evaluate the function at.

If you wish, you can write the function to take further arguments that are passed on as additional arguments to the function being integrated, as in the `trapezoidal.rule` example, but this isn't required. You could write `midpoint.rule` so that instead of using a fixed number of points, it increases the number of points used until some convergence criterion is reached, but this also is not required.

Test your function by integrating some simple functions for which you know the answer, and hand in a few such tests.

Part 2: A finance researcher studying the Canadian dollar versus US dollar exchange rate has concluded that days when the rate goes up and days when the rate goes down occur equally often. However, the researcher thinks that the direction of movement may depend to some extent on which way the rate moved the previous day.

Data on these changes is available in which a $+1$ indicates an upward move for the Canadian dollar, and a -1 indicates a downward move (we'll assume that exchange rates are kept to such high precision that no change at all is very unlikely). For instance, the data might be

$$+1 \quad +1 \quad -1 \quad +1 \quad +1 \quad +1 \quad -1 \quad -1 \quad -1 \quad +1 \quad +1$$

indicating that the Canadian dollar when up the first two days, then it went down, etc.

We might model these data values, $x_0, x_1, \ldots$, as a Markov process, with

$$P(X_0 = +1) \ = \ P(X_0 = -1) \ = \ 1/2$$

$$P(X_{i+1} = +1 \mid X_i = +1) \ = \ P(X_{i+1} = -1 \mid X_i = -1) \ = \ \theta$$
$$P(X_{i+1} = -1 \mid X_i = +1) \ = \ P(X_{i+1} = +1 \mid X_i = -1) \ = \ 1 - \theta$$

for $i = 0, 1, \ldots$, with θ being an unknown model parameter between 0 and 1.

The researcher thinks that θ might be close to $1/2$, or that it might be close to 0 or 1, or with somewhat smaller probability, it might be anything else. To express this, they decide to use a prior distribution for θ whose density is proportional to the following:

$$\theta^{-1/2}\,(1-\theta)^{-1/2}\,(0.2+\exp(-100(\theta-1/2)^2))$$

Here is a plot of this unnormalized density, showing that it favours values near 0, 1/2, and 1:

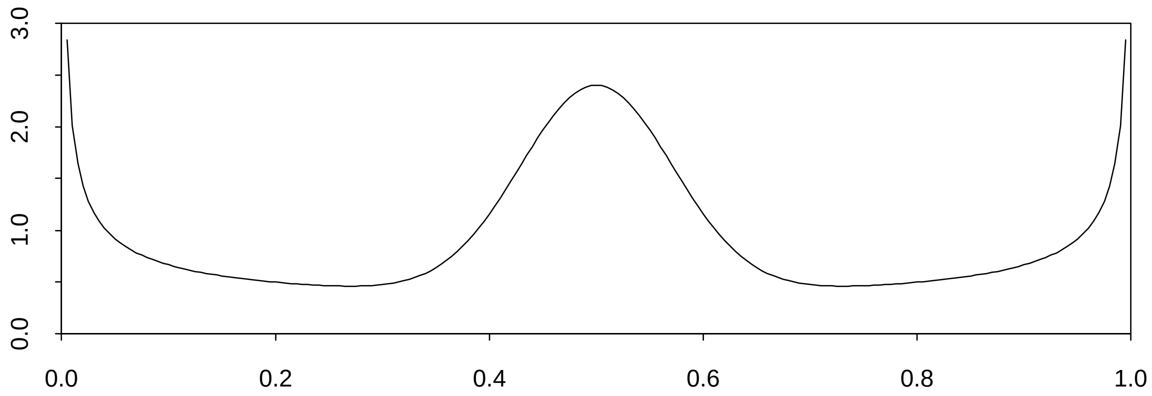

Note that the density goes to infinity as θ approaches 0 or 1.

Use the `midpoint.rule` procedure you wrote for Part 1 to answer the following questions:

1. What is the normalizing constant needed to convert the expression proportional to the prior density above to an actual prior density?

2. Given the eleven data points above, what is the Bayesian prediction for the probability that the next data point will be $+1$?

3. What is the Bayes factor giving the ratio of the probability of the eleven data points under this model and prior to the probability of these data points under the simple model in which every data point is independent, and equally likely to be $+1$ or -1?

You should decide on an appropriate number of points to use for the midpoint rule by seeing at what point increasing the number of points makes little difference.

If possible, you should summarize the data using some simple sufficient statistic, which you compute once at the beginning.

Discuss how fast the midpoint rule seems to be converging for these problems.