

A demonstration of the Bootstrap Filter

The purpose of this note is to provide an illustration of the Bootstrap Filter algorithm. We observe 5 independent observations from a Poisson distribution with mean θ and use a Gamma(10,1) prior on θ . We know the posterior in this case is again a Gamma(58,6) distribution with mean 9.666 and standard deviation 1.269.

```
set.seed(1857)

# data
x = c(8,11,10,14,5)

# prior on theta is gamma(10,1)
# posterior distribution is gamma(58,6)
# posterior mean = 9.666, posterior sd = 1.269
```

We know run the Bootstrap filter with 10000 prior draws.

```
# -- Bootstrap filter -- #
T = 10000 # particles
theta = rgamma(T,10,1) # sample from prior
lgprob = -5*theta + sum(x)*log(theta) # calculate log-lik
lgprob = lgprob - max(lgprob) # for better conditioning
postprob = exp(lgprob)/sum(exp(lgprob)) # calculate weights
postme = sum(postprob*theta) # posterior mean
postv = sum(postprob*theta^2) - postme^2 # posterior variance
postsd = sqrt(postv) # posterior sd
```

Let's print the posterior mean and standard deviation obtained from the BF.

```
c(postme,postsd)
```

```
## [1] 9.678683 1.260589
```

We have obtained a discrete approximation to the posterior with support points theta and weights in postprob. If we want to draw a histogram of the samples, we need to keep in mind that we have weighted samples here. There are some R functions which can take weighted samples to plot histograms. I am simply going to resample a large number of times from this discrete distribution and plot a histogram of these independent samples

```
numsamp = 5000 # number of samples to draw
thetasamp = sample(theta,numsamp,prob=postprob) # draw samples
numsampu = length(unique(thetasamp)) # number of unique samples
hist(thetasamp,prob="TRUE") # draw histogram
```

