

Nested loops with R

A *loop* in a program is a block of code that gets *executed* many times. This block of code is the *body* of the loop. The body of a loop can itself include a loop. If it does, the whole thing is a *nested loop*, with the *inner loop* being *nested* inside the *outer loop*. This would be a *double loop* if the inner loop doesn't contain a loop. It is possible to have *triple loops*, with an inner-inner loop in the inner loop, in the outer loop, and so on.

The R program posted with this assignment starts with a double loop. The control statement for the outer loop is line 4: `for (i in 1:n){ ...`. Lines 5 through 10 are the body of the outer loop. Each of them is indented at least 4 spaces, which allows you to see the loop body easily. Line 6 is the control statement for the inner loop. Line 7 is the body of the inner loop. Each line of the inner loop should be indented at least 4 more spaces, to make the inner loop easy to distinguish from the outer loop. Line 9 is the curly brace that closes the inner loop. It is indented 2 spaces in from the outer loop, for a total of 6. This makes it stand out both from the outer loop and the inner loop body. Line 11 is the curly that ends the outer loop. It is indented 2 spaces, which makes it stand out from the body of the outer loop and from the rest of the code.

The programming language R does not force you to indent loops, and you don't have to be exactly like this. You can find lots of examples of un-indented R code on the web, even in R tutorials and examples, sadly. But good programmers (and students hoping to get full credit for their codes in Mathematics of Finance) indent loops in some systematic way. These are the programmers who spend 10 minutes more time typing the code then spend 10 hours less time debugging. And these are the programmers you want on your team, because their code is *readable*. If you write *good code*, not just working code, people will hire you on their team. The Resources page on the class web site has more on programming style.

Nested loops are common in most codes. They come up, for example, in matrix operations. Suppose A and B are $n \times n$ matrices with elements a_{ij} and b_{ij} respectively. The product is $C = AB$. Its entries are

$$c_{ij} = \sum_{k=1}^n a_{ik}b_{kj}.$$

This formula can be implemented in R (or most other programming languages) using a triple loop. The formula holds for $i = 1, 2, \dots, n$ (the outer loop), and for $j = 1, 2, \dots, n$ (the middle loop, inside the outer loop and outside the inner-inner loop). The outer loop is the loop over i , and the middle loop is the loop over j . The sum is evaluated in an inner-inner loop over k . In R, the loop structure looks like:

A triple loop, indented four spaces each level

```
for ( i in 1:n){
  -----
  --- some code ---      # outer loop body
  -----
  for ( j in 1:n){
    -----
    --- some code ---    # middle loop body
    -----
    for (k in 1:n){
      -----
      --- some code --- # inner loop body
      -----
    }                    # end of the inner loop over k
    -----
    --- some code ---    # middle loop body
    -----              # aligned with the first middle loop code
  }                      # end of the middle loop over j
  -----
  --- some code ---      # outer loop body
  -----                # aligned with the first outer loop code
}                        # end of the outer loop over i
```