STA 360/602L: Module 2.4

Truncated priors and the inverse cdf method

Dr. Olanrewaju Michael Akande

STA 602L

Priors with restricted support

- As we have seen, when dealing with rare events, we might expect the true proportion to be very small.
- In that case, we might want to try a restricted prior, e.g. $\text{Unif}(0, 0.1)$.
- Even when we don't have rare events, we might still desire some restriction if we are certain the true proportion lies within (a, b) with $0 < a < b < 1$.
- It is thus often really useful to explore "truncation" on priors.
- Let $\theta =$ probability of a randomly-selected student making an A in this course.
- You may want to rule out very low & very high values -- perhaps $\theta \in [0.35, 0.6]$ with probability one.
- How to choose a prior restricted to this interval?

STA 602L

Uniform priors

- One possibility is to just choose a uniform prior.
- When the parameter θ is a probability, the typical uniform prior would correspond to $\mathrm{Beta}(1,1)$.
- This is uniform on the entire $(0,1)$ interval.
- However, we can just as easily choose a uniform prior on a narrower interval $\mathrm{Unif}(a,b)$ with $0 < a < b < 1$.
- Perhaps not flexible enough.
- Would be nice if we could pick a flexible beta density and then truncate it to (a,b).

Truncated random variables

- Suppose we have some arbitrary random variable $\theta \sim f(\theta)$ with support Θ.
- For example, $\theta \sim \text{Beta}(c, d)$ has support on $(0, 1)$.
- Then, we can modify the density $f(\theta)$ to have support on a sub-interval $[a, b] \in \Theta$.
- The density $f(\theta)$ **truncated** to $[a, b]$ is

$$f_{[a,b]}(\theta) = \frac{f(\theta) 1[\theta \in [a, b]]}{\int_a^b f(\theta^\star) \mathrm{d}\theta^\star},$$

with $1[A]$ being the indicator function that returns 1 if A is true & 0 otherwise.

STA 602L

Truncated beta density

- Suppose to characterize the prior probability of earning an A, you poll a sample of students from a former STA 602 course and find that 10 earned an A and 10 earned a B or lower.
- Therefore, you go with a $\mathrm{Beta}(10, 10)$ prior truncated to $[0.35, 0.6]$.
- In R we can calculate the truncated beta density at p via

```
p <- seq(0,1,length=1000)
f1 <- dbeta(p,10,10)
f2 <- dbeta(p,10,10)*as.numeric(p>0.35 & p<0.6)/(pbeta(0.6,10,10) - pbeta(0.3,10,10))
f3 <- dunif(p,0.35,.6)
plot(p,f2,type='l',col='green4',xlim=c(0,1),ylab='Density', xlab=expression(theta),
  ylim=c(0,6))
lines(p,f1,type='l',col='blue')
lines(p,f3,type='l',col='red4')
labels <- c("beta(10,10)", "truncated beta","unif(0.35,.6)")
legend("topright", inset=.05, labels, lwd=2, lty=c(1,1,1), col=c('blue4','green4','re
```

STA 602L

Truncated beta density

What would that look like?

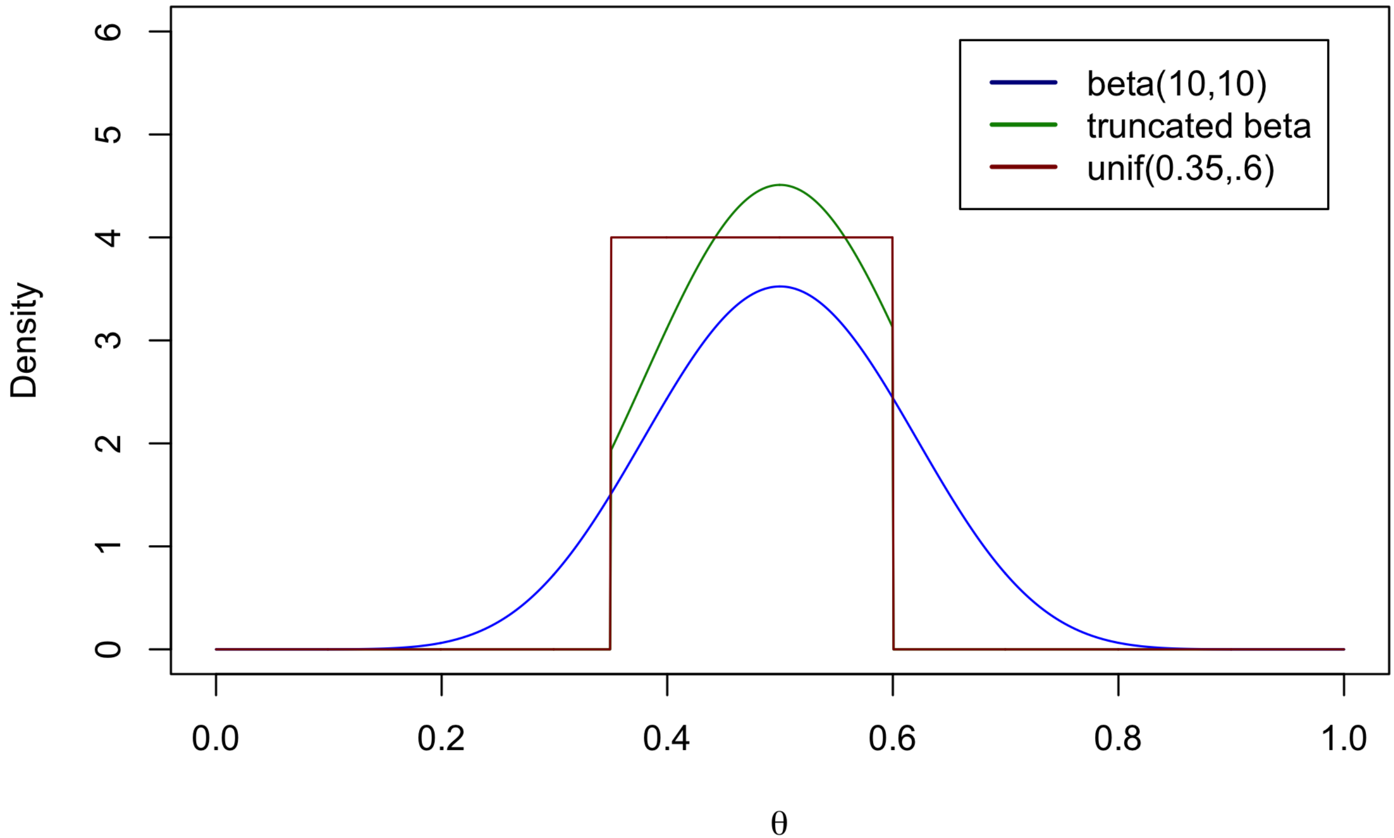

STA 602L

Truncated beta density

The truncated density by itself would look like

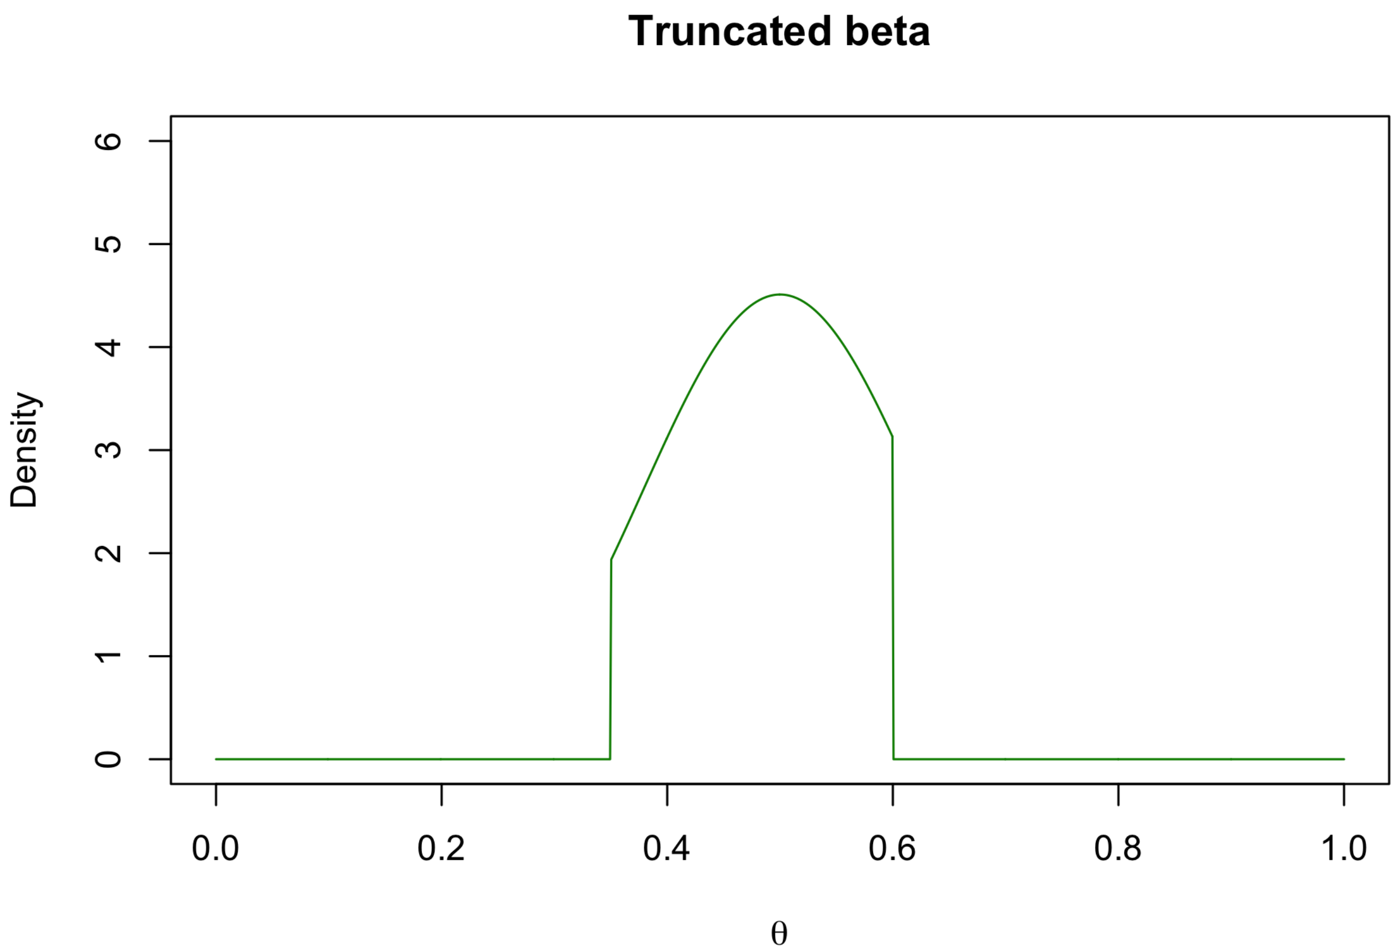

THE INVERSE CDF METHOD

- How to sample truncated random variables?
- First start with the pdf for an untruncated distribution such as $\theta \sim \text{Beta}(c, d)$.
- Suppose we then want to sample $\theta \sim \text{Beta}_{[a,b]}(c, d)$. How can we do that? One popular method is the inverse-cdf method.
- The inverse cdf is useful for generating random variables in general, especially for generating truncated random variables.
- Suppose we have $\theta \sim f(\theta)$, for some arbitrary continuous density f.
- According to probability integral transform, for any continuous random variable θ, the random variable $U = F(\theta)$ has a $\text{Unif}(0, 1)$ distribution. Note that F is the cdf.
- Thus, to use the inverse-cdf method to sample $\theta \sim f$, first sample $u \sim \text{Unif}(0, 1)$, then set $\theta = F^{-1}(u)$.

STA 602L

The inverse cdf method

- As an example, suppose we want to sample $\theta \sim \text{Beta}(c, d)$ through the inverse cdf method.
- Very easy. Just do the following in R.

```
u <- runif (1, 0, 1)
theta <- qbeta(u,c,d)
```

- That is, first sample from a uniform distribution.
- Then, transform it using the inverse cdf of the $\text{Beta}(c, d)$ distribution.
- Viola!

The inverse cdf method

- Back to the original problem: how to sample $\theta \sim \text{Beta}_{[a,b]}(c, d)$?
- If we had the inverse cdf of $\text{Beta}(c, d)$ truncated to $[a, b]$, then we could use the inverse cdf method. Easy enough! Let's find that inverse cdf.
- Let f, F and F^{-1} denote the pdf, cdf and inverse-cdf without truncation and let $A = [a, b]$.
- Recall that the density $f(\theta)$ **truncated** to $[a, b]$ is

$$f_A(\theta) = f_{[a,b]}(\theta) = \frac{f(\theta)1[\theta \in [a,b]]}{\int_a^b f(\theta^\star)\mathrm{d}\theta^\star} = \frac{f(\theta)1[\theta \in [a,b]]}{F(b) - F(a)}.$$

- Therefore, the truncated cdf

$$F_A(z) = \Pr[\theta \leq z] = \frac{F(z) - F(a)}{F(b) - F(a)}.$$

- Not enough though. We need the truncated inverse cdf.

STA 602L

The inverse cdf method

- To find the inverse cdf $F_A^{-1}(u)$, let $F_A(z) = u$. That is, set

$$u = F_A(z) = \frac{F(z) - F(a)}{F(b) - F(a)}$$

and solve for z as a function of u.

- Re-expressing as a function of $F(z)$,

$$F(z) = \{F(b) - F(a)\}u + F(a).$$

- Applying the untruncated inverse cdf F^{-1} to both sides, we have

$$z = F^{-1}[\{F(b) - F(a)\}u + F(a)] = F_A^{-1}(u).$$

STA 602L

The inverse cdf method

- We now have all the pieces to use the inverse-cdf method to sample $\theta \sim f_A$, that is, f truncated to A.
- First draw a $\text{Unif}(0, 1)$ random variable

```
u <- runif (1, 0, 1)
```

- Next, apply the linear transformation:

$$u^\star = \{F(b) - F(a)\}u + F(a).$$

- Finally, plug $u^\star$ into the untruncated cdf $\theta = F^{-1}(u^\star)$.
- Note we can equivalently sample $u^\star \sim runif(1, F(a), F(b))$.

STA 602L

What's next?

Move on to the readings for the next module!