

Figure 11.1. Upper envelope for acceptance/rejection. Solid line is rescaled t density with five degrees of freedom; dashed line is upper envelope density function $g(x)$.

Figures 11.1 and 11.2 show the use of acceptance/rejection for generating from a Student's t distribution with five degrees of freedom (Algorithm C2 in the next section). In Figure 11.1, we see the rescaled density $\beta f(x)$ with the upper envelope density $g(x)$ lying above it. One view of acceptance/rejection is to produce a point uniformly distributed in the striped region under the target density's curve, so that the marginal density of X is as desired. Notice that the two curves follow each other's shapes, but not particularly well. In Figure 11.2, however, we see the effectiveness of the inner and outer bounds. The region below the inner bound is shown in solid gray, and the outer bound cuts off most of the clear region below the envelope curve $g(x)$ so that the (unnormalized) density $\beta f(x)$ is computed only in the region between these two bounds.

The number of trials needed before acceptance obviously follows a geometric distribution with probability β ; hence the expected number of trials is $1/\beta$, the inflation factor in the envelope $g(x)/\beta$. The use of inner and outer bounds doesn't change this, but it speeds up each trial. When β is near 1, the inner bound $b(x)$ becomes more important; as β decreases, the quick-reject outer bound $B(x)$ becomes more important.




Figure 11.2. Inner and outer bounds for acceptance/rejection. Dashed line is the upper envelope density; solid line is quick-reject upper bound. Triangular lower bound forms solid quick-accept region.

The t distribution has the density $f_\alpha(x) = c_\alpha u_\alpha(x)$, where

$$c_\alpha = \Gamma\left(\frac{\alpha+1}{2}\right) / \left[\sqrt{\pi\alpha} \Gamma\left(\frac{\alpha}{2}\right) \right] \quad \text{and} \quad u_\alpha(x) = \left(1 + \frac{x^2}{\alpha}\right)^{-(\alpha+1)/2},$$

which looks quite complicated. Nonetheless, for $\alpha \geq 1$, we can show that the unnormalized density $u_\alpha(x)$ lies below a simple rescaled density function:

$$u_\alpha(x) \leq \min(1, |x|^{-2}) \equiv 4g(x),$$

leading to the envelope inequality $\beta f_\alpha(x) \leq g(x)$ with $\beta = 1/(4c_\alpha)$. For acceptance/rejection, then, we need to generate from the density

$$g(x) = (1/4) \min(1, |x|^{-2}), \quad (11.3.6)$$

which can be done rather easily: Generate $Y \sim \text{uniform}(-1, 1)$; with probability $1/2$ return $X = Y$, else return $X = 1/Y$. The main acceptance/rejection test takes on the simplified form

$$U \leq \beta \frac{f(X)}{g(X)} = \left(\frac{1}{4c_\alpha}\right) \frac{c_\alpha u_\alpha(X)}{g(X)} = \frac{u_\alpha(X)}{\min(1, |X|^{-2})} = \frac{u_\alpha(X)}{4g(X)}.$$

This does not involve the ugly constant c_α , which therefore need not be computed. Pushing further, the density kernel $u_\alpha(x)$ can be avoided using the inner (quick-accept) bound

$$1 - |x|/2 \leq u_\alpha(x)$$

and the upper bounds

$$u_\alpha(x) \leq [2u_\alpha(1)]u_1(x) \leq [2u_\infty(1)]u_1(x) = 2e^{-1/2}/(1+x^2). \quad (11.3.7)$$

The setup version of the acceptance/rejection algorithm for the t distribution takes the following form.

Algorithm C2 (Simplified TIR from Kinderman, Monahan, and Ramage 1977)

(gttir)

- (0) (Setup) Compute $d_\alpha = 2u_\alpha(1) = 2(1 + 1/\alpha)^{-(\alpha+1)/2}$.
- (1) Generate $V \sim \text{uniform}(0, 1)$ and $U \sim \text{uniform}(0, 1)$.
- (2) If $V \leq 1/2$, then set $X = (1/4)/(V - (1/4))$ and $W = U/X^2$;
else $X = 4 * V - 3$ and $W = U$
(now $X \sim g(x)$ and $W = U * (4g(X)) \sim \text{uniform}(0, 4g(X))$).
- (3) If $W \leq 1 - |X|$ then deliver X .
- (4) If $W > d_\alpha/(1 + X^2)$ then go to (1).
- (5) If $W \leq u_\alpha(X)$ then deliver X ; else go to (1).

The non-setup version of Algorithm C2 simply replaces d_α with $d = 2e^{-1/2}$, which works for all α . However, further improvements exploit details avoided here for simplicity – namely, that the inner bound applies only for $|X| < 1$ and that the upper bound is useful only for $|X|$ in the interval $[(2u_\alpha(1) - 1)^{1/2}, (2u_\alpha(1) - 1)^{-1/2}]$. One of the critical concerns of an algorithm for the range $\alpha \geq 1$ is that the performance not degrade for large values of α . Indeed, the amount of computation does not vary greatly with α , and the expected number of trials is confined to a narrow range:

$$1.27 \leq 1/\beta = 4c_\alpha \leq 1.60.$$

In examining the upper bounds, one may also consider using the Cauchy distribution for the upper envelope $g(x)$. However, the Cauchy density is not as easily computed as a uniform and a multiply or divide (see Exercise 11.8).