Simulating Signal Detection Theory
Psychology 465A
Dr. Lawrence Ward

1. Make the usual assumptions (see previous handout). Note that sensory processing of the stimulus is modeled by sampling from N or SN distributions.

2. You are simulating both an experimenter and a subject.
    2.1 Experimenter: selects and presents stimuli and records responses, tabulates indices, calculates measures
    2.2 Subject: processes stimulus, makes responses according to decision rule

3. Below is a flow diagram of the major processes. Note that part of the "Sample from…" box and the "Samp>tisubc?" decision diamond are subject processes, the remainder are experimenter processes. This is a very simple simulation indeed.

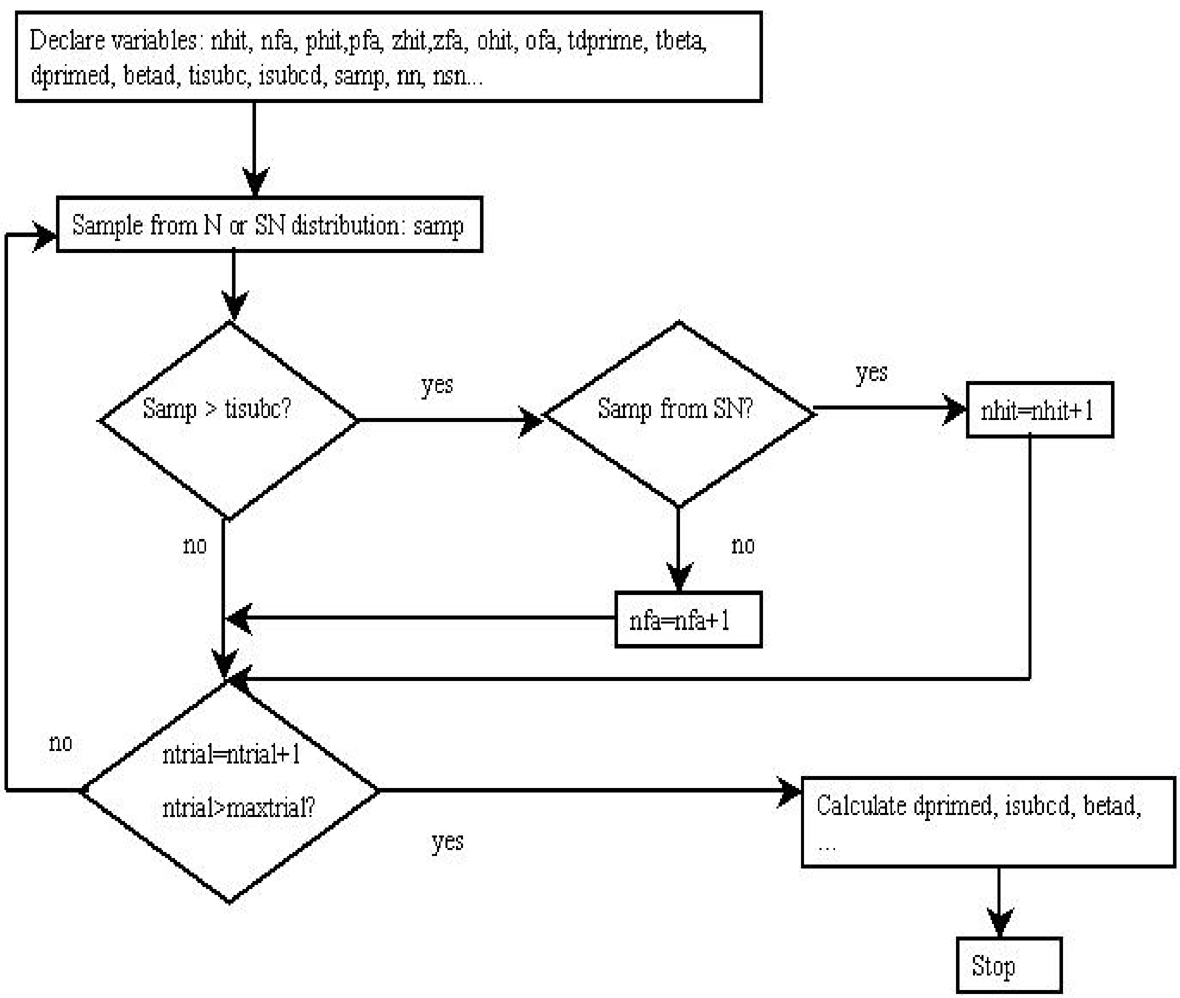

4. The assignment

4.1 Verify already known results: verify that when theoretical *d'* changes, *d'* calculated from the simulation data changes and moreover converges on the theoretical *d',* and that when theoretical β changes the calculated β changes appropriately but the change has no effect on calculated *d'*. To do this you need at least a "canonical" run using, say, *d'*=1 and β=1. Then to verify the *d'* change, do a run with *d'*=2 and β=1, say, and to verify the non-effect of β change do a run with *d'*=1 and β=3, say (you choose the values).

4.2 Derive a prediction: change one of the assumptions of the model and work out the consequences using one or more additional runs (depends on the change how many you would need).

5. Simulation strategies

5.1 "The long run": do one or a few very long runs in each "condition" because by the law of large numbers the results should converge to some asymptote. What is a very long run? Decide by recalculating the relevant measure (e.g., *d'*) every 1000 trials, say, until it doesn't seem to be changing any more. This has the advantage that a plot of the measure versus the number of trials in the run is evidence for the sought-after convergence. It has the disadvantage that your one run could be a fluke.

5.2 "Intro stats": you don't trust the law of large numbers because there can be idiosyncratic runs. Thus you do several (say 100) shorter runs (say 1000 trials) and perform ordinary statistical tests on the results. For *d'*, e.g., you could calculate *d'* for each of 100, 1000-trial runs, and take the mean and sd over this set. Then, using a *z*-test you could compare the mean *d'* for one run with the mean for another and draw the indicated conclusions. For fewer runs you could do a *t*-test. This has the disadvantage of requiring more work (could be automated) and also that the conclusion is still statistical, with some nonzero probability of being wrong.