1 Caution!

The name “negative binomial” is commonly used to refer to two different distributions, which may lead to confusion.

Both distributions have 2 parameters, p and k, and are related to the following experiment: perform independent Bernoulli trial with success probability p until exactly k successes have been observed.

In the first variation, we let X be the total number of trials until the k-th success. This leads to the following probability mass function:

$$p_X(n) = \binom{n-1}{k-1} p^k (1-p)^{n-k} \qquad \text{for } n \in \{k, k+1, \ldots\}$$

The mean of this distribution is:

$$\mu = \frac{k}{p}$$

In the second variation, we let X be the number of *failures* until the k-th success. This leads to the following probability mass function:

$$p_X(n) = \binom{n+k-1}{k-1} p^k (1-p)^{n} \qquad \text{for } n \in \{0, 1, \ldots\}$$

This is the formulation you will find for example in *mathworld* (http://mathworld.wolfram.com/NegativeBinomialDistribution.html), and is also the formulation used for the function `dnbinom()` in R.

The mean of this distribution is:

$$\mu = \frac{k(1-p)}{p}$$

Suppose we assume that the lengths X of a coding sequence has the the first negative binomial distribution, for a fixed k. I want to use the second formulation, and need to estimate the parameter p. I use the method of moments, noting that $X - k$ has the second distribution:

$$\bar{x} - k = \frac{k(1-p)}{p} \implies \hat{p} = \frac{k}{\bar{x}}$$

On the other hand, if I want to plot the first distribution, I simply need to shift the second distribution to the right (0 failures $\iff$ k trials, 1 failure $\iff$ $k+1$ trials, ...).