

7.3 Rejection Method: Gamma, Poisson, Binomial Deviates

The *rejection method* is a powerful, general technique for generating random deviates whose distribution function $p(x)dx$ (probability of a value occurring between x and $x + dx$) is known and computable. The rejection method does *not* require that the cumulative distribution function [indefinite integral of $p(x)$] be readily computable, much less the inverse of that function — which was required for the transformation method in the previous section.

The rejection method is based on a simple geometrical argument:

Draw a graph of the probability distribution $p(x)$ that you wish to generate, so that the area under the curve in any range of x corresponds to the desired probability of generating an x in that range. If we had some way of choosing a random point *in two dimensions*, with uniform probability in the *area* under your curve, then the x value of that random point would have the desired distribution.

Now, on the same graph, draw any other curve $f(x)$ which has finite (not infinite) area and lies everywhere *above* your original probability distribution. (This is always possible, because your original curve encloses only unit area, by definition of probability.) We will call this $f(x)$ the *comparison function*. Imagine now that you have some way of choosing a random point in two dimensions that is uniform in the area under the comparison function. Whenever that point lies outside the area under the original probability distribution, we will *reject* it and choose another random point. Whenever it lies inside the area under the original probability distribution, we will *accept* it. It should be obvious that the accepted points are uniform in the accepted area, so that their x values have the desired distribution. It should also be obvious that the fraction of points rejected just depends on the ratio of the area of the comparison function to the area of the probability distribution function, not on the details of shape of either function. For example, a comparison function whose area is less than 2 will reject fewer than half the points, even if it approximates the probability function very badly at some values of x , e.g., remains finite in some region where x is zero.

It remains only to suggest how to choose a uniform random point in two dimensions under the comparison function $f(x)$. A variant of the transformation method (§7.2) does nicely: Be sure to have chosen a comparison function whose indefinite integral is known analytically, and is also analytically invertible to give x as a function of “area under the comparison function to the left of x .” Now pick a uniform deviate between 0 and A , where A is the total area under $f(x)$, and use it to get a corresponding x . Then pick a uniform deviate between 0 and $f(x)$ as the y value for the two-dimensional point. You should be able to convince yourself that the point (x, y) is uniformly distributed in the area under the comparison function $f(x)$.

An equivalent procedure is to pick the second uniform deviate between zero and one, and accept or reject according to whether it is respectively less than or greater than the ratio $p(x)/f(x)$.

So, to summarize, the rejection method for some given $p(x)$ requires that one find, once and for all, some reasonably good comparison function $f(x)$. Thereafter, each deviate generated requires two uniform random deviates, one evaluation of f (to get the coordinate y), and one evaluation of p (to decide whether to accept or reject

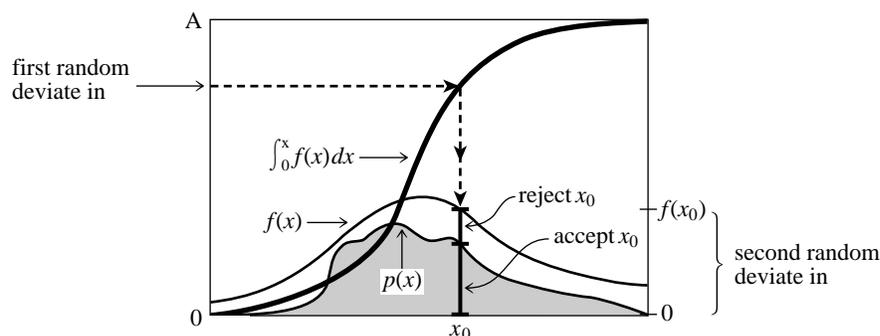


Figure 7.3.1. Rejection method for generating a random deviate x from a known probability distribution $p(x)$ that is everywhere less than some other function $f(x)$. The transformation method is first used to generate a random deviate x of the distribution f (compare Figure 7.2.1). A second uniform deviate is used to decide whether to accept or reject that x . If it is rejected, a new deviate of f is found; and so on. The ratio of accepted to rejected points is the ratio of the area under p to the area between p and f .

the point x, y). Figure 7.3.1 illustrates the procedure. Then, of course, this procedure must be repeated, on the average, A times before the final deviate is obtained.

Gamma Distribution

The gamma distribution of integer order $a > 0$ is the waiting time to the a th event in a Poisson random process of unit mean. For example, when $a = 1$, it is just the exponential distribution of §7.2, the waiting time to the first event.

A gamma deviate has probability $p_a(x)dx$ of occurring with a value between x and $x + dx$, where

$$p_a(x)dx = \frac{x^{a-1}e^{-x}}{\Gamma(a)}dx \quad x > 0 \quad (7.3.1)$$

To generate deviates of (7.3.1) for small values of a , it is best to add up a exponentially distributed waiting times, i.e., logarithms of uniform deviates. Since the sum of logarithms is the logarithm of the product, one really has only to generate the product of a uniform deviates, then take the log.

For larger values of a , the distribution (7.3.1) has a typically “bell-shaped” form, with a peak at $x = a$ and a half-width of about \sqrt{a} .

We will be interested in several probability distributions with this same qualitative form. A useful comparison function in such cases is derived from the *Lorentzian distribution*

$$p(y)dy = \frac{1}{\pi} \left(\frac{1}{1+y^2} \right) dy \quad (7.3.2)$$

whose inverse indefinite integral is just the tangent function. It follows that the x -coordinate of an area-uniform random point under the comparison function

$$f(x) = \frac{c_0}{1+(x-x_0)^2/a_0^2} \quad (7.3.3)$$

for any constants a_0, c_0 , and x_0 , can be generated by the prescription

$$x = a_0 \tan(\pi U) + x_0 \quad (7.3.4)$$

where U is a uniform deviate between 0 and 1. Thus, for some specific “bell-shaped” $p(x)$ probability distribution, we need only find constants a_0, c_0, x_0 , with the product $a_0 c_0$ (which determines the area) as small as possible, such that (7.3.3) is everywhere greater than $p(x)$.

Ahrens has done this for the gamma distribution, yielding the following algorithm (as described in Knuth [1]):

```

FUNCTION gamdev(ia,idum)
INTEGER ia,idum
REAL gamdev
C USES ran1
  Returns a deviate distributed as a gamma distribution of integer order ia, i.e., a waiting
  time to the iath event in a Poisson process of unit mean, using ran1(idum) as the source
  of uniform deviates.
INTEGER j
REAL am,e,s,v1,v2,x,y,ran1
if(ia.lt.1)pause 'bad argument in gamdev'
if(ia.lt.6)then          Use direct method, adding waiting times.
  x=1.
  do 11 j=1,ia
    x=x*ran1(idum)
  enddo 11
  x=-log(x)
else                    Use rejection method.
1  v1=ran1(idum)         These four lines generate the tangent of a random angle, i.e.,
  v2=2.*ran1(idum)-1.   are equivalent to y = tan(3.14159265 * ran1(idum)).
  if(v1**2+v2**2.gt.1.)goto 1
  y=v2/v1
  am=ia-1
  s=sqrt(2.*am+1.)
  x=s*y+am              We decide whether to reject x:
  if(x.le.0.)goto 1     Reject in region of zero probability.
  e=(1.+y**2)*exp(am*log(x/am)-s*y)   Ratio of prob. fn. to comparison fn.
  if(ran1(idum).gt.e)goto 1     Reject on basis of a second uniform de-
                                viate.
endif
gamdev=x
return
END

```

Poisson Deviates

The Poisson distribution is conceptually related to the gamma distribution. It gives the probability of a certain integer number m of unit rate Poisson random events occurring in a given interval of time x , while the gamma distribution was the probability of waiting time between x and $x + dx$ to the m th event. Note that m takes on only integer values ≥ 0 , so that the Poisson distribution, viewed as a continuous distribution function $p_x(m)dm$, is zero everywhere except where m is an integer ≥ 0 . At such places, it is infinite, such that the integrated probability over a region containing the integer is some finite number. The total probability at an integer j is

$$\text{Prob}(j) = \int_{j-\epsilon}^{j+\epsilon} p_x(m)dm = \frac{x^j e^{-x}}{j!} \quad (7.3.5)$$

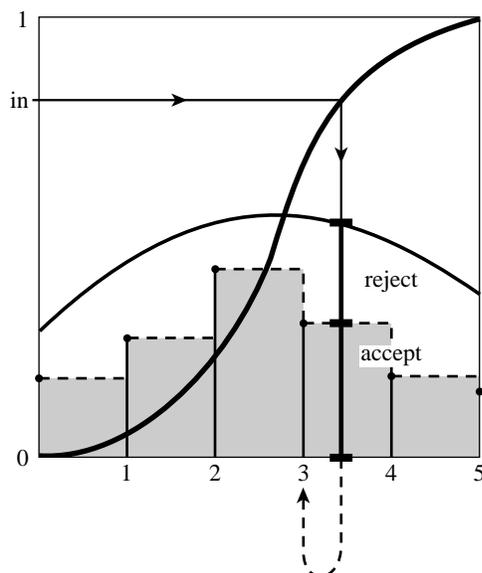


Figure 7.3.2. Rejection method as applied to an integer-valued distribution. The method is performed on the step function shown as a dashed line, yielding a real-valued deviate. This deviate is rounded down to the next lower integer, which is output.

At first sight this might seem an unlikely candidate distribution for the rejection method, since no continuous comparison function can be larger than the infinitely tall, but infinitely narrow, *Dirac delta functions* in $p_x(m)$. However, there is a trick that we can do: Spread the finite area in the spike at j uniformly into the interval between j and $j + 1$. This defines a continuous distribution $q_x(m)dm$ given by

$$q_x(m)dm = \frac{x^{[m]}e^{-x}}{[m]!} dm \quad (7.3.6)$$

where $[m]$ represents the largest integer less than m . If we now use the rejection method to generate a (noninteger) deviate from (7.3.6), and then take the integer part of that deviate, it will be as if drawn from the desired distribution (7.3.5). (See Figure 7.3.2.) This trick is general for any integer-valued probability distribution.

For x large enough, the distribution (7.3.6) is qualitatively bell-shaped (albeit with a bell made out of small, square steps), and we can use the same kind of Lorentzian comparison function as was already used above. For small x , we can generate independent exponential deviates (waiting times between events); when the sum of these first exceeds x , then the number of events that would have occurred in waiting time x becomes known and is one less than the number of terms in the sum.

These ideas produce the following routine:

```

FUNCTION poidev(xm,idum)
  INTEGER idum
  REAL poidev,xm,PI
  PARAMETER (PI=3.141592654)
C  USES gammln,ran1
  Returns as a floating-point number an integer value that is a random deviate drawn from a
  Poisson distribution of mean xm, using ran1(idum) as a source of uniform random deviates.

```

```

REAL alxm,em,g,oldm,sq,t,y,gammln,ran1
SAVE alxm,g,oldm,sq
DATA oldm /-1./
if (xm.lt.12.)then
  if (xm.ne.oldm) then
    oldm=xm
    g=exp(-xm)
  endif
  em=-1
  t=1.
  2 em=em+1.
  t=t*ran1(idum)
  if (t.gt.g) goto 2
else
  if (xm.ne.oldm) then
    oldm=xm
    sq=sqrt(2.*xm)
    alxm=log(xm)
    g=xm*alxm-gammln(xm+1.)
  1 endif
  y=tan(PI*ran1(idum))
  em=sq*y+alxm
  if (em.lt.0.) goto 1
  em=int(em)
  t=0.9*(1.+y**2)*exp(em*alxm-gammln(em+1.))-g
  if (ran1(idum).gt.t) goto 1
endif
poidev=em
return
END

```

Flag for whether xm has changed since last call.
Use direct method.
If xm is new, compute the exponential.
Instead of adding exponential deviates it is equivalent to multiply uniform deviates. We never actually have to take the log, merely compare to the pre-computed exponential.
Use rejection method.
If xm has changed since the last call, then precompute some functions that occur below.
The function gammln is the natural log of the gamma function, as given in §6.1.
y is a deviate from a Lorentzian comparison function. em is y, shifted and scaled.
Reject if in regime of zero probability.
The trick for integer-valued distributions.
The ratio of the desired distribution to the comparison function; we accept or reject by comparing it to another uniform deviate.
The factor 0.9 is chosen so that t never exceeds 1.

Binomial Deviates

If an event occurs with probability q , and we make n trials, then the number of times m that it occurs has the binomial distribution,

$$\int_{j-\epsilon}^{j+\epsilon} p_{n,q}(m) dm = \binom{n}{j} q^j (1-q)^{n-j} \quad (7.3.7)$$

The binomial distribution is integer valued, with m taking on possible values from 0 to n . It depends on *two* parameters, n and q , so is correspondingly a bit harder to implement than our previous examples. Nevertheless, the techniques already illustrated are sufficiently powerful to do the job:

```

FUNCTION bnldev(pp,n,idum)
INTEGER idum,n
REAL bnldev,pp,PI
  C USES gammln,ran1
PARAMETER (PI=3.141592654)
  Returns as a floating-point number an integer value that is a random deviate drawn from
  a binomial distribution of n trials each of probability pp, using ran1(idum) as a source
  of uniform random deviates.
INTEGER j,nold
REAL am,em,en,g,oldg,p,pc,pclog,plog,pold,sq,t,y,gammln,ran1

```

Sample page from NUMERICAL RECIPES IN FORTRAN 77: THE ART OF SCIENTIFIC COMPUTING (ISBN 0-521-43064-X)
 Copyright (C) 1986-1992 by Cambridge University Press. Programs Copyright (C) 1986-1992 by Numerical Recipes Software.
 Permission is granted for internet users to make one paper copy for their own personal use. Further reproduction, or any copying of machine-readable files (including this one), to any server computer, is strictly prohibited. To order Numerical Recipes books, diskettes, or CDROMs visit website <http://www.nr.com> or call 1-800-872-7423 (North America only), or send email to trade@cup.cam.ac.uk (outside North America).

```

SAVE nold,pold,pc,plog,pclog,en,oldg
DATA nold /-1/, pold /-1./ Arguments from previous calls.
if(pp.le.0.5)then           The binomial distribution is invariant under changing pp to
    p=pp                    1.-pp, if we also change the answer to n minus itself;
else                          we'll remember to do this below.
    p=1.-pp
endif
am=n*p                       This is the mean of the deviate to be produced.
if (n.lt.25)then            Use the direct method while n is not too large. This can
    bnldev=0.                require up to 25 calls to ran1.
    do 11 j=1,n
        if(ran1(idum).lt.p) bnldev=bnldev+1.
    enddo 11
else if (am.lt.1.) then     If fewer than one event is expected out of 25 or more tri-
    g=exp(-am)              als, then the distribution is quite accurately Poisson. Use
    t=1.                    direct Poisson method.
    do 12 j=0,n
        t=t*ran1(idum)
        if (t.lt.g) goto 1
    enddo 12
    j=n
    bnldev=j
1 else                       Use the rejection method.
    if (n.ne.nold) then     If n has changed, then compute useful quantities.
        en=n
        oldg=gammln(en+1.)
        nold=n
    endif
    if (p.ne.pold) then     If p has changed, then compute useful quantities.
        pc=1.-p
        plog=log(p)
        pclog=log(pc)
        pold=p
    endif
    sq=sqrt(2.*am*pc)       The following code should by now seem familiar: rejection
    y=tan(PI*ran1(idum))    method with a Lorentzian comparison function.
    em=sq*y+am
    if (em.lt.0..or.em.ge.en+1.) goto 2      Reject.
    em=int(em)                Trick for integer-valued distribution.
    t=1.2*sq*(1.+y**2)*exp(oldg-gammln(em+1.)
    * -gammln(en-em+1.)+em*plog+(en-em)*pclog)
    if (ran1(idum).gt.t) goto 2      Reject. This happens about 1.5 times per deviate, on
    bnldev=em                    average.
endif
if (p.ne.pp) bnldev=n-bnldev    Remember to undo the symmetry transformation.
return
END

```

See Devroye [2] and Bratley [3] for many additional algorithms.

CITED REFERENCES AND FURTHER READING:

- Knuth, D.E. 1981, *Seminumerical Algorithms*, 2nd ed., vol. 2 of *The Art of Computer Programming* (Reading, MA: Addison-Wesley), pp. 120ff. [1]
- Devroye, L. 1986, *Non-Uniform Random Variate Generation* (New York: Springer-Verlag), §X.4. [2]
- Bratley, P., Fox, B.L., and Schrage, E.L. 1983, *A Guide to Simulation* (New York: Springer-Verlag). [3].

7.4 Generation of Random Bits

This topic is not very useful for programming in high-level languages, but it can be quite useful when you have access to the machine-language level of a machine or when you are in a position to build special-purpose hardware out of readily available chips.

The problem is how to generate single random bits, with 0 and 1 equally probable. Of course you can just generate uniform random deviates between zero and one and use their high-order bit (i.e., test if they are greater than or less than 0.5). However this takes a lot of arithmetic; there are special-purpose applications, such as real-time signal processing, where you want to generate bits very much faster than that.

One method for generating random bits, with two variant implementations, is based on “primitive polynomials modulo 2.” The theory of these polynomials is beyond our scope (although §7.7 and §20.3 will give you small tastes of it). Here, suffice it to say that there are special polynomials among those whose coefficients are zero or one. An example is

$$x^{18} + x^5 + x^2 + x^1 + x^0 \quad (7.4.1)$$

which we can abbreviate by just writing the nonzero powers of x , e.g.,

$$(18, 5, 2, 1, 0)$$

Every primitive polynomial modulo 2 of order n (=18 above) defines a recurrence relation for obtaining a new random bit from the n preceding ones. The recurrence relation is guaranteed to produce a sequence of maximal length, i.e., cycle through all possible sequences of n bits (except all zeros) before it repeats. Therefore one can seed the sequence with any initial bit pattern (except all zeros), and get $2^n - 1$ random bits before the sequence repeats.

Let the bits be numbered from 1 (most recently generated) through n (generated n steps ago), and denoted a_1, a_2, \dots, a_n . We want to give a formula for a new bit a_0 . After generating a_0 we will shift all the bits by one, so that the old a_n is finally lost, and the new a_0 becomes a_1 . We then apply the formula again, and so on.

“Method I” is the easiest to implement in hardware, requiring only a single shift register n bits long and a few XOR (“exclusive or” or bit addition mod 2) gates. For the primitive polynomial given above, the recurrence formula is

$$a_0 = a_{18} \text{ XOR } a_5 \text{ XOR } a_2 \text{ XOR } a_1 \quad (7.4.2)$$

The terms that are XOR’d together can be thought of as “taps” on the shift register, XOR’d into the register’s input. More generally, there is precisely one term for each nonzero coefficient in the primitive polynomial except the constant (zero bit) term. So the first term will always be a_n for a primitive polynomial of degree n , while the last term might or might not be a_1 , depending on whether the primitive polynomial has a term in x^1 .

It is rather cumbersome to illustrate the method in FORTRAN. Assume that `iand` is a bitwise AND function, `not` is bitwise complement, `ishft(, 1)` is leftshift by one bit, `ior` is bitwise OR. (These are available in many FORTRAN implementations.) Then we have the following routine.

Sample page from NUMERICAL RECIPES IN FORTRAN 77: THE ART OF SCIENTIFIC COMPUTING (ISBN 0-521-43064-X)
 Copyright (C) 1986-1992 by Cambridge University Press. Programs Copyright (C) 1986-1992 by Numerical Recipes Software.
 Permission is granted for internet users to make one paper copy for their own personal use. Further reproduction, or any copying of machine-readable files (including this one), to any server computer, is strictly prohibited. To order Numerical Recipes books, diskettes, or CDROMs visit website <http://www.nr.com> or call 1-800-872-7423 (North America only), or send email to trade@cup.cam.ac.uk (outside North America).