

Spreadsheet signal detection

ROBERT D. SORKIN
University of Florida, Gainesville, Florida

A number of studies in perception, attention, and memory employ signal detection theory (SDT) to assess the accuracy of an observer's detection or discrimination performance. Some of the problems that students have with understanding and using SDT are associated with the calculations needed to obtain SDT parameters and predictions. All of these calculations, plus the simulation of SDT processes, can be performed using a spreadsheet application program, such as Excel or Quattro Pro. This paper offers a short tutorial on how to use a spreadsheet program to increase your students' knowledge and understanding of SDT.

Signal detection theory (SDT) is a psychophysical theory that allows measurement of a subject's ability to detect or discriminate sensory stimuli. SDT emerged from the application of statistical decision theory to the problem of detecting weak signals in noise. The first application of the theory to problems in sensory psychophysics was by Tanner and Swets (1954). An important aspect of the theory is its ability to separate the effects of sensory from nonsensory factors. Sensory factors include the intensity and duration of the stimuli; nonsensory factors include motivational aspects, such as the observer's expectancies about the possible stimuli and the costs and benefits of the possible responses. The theory enables one to assess these two distinct aspects of behavior: the observer's sensory sensitivity (d') and the observer's bias (β or c) for choosing one response over another.

SDT is an important component in the cognitive psychology toolbox. SDT's influence is shown by the large number of studies in perception, attention, and memory that employ the d' index as a measure of the accuracy of an observer's detection or discrimination performance.¹ A basic treatment of SDT can usually be found in textbooks in sensation and perception (e.g., Goldstein, 1996; Schiffman, 1996), in human factors psychology (Kantowitz & Sorkin, 1983; Proctor & Van Zandt, 1994; Wickens, 1992), sometimes in cognitive psychology (Haberlandt, 1994; Kellogg, 1995; Solso, 1998), and occasionally in experimental methods (Levine & Parkinson, 1994; Snodgrass, Levy-Berger, & Haydon, 1985).

Because SDT is based on the application of statistical decision theory to psychophysics, it is basically a mathematical theory. As a result, many students (and some instructors) have problems understanding (or teaching) it. Some of these problems are associated with the computations that are required to obtain the SDT model's parameters and predictions. Powerful mathematical programs,

such as MATLAB, MATHCAD, and Mathematica, are available for performing these functions, but most psychology students are not familiar with such programs. On the other hand, these calculations can be performed using commonly available spreadsheet application programs, such as Excel or Quattro Pro.² This paper offers a short tutorial on how to use a spreadsheet to (1) calculate the behavior of the SDT model given assumed parameter values, (2) obtain an observer's parameters from obtained data, and (3) simulate the behavior of hypothetical signal detection systems.³ Since moving to a spreadsheet as an instructional tool, I have found students to be more enthusiastic about learning and doing signal detection.⁴

USING A SPREADSHEET TO EXPLAIN THE BASIC THEORY

The basic idea of SDT is that the observer (or listener, decision-maker, etc.) receives an input and then must make a decision about which one of two events—the signal-plus-noise event or the noise-alone event—gave rise to that input. The input may be a multidimensional stimulus, (x_1, x_2, \dots, x_k) , but the observer's perceptual system somehow maps the stimulus input to a single number, z , on a one-dimensional decision axis. After the observer has transformed the input to a number on the decision axis, that number is compared with a criterion, z_c . If z is greater than or equal to z_c , the observer responds "yes," indicating acceptance of the signal-plus-noise hypothesis. If z is less than z_c , the observer responds "no."

This process is illustrated in Figure 1. The observer calculates the statistic z using information about the signal and noise characteristics stored in memory and knowledge of when and where the signal may occur.⁵ The specific setting for the response criterion, z_c , depends on the prior probabilities of signal-plus-noise and noise-alone and on the payoffs and penalties associated with the different trial outcomes (hits, misses, false alarms, and correct rejections).⁶

It is clear that the properties of the signal-plus-noise and noise-alone distributions on z will determine how well the observer can discriminate between these two

The author thanks Stephen Link and an anonymous reviewer for their excellent suggestions for improving an earlier draft of the article. Correspondence should be addressed to R. D. Sorkin, Department of Psychology, P.O. Box 112250, University of Florida, Gainesville, FL 32611 (e-mail: sorkin@psych.ufl.edu).

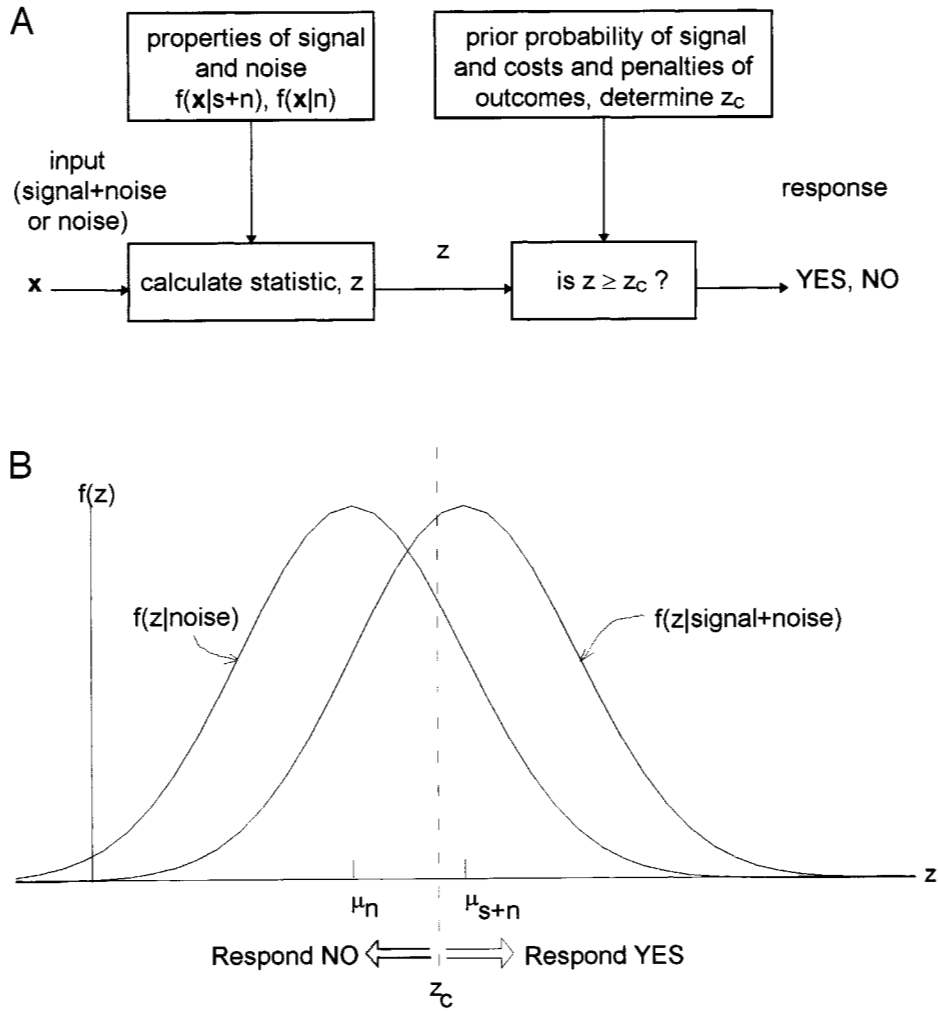


Figure 1. The basic signal detection process. (A) A block diagram of the signal detection system. (B) The probability distributions of the decision statistic, z, given noise-alone and signal-plus-noise (see text).

events. If the means of the distributions are near each other, the probability of an error will be high, and, thus, performance accuracy will be low. The usual SDT assumption is that the distributions are equal-variance, normal distributions. If so, performance accuracy will be completely determined by the normalized distance between the means of the distributions:

$$d' = (\mu_{s+n} - \mu_n) / \sigma. \tag{1}$$

The parameter d' defines the observer's index of sensitivity or discriminability. The value for d' can vary between 0 at chance performance and above 4 when discrimination performance is essentially errorless.

The observer's bias toward making a "yes" response has been defined in several different ways. Two common measures are β and c . β is the ratio of the ordinates of the signal and noise distributions at the decision criterion, z_c ; that is, $\beta = f(z_c | s + n) / f(z_c | n)$. The other bias measure is c , the

relative position of the decision criterion, z_c , on the decision axis (Macmillan & Creelman, 1991). The bias parameters, β and c , are related to z_c and d' in the following way:

$$\beta = e^{[z_c d' - (d')^2 / 2]} \tag{2a}$$

and

$$c = z_c - d' / 2. \tag{2b}$$

Values for $\beta = 1$ or $c = 0$ indicate a zero or neutral bias; values of $\beta > 1$ or $c > 0$ indicate a conservative bias (a bias against responding "yes").

The hit probability, $p(\text{hit})$, is the probability that the observer responds "yes" given signal-plus-noise; the false-alarm probability, $p(\text{fa})$, is the probability that the observer responds "yes" given noise (see Figure 1). The hit and false-alarm probabilities that will be generated by an observer who has a specific d' and decision criterion, z_c , are given by the following equations:

$$p(\text{hit}) = p(\text{"yes"} | \text{signal} + \text{noise}) = \int_{z_c}^{\infty} f(z | s + n) dz$$

and

$$p(\text{fa}) = p(\text{"yes"} | \text{noise}) = \int_{z_c}^{\infty} f(z | n) dz.$$

For simplicity, let the mean of the noise distribution, μ_n , be equal to zero, and let the variance of both distributions be 1. Then, the mean of the signal-plus-noise distribution, μ_{s+n} , will be equal to d' . Furthermore, let the distribution $f(z | n)$ be designated by the expression, $f(z)$. Then, the hit and false-alarm probabilities (using spreadsheet terminology in Equations 3b and 4b) are

$$\begin{aligned} p(\text{hit}) &= \int_{z_c}^{\infty} f[z | s + n] dz = \int_{z_c}^{\infty} f[(z - d') | n] dz \\ &= 1 - \int_{-\infty}^{z_c} f(z - d') dz \end{aligned} \tag{3a}$$

$$= 1 - \text{NORMDIST}(z_c, d', 1, 1), \tag{3b}$$

where $\text{NORMDIST}(z_c, d', 1, 1)$ is the value of the area between $-\infty$ and z_c for a normal probability distribution with mean d' and variance $\sigma^2 = 1$.

$$\begin{aligned} p(\text{fa}) &= \int_{z_c}^{\infty} f(z | n) dz = \int_{z_c}^{\infty} f(z) dz \\ &= 1 - \int_{-\infty}^{z_c} f(z) dz \end{aligned} \tag{4a}$$

$$= 1 - \text{NORMDIST}(z_c, 0, 1, 1), \tag{4b}$$

where $\text{NORMDIST}(z_c, 0, 1, 1)$ is the value of the area between $-\infty$ and z_c for a normal probability distribution with mean 0 and variance $\sigma^2 = 1$. (Equations 3b and 4b give the desired probabilities using Excel or Quattro Pro spreadsheet terminology. The second and third terms in the NORMDIST argument are, respectively, the distribution's mean and standard deviation; the last term indicates whether the density function or the cumulative function is desired.⁷)

It is often the case that we know $p(\text{hit})$ and $p(\text{fa})$, but we wish to calculate the values of d' and z_c that gave rise to those probabilities. Students often have difficulty with these calculations. The spreadsheet simplifies the calculations because it provides a function for evaluating the inverse of the cumulative normal distribution. Let p be the area from $-\infty$ to z^* under the standard normal distribution, $f(z)$. That is,

$$p = \int_{-\infty}^{z^*} f(z) dz.$$

Knowing p , the value of z^* can be computed using the spreadsheet's NORMSINV function:

$$z^* = \text{NORMSINV}(p). \tag{5}$$

For example, the unknown criterion value can be determined from the false-alarm probability by

$$z_c = \text{NORMSINV}(1 - p(\text{fa})). \tag{6}$$

The distribution representing signal + noise is identical to the distribution for noise except for a shift of amount d' on the z -axis. This shift can be determined by computing

$$\text{NORMSINV}(p(\text{hit})) - \text{NORMSINV}(p(\text{fa})) = d'. \tag{7}$$

Also,

$$-(\text{NORMSINV}(p(\text{hit})) + \text{NORMSINV}(p(\text{fa}))) / 2 = c. \tag{8}$$

The value of β is determined by

$$\text{NORMDIST}(z_c, d', 1, 0) / \text{NORMDIST}(z_c, 0, 1, 0) = \beta. \tag{9}$$

RECEIVER OPERATING CHARACTERISTIC CURVES

If d' is fixed and the observer is induced to shift criterion to a value that is more conservative or more liberal than the initial value, a new hit rate and false-alarm rate will result. The set of all the possible hit and false alarms that would result from this procedure defines a curve called the *receiver operating characteristic* (ROC). This is shown in Figure 2 for several different d' values. The higher the d' , the nearer the resulting hit and false alarm values are to the upper left corner of the plot.

It is instructive for the student to try different values for d' and observe the effect on the obtained ROC curves or, alternatively, to fix the d' and observe the shift in operating point as c is made more conservative or liberal.




Figure 2. Receiver operating characteristic (ROC) curves for detection systems with d' values of 1, 2, and 3, computed by the spreadsheet program discussed in the text.

Either of these “experiments” may be performed simply by changing the appropriate cell in the spreadsheet and then observing the spreadsheet graph. Equations 3 and 4 were used to implement these functions with a spreadsheet. Figure 3A shows a subset of the spreadsheet cells used in this example. Figure 3B portrays the same cells but shows the spreadsheet formulas needed to perform the spreadsheet calculations. Column A contains the values of z_c used to calculate all the values of $f(z)$ and the integral of $f(z)$. Columns B and C are the ordinates of the signal and noise distributions. These values were used to plot the hypothesis distributions shown in Figure 1. Column D contains the false-alarm probabilities corresponding to the z_c values in column A. Columns E, F, and G contain the hit probabilities calculated for the z_c values in column A and the d' value in cell 1 of their respective columns. Thus, column D contains the false-alarm rates that are common to all the ROC curves, whereas columns E, F, and G contain the hit rates produced by d' values of 1, 2, and 3, respectively. To plot multiple ROC curves, one directs the program to plot the values in columns E, F, and G as the Y values and the values in column D as the X values. You can change the d' value (i.e., cell E1) and observe the immediate change in the plotted ROC curve.

CALCULATING PARAMETERS FROM A TYPICAL EXPERIMENT

In the next example, we assume that a detection experiment has been performed and that the resulting data have been successfully imported into a spreadsheet file in the form shown in Figure 4. We then use the spreadsheet (and Equations 6–9) to automate the process of analyzing the obtained data.

In this example, there was one block of 20 trials, containing 11 signal events (sig) and 9 noise (noi) events. The

participant responded either “yes” or “no” on each trial. We can ignore the block and parameter columns, since these were fixed throughout the 20-trial block. There are many ways to sort the data from such an experiment, but a simple technique is to use the spreadsheet IF function. The IF mathematical function performs logical tests on the entries of different cells. Thus, you can test for all the specific conditions of interest in an experiment and place the result in a cell in a new column. In the example, columns F and G were used to calculate whether the stimulus on each trial was a signal or a noise, and columns H and I were used to indicate whether the outcome was a hit or a false alarm. These events were then summed, the hit and false-alarm probabilities calculated, and the SDT parameters computed. The italicized headings in Figure 4 indicate calculated values. Figure 4B shows the formulas used for the calculations. Of course, in a “real” experiment, we would like to have many more than 20 trials.⁸

SIMULATING SDT PROCESSES

The spreadsheet program can also be used to simulate signal detection processes by virtue of its ability to generate random numbers according to specified probability distributions. For example, the random number generating function could generate 25 “noise” and 25 “signal” inputs to an observer. We could then have the observer behave according to Figure 1, by comparing each of those “inputs” to a criterion value and then making a “yes” or “no” decision. The resulting “hits” and “false alarms” could be aggregated, and we could analyze the resulting “data” using the spreadsheet procedure described in the last section. If we repeated this entire process 200 times, we would obtain a frequency distribution of obtained d' values. That is, we would have obtained 200 estimates of d' , each based on a sample of 50 trials. The variance of

A

	A	B	C	D	E	F	G
1	dp		2		1	2	3
2	zc	f(z n)	f(z s+n)	fa's	hits	hits	hits
3	-3.0	0.004	0.000	0.998	1.000	1.000	1.000
4	-2.8	0.008	0.000	0.997	1.000	1.000	1.000
5	-2.6	0.014	0.000	0.995	1.000	1.000	1.000
6	-2.4	0.022	0.000	0.991	0.999	1.000	1.000

B

	A	B	C	D	E
1	dp		2		1
2	zc	f(z n)	f(z s+n)	fa's	hits
3	-3	=NORMDIST(A3,0,1,0)	=NORMDIST(A3,\$C\$1,1,0)	=1-NORMDIST(\$A3,0,1,1)	=1-NORMDIST(\$A3,E\$1,1,1)
4	=A3+0.2	=NORMDIST(A4,0,1,0)	=NORMDIST(A4,\$C\$1,1,0)	=1-NORMDIST(\$A4,0,1,1)	=1-NORMDIST(\$A4,E\$1,1,1)
5	=A4+0.2	=NORMDIST(A5,0,1,0)	=NORMDIST(A5,\$C\$1,1,0)	=1-NORMDIST(\$A5,0,1,1)	=1-NORMDIST(\$A5,E\$1,1,1)
6	=A5+0.2	=NORMDIST(A6,0,1,0)	=NORMDIST(A6,\$C\$1,1,0)	=1-NORMDIST(\$A6,0,1,1)	=1-NORMDIST(\$A6,E\$1,1,1)
7	=A6+0.2	=NORMDIST(A7,0,1,0)	=NORMDIST(A7,\$C\$1,1,0)	=1-NORMDIST(\$A7,0,1,1)	=1-NORMDIST(\$A7,E\$1,1,1)

Figure 3. (A) A portion of the first Excel spreadsheet showing the ordinates of the hypothesis distributions and the hit and false alarm areas for d' 's of 1, 2, and 3 (see text). (B) A portion of the Excel spreadsheet showing the spreadsheet formulas for the ordinates of the hypothesis distributions and the hit and false-alarm areas (see text).

A

	A	B	C	D	E	F	G	H	I
1	block	parameter	trial	stim	resp	sigs	nois	hits	fas
2	1	330	1	sig	yes	1	0	1	0
3	1	330	2	sig	no	1	0	0	0
4	1	330	3	noi	no	0	1	0	0
5	1	330	4	noi	yes	0	1	0	1
6	1	330	5	sig	yes	1	0	1	0
7	1	330	6	noi	yes	0	1	0	1
8	1	330	7	sig	yes	1	0	1	0
9	1	330	8	noi	yes	0	1	0	1
10	1	330	9	sig	yes	1	0	1	0
11	1	330	10	sig	yes	1	0	1	0
12	1	330	11	noi	no	0	1	0	0
13	1	330	12	noi	no	0	1	0	0
14	1	330	13	sig	no	1	0	0	0
15	1	330	14	sig	yes	1	0	1	0
16	1	330	15	noi	no	0	1	0	0
17	1	330	16	sig	yes	1	0	1	0
18	1	330	17	sig	yes	1	0	1	0
19	1	330	18	noi	no	0	1	0	0
20	1	330	19	sig	yes	1	0	1	0
21	1	330	20	noi	no	0	1	0	0
22	count					11	9	9	3
23									
24					p(hit)			0.818	
25					p(fa)			0.333	
26					zc			0.431	
27					d'			1.339	
28					c			-0.24	
29					beta			0.726	

B

	E	F	G	H	I
1	resp	sigs	nois	hits	fas
2	yes	=IF(D2="sig",1,0)	=IF(D2="noi",1,0)	=IF(AND(D2="sig",E2="yes"),1,0)	=IF(AND(D2="noi",E2="yes"),1,0)
3	no	=IF(D3="sig",1,0)	=IF(D3="noi",1,0)	=IF(AND(D3="sig",E3="yes"),1,0)	=IF(AND(D3="noi",E3="yes"),1,0)
19	no	=IF(D19="sig",1,0)	=IF(D19="noi",1,0)	=IF(AND(D19="sig",E19="yes"),1,0)	=IF(AND(D19="noi",E19="yes"),1,0)
20	yes	=IF(D20="sig",1,0)	=IF(D20="noi",1,0)	=IF(AND(D20="sig",E20="yes"),1,0)	=IF(AND(D20="noi",E20="yes"),1,0)
21	no	=IF(D21="sig",1,0)	=IF(D21="noi",1,0)	=IF(AND(D21="sig",E21="yes"),1,0)	=IF(AND(D21="noi",E21="yes"),1,0)
22		=SUM(F2:F21)	=SUM(G2:G21)	=SUM(H2:H21)	=SUM(I2:I21)
23					
24	p(hit)			=H22/F22	
25	p(fa)			=I22/G22	
26	zc			=NORMSINV(1-H25)	
27	d'			=NORMSINV(H24)-NORMSINV(H25)	
28	c			=-0.5*(NORMSINV(H24)+NORMSINV(H25))	
29	beta			=EXP((-0.5*(((NORMSINV(H24))^2-(NORMSINV(H25))^2)))	

Figure 4. (A) A portion of the Excel spreadsheet showing 20 trials of data from a detection experiment and the calculated detection parameters. (B) The same spreadsheet as in Figure 4A, with the formulas for computing the detection parameters shown.

this sampling distribution of estimated d' would depend on the sample size (e.g., how many signal and noise “inputs” had been generated for each “experiment”). The more trials per experiment, the smaller the variability in the estimated value of d' . It would be useful for students to go through this process in order to get a feeling for how the sampling distribution of d' changes with the number of trials and the observer’s detection parameters.

There often is a practical reason for performing a simulation similar to the one just described: because it is al-

most inevitable that some blocks of trials in an experiment will result in either zero false alarms or 100% hits. Since the SDT model cannot handle such extreme hit or false-alarm rates, the question is, What does one do with such results? The common convention is to apply an arbitrary correction procedure to the miscreant trial blocks. For example, we could redo the problematic experimental condition, or we could simply assign a false-alarm rate of .0001 or a hit rate of .9999 to that condition. Either procedure will have an effect on our estimate of d' , and we

would like to know what that effect will be. If the correction biases the estimate of d' , we would like to have some idea about the direction and magnitude of that bias and how it depends on the sample size and the true values of d' and criterion. Miller (1996) considered the consequences of several of these correction procedures on the mean and variance of the estimated d' . He pointed out that, while the distribution of sample d' 's may be mathe-

matically intractable, it can be determined easily by a small computer performing a Monte Carlo simulation of the process. The last example in this tutorial uses a spreadsheet simulation to observe some of the properties of such an arbitrary correction procedure.

We consider the following correction procedure (Anscombe, 1956): If there are n no-signal trials and zero "yes" responses, we will set the false-alarm rate equal to

A

	A	B	C	D	E	F
1	d'		2.108			
2	2					
3	zc					
4	1.5					
5	hitp	hitp	hitp	hitp	hitp	hitp
6	0.574	0.684	0.729	0.646	0.588	0.676
7	fap	fap	fap	fap	fap	fap
8	0.140	0.080	0.100	0.240	0.160	0.160
9	chitp	chitp	chitp	chitp	chitp	chitp
10	0.574	0.684	0.729	0.646	0.588	0.676
11	cfap	cfap	cfap	cfap	cfap	cfap
12	0.140	0.080	0.100	0.240	0.160	0.160
13	pred'	pred'	pred'	pred'	pred'	pred'
14	1.268	1.885	1.892	1.081	1.217	1.450
15	hit/fa	hit/fa	hit/fa	hit/fa	hit/fa	hit/fa
16	1	1	1	1	1	1
17	1	1	1	1	1	1
18	1	1	1	1	1	1
19	1	1	1	1	1	1
20	1	1	1	1	1	1
21	1	1	1	1	1	1
22	1	1	0	1	0	1
23	0	1	1	0	0	0
24	1	0	0	1	1	1

B

	A	B
1	d'	
2	2	
3	zc	
4	1.5	
5	hitp	hitp
6	=SUM(Sheet1!A16:A40)/25	=SUM(Sheet1!B16:B40)/25
7	fap	fap
8	=SUM(Sheet1!A41:A65)/25	=SUM(Sheet1!B41:B65)/25
9	chitp	chitp
10	=IF(A6=1,0.9999,A6)	=IF(B6=1,0.9999,B6)
11	cfap	cfap
12	=IF(A8=0,0.0001,A8)	=IF(B8=0,0.0001,B8)
13	pred'	pred'
14	=NORMSINV(A10)-NORMSINV(A12)	=NORMSINV(B10)-NORMSINV(B12)
15	hit/fa	hit/fa
16	=IF(Sheet2!A16>=Sheet1!\$A\$4-Sheet1!\$A\$2,1,0)	=IF(Sheet2!B16>=Sheet1!\$A\$4-Sheet1!\$A\$2,1,0)
17	=IF(Sheet2!A17>=Sheet1!\$A\$4-Sheet1!\$A\$2,1,0)	=IF(Sheet2!B17>=Sheet1!\$A\$4-Sheet1!\$A\$2,1,0)
39	=IF(Sheet2!A39>=Sheet1!\$A\$4-Sheet1!\$A\$2,1,0)	=IF(Sheet2!B39>=Sheet1!\$A\$4-Sheet1!\$A\$2,1,0)
40	=IF(Sheet2!A40>=Sheet1!\$A\$4-Sheet1!\$A\$2,1,0)	=IF(Sheet2!B40>=Sheet1!\$A\$4-Sheet1!\$A\$2,1,0)
41	=IF(Sheet2!A41>=Sheet1!\$A\$4,1,0)	=IF(Sheet2!B41>=Sheet1!\$A\$4,1,0)
42	=IF(Sheet2!A42>=Sheet1!\$A\$4,1,0)	=IF(Sheet2!B42>=Sheet1!\$A\$4,1,0)
64	=IF(Sheet2!A64>=Sheet1!\$A\$4,1,0)	=IF(Sheet2!B64>=Sheet1!\$A\$4,1,0)
65	=IF(Sheet2!A65>=Sheet1!\$A\$4,1,0)	=IF(Sheet2!B65>=Sheet1!\$A\$4,1,0)

Figure 5. (A) A portion of the spreadsheet (Sheet 1) used for the simulation demonstration. (B) A portion of the simulation spreadsheet (Sheet 1) showing the formulas used to make the calculations of hit rate, false alarm rate, d' , and so on (see text).

A

	B	D	E
1	1.268	0	0
2	1.885	0.2	0
3	1.892	0.4	0
4	1.081	0.6	0
5	1.217	0.8	0
6	1.450	1.0	0
7	2.197	1.2	4
8	2.320	1.4	3
9	1.910	1.6	23
10	1.947	1.8	22
11	2.177	2.0	34
12	1.688	2.2	40
13	2.400	2.4	30
14	2.400	2.6	20
15	2.161	2.8	5
16	1.836	3.0	7
17	1.836	3.2	8
18	2.580	3.4	0
19	2.169	3.6	3
20	1.688	3.8	0
21	1.989	4	0

B

	B	D	E
1	{=TRANSPOSE(Sheet1!A14:GR14)}	0	{=FREQUENCY(B1:B199,D1:D21)}
2	{=TRANSPOSE(Sheet1!A14:GR14)}	=D1+0.2	{=FREQUENCY(B1:B199,D1:D21)}
3	{=TRANSPOSE(Sheet1!A14:GR14)}	=D2+0.2	{=FREQUENCY(B1:B199,D1:D21)}
4	{=TRANSPOSE(Sheet1!A14:GR14)}	=D3+0.2	{=FREQUENCY(B1:B199,D1:D21)}
5	{=TRANSPOSE(Sheet1!A14:GR14)}	=D4+0.2	{=FREQUENCY(B1:B199,D1:D21)}
6	{=TRANSPOSE(Sheet1!A14:GR14)}	=D5+0.2	{=FREQUENCY(B1:B199,D1:D21)}

Figure 6. (A) A portion of the simulation spreadsheet (Sheet 2) showing the results of transposing the (row of) estimated d' values from Sheet 1 and their frequency distribution in columns F and G (see text). (B) A portion of the simulation spreadsheet (Sheet 2) showing the formulas needed to transpose the (row of) estimated d' values and calculate their frequency distribution.

$1/n$. If there are n signal trials and n “yes” responses, we will set the hit rate to be $(n - 1)/n$. We are interested in the following question: For sample sizes of 50 and d' 's < 3 , does this procedure produce a large bias in the estimated d' ?

Figures 5 and 6 illustrate portions of the spreadsheet formulas needed to demonstrate the simulation process. Obviously, we can show only a portion of the entire spreadsheet in the figures. Sheet 2 of the worksheet is not shown in the figures and is used only for the random numbers. We used the normal random number generator to fill the columns of Sheet 2 with 50 random numbers in each column (from row 16 to row 115) and across 200 columns from column A through column GR. So, our sample size is 50, and we have taken 200 Monte Carlo “runs.”

Sheet 1 contains the important calculations for the simulation, shown in Figures 5 and 6. Each column of rows 16–65 contains the outcome of a “yes” or a “no” trial. We consider the entries in the first 25 rows, from row 16 through row 40, of each column on Sheet 2 (and the corresponding columns on Sheet 1) as “signal” trials. For each of the observations in the cells on Sheet 2, we test whether that cell value (plus d') exceeded the criterion. The value of d' is contained in cell A2 of Sheet 1,

and the criterion is contained in cell A4. If the observation on Sheet 2 exceeds the criterion, a 1 is put in the corresponding cell on Sheet 1; if not, a 0 is put in the corresponding cell on Sheet 1. The process is the same for rows 41–65, except that these are considered “nonsignal” trials so that we do not add d' to the value of the normal sample. The number of 1s in rows 16–40 then indicates the number of hits, and the number of 1s in rows 41–65 indicates the number of false alarms.

The hit and false-alarm rates are calculated in cells 6 and 8, respectively, of each column of Sheet 1. If the resulting hit rate is equal to 1 or the false-alarm rate is equal to 0, the rates are corrected to rates of .96 and .04, respectively, in cells 10 and 12. These rates are then used to calculate the estimated d' for that sample of 50 trials, placing the result in a cell in row 14. The average of the 200 d' estimates is calculated in cell C1 of Sheet 1 and is equal to 2.108. The standard deviation of the estimates is calculated in cell C3 and is equal to 0.487.

We can have the spreadsheet calculate and then plot the frequency distribution of our 200 estimates of d' , but we first need to transpose the row vector of 200 estimates to a column vector. This is done in column A of Sheet 3, shown in Figure 6. Columns D and E, respectively, of




Figure 7. Frequency distribution of the d' values obtained using the simulation spreadsheet, sample sizes of 50 trials and a value of $d'_{\text{true}} = 2$. The mean and standard deviation of d'_{estimate} are 2.108 and 0.487, respectively.

Sheet 3 are used for the category bins of the frequency distribution and the resulting count of d' estimates (in each category). Figure 7 is a plot of columns D and E of Sheet 3 (i.e., the frequency distribution of the 200 d' estimates).

Trying different values of d' and criterion (in cells A2 and A4, respectively, of Sheet 1) reveals some interesting things about the distribution of estimated d' . We set the criterion to be halfway between the noise and signal distributions (by setting cell A4 equal to $\$A\$2/2$) and evaluated different values of d' . Table 1 shows the results of this informal analysis. There is a small positive bias in the estimated d' , and this bias gets larger as the true d' increases to a value of 2.5. For true d' 's of less than 3, the maximum bias is less than 0.2. If the sample size is increased to 100, the bias decreases to a worst-case maximum of 0.06. These values are approximately consistent with the more complete analysis performed by Miller (1996).

SUMMARY

We have provided simple examples of how a spreadsheet application program, such as Excel or Quattro Pro, can be used to describe the basic model of SDT, to gen-

erate ROC curves, to analyze the data obtained in an SDT experiment, and to simulate detection processes that may be too awkward to address analytically. The advantages of simulation should be especially important for students who do not have strong backgrounds in statistics or probability theory. The spreadsheet empowers these students to teach themselves how the processes work and thereby to build up an intuition about the underlying theory. If you teach statistics, you probably already know how useful the spreadsheet is for conveying concepts involving functions, correlations, and inferential statistics. My evidence for the instructional advantages of using spreadsheet programs for learning detection theory is based on informal data. But my experience is that, once students have seen spreadsheet signal detection, they never return to the old "look-up" procedures.

REFERENCES

- Anscombe, F. J. (1956). On estimating binomial response relations. *Biometrika*, *43*, 461-464.
- Golstein, E. B. (1996). *Sensation and perception*. Pacific Grove, CA: Brooks/Cole.
- Haberlandt, K. (1994). *Cognitive psychology*. Boston: Allyn & Bacon.
- Kantowitz, B. H., & Sorkin, R. D. (1983). *Human factors: Understanding people-system relationships*. New York: Wiley.
- Kellogg, R. T. (1995). *Cognitive psychology*. Thousand Oaks, CA: Sage.
- Levine, G., & Parkinson, S. (1994). *Experimental methods in psychology*. Hillsdale, NJ: Erlbaum.
- Macmillan, N. A., & Creelman, C. D., Jr. (1991). *Detection theory: A user's guide*. Cambridge: Cambridge University Press.
- Miller, J. (1996). The sampling distribution of d' . *Perception & Psychophysics*, *58*, 65-72.
- Proctor, R. W., & Van Zandt, T. (1994). *Human factors in simple and complex systems*. Boston: Allyn & Bacon.
- Schiffman, H. R. (1996). *Sensation and perception: An integrated approach*. New York: Wiley.
- Snodgrass, J. G., Levy-Berger, G., & Haydon, M. (1985). *Human experimental psychology*. New York: Oxford University Press.

Table 1

True and Estimated d' Values and Standard Deviations for 200 Simulated Runs of 50-Trial Samples

True d'	Mean d'	$SD d'$	Bias ($d'_{\text{est}} - d'_{\text{true}}$)
0.50	0.545	0.364	0.045
1.00	1.074	0.401	0.074
1.50	1.575	0.449	0.075
2.00	2.108	0.487	0.108
2.50	2.654	0.441	0.154
2.75	2.875	0.414	0.125
3.00	3.053	0.371	0.053

- Solso, R. L. (1998). *Cognitive psychology*. Boston: Allyn & Bacon.
- Tanner, W. P., Jr., & Swets, J. A. (1954). A decision-making theory of visual discrimination. *Psychological Review*, **61**, 401-409.
- Wickens, C. D. (1992). *Engineering psychology and human performance*. New York: HarperCollins.

NOTES

1. A surprising number of these articles contain no citations of the SDT literature. This is probably indicative of the general acceptance of the SDT methodology.

2. The following description is offered for the reader who is unfamiliar with spreadsheets and their operation: A spreadsheet is a matrix of cells organized as rows (row 1,2,3 . . .) and columns (column A,B,C . . .). A cell entry can be a text label, a number, or a formula. A formula can contain the entries of other cells. Cell entries, whether they are labels, numbers, or formulas, can be easily copied from one cell to another or to multiple cells. For example, a cell entry in one column can be easily copied into all the other cells in that or another column. The usual method for copying the contents of cells from one row or column to another is relative to the original column and row. For example, suppose that cell A101 contains the formula for summing cells A1–A100; in Excel notation, this would be =sum(A1:A100). Copying the contents of cell A101 to cell B101 would produce in B101 a formula to compute the sum of cells B1–B100; in Excel notation, =sum(B1:B100). The spreadsheet is designed to allow a person to accomplish these copy operations with a minimum number of mouse movements or keystrokes; it is rarely necessary to manually enter or copy more than a small number of cell entries in order to set up a spreadsheet to perform very powerful operations. Basically, a spreadsheet is a system designed to easily perform vector and matrix operations on column and row vectors. Modern spreadsheets can also perform database and statistical operations, such as cross-tabular analyses.

3. Excel files (version 7.0) containing these demonstrations are available on the journal website for downloading.

4. An informal evaluation of student attitudes toward the use of the spreadsheet was made during an upper level sensory laboratory course. In a 4-week period during this course, the students performed a preprogrammed auditory detection experiment, analyzed the results, and wrote up an APA-style report of the experiment. The detection experiment was performed at two signal levels (in noise) and three response bias instructions. The analysis compared the predictions of classical threshold theory and SDT. A laboratory manual and lectures provided instructions on how to use a spreadsheet to perform the needed calculations and graphs. No nonspreadsheet control group was run. Many students had learned about SDT in an earlier sensory processes course. A few students were already familiar with spreadsheets, and, at the beginning of the course, several students had expressed an interest in learning how to use

a spreadsheet program. None of the students had prior experience with a numerical analysis program, such as MATLAB, MATHCAD, or Mathematica. After the project was completed, the students filled out a questionnaire that asked them about the help provided by the lecture, textbook, spreadsheet, and lab manual in understanding the theory and in calculating theoretical parameters. In both cases, students found the lab manual most helpful, the lectures next best, the spreadsheet program next, and the textbook least helpful. Many students suggested that, in future classes, everyone (including the instructors) should use the same spreadsheet program. Ten days after the project was completed, students were asked what they thought of a proposal to require the use of tables rather than a spreadsheet for calculating SDT indices and predictions. Students expressed disbelief and some indignation that such a requirement might be considered.

5. In some well-defined situations, it is possible to define an “ideal observer.” The ideal observer is an observer who is assumed to use the optimal statistic, the likelihood ratio, $z = f(x|s+n)/f(x|n)$ for converting the multidimensional input to a point on the decision axis.

6. An optimal strategy for setting z_c is to set it to maximize the expected value of the trial outcome. That is, set

$$\beta = \frac{p(n)}{p(s+n)} \times \frac{\text{value(no and noise)} - \text{cost(yes and noise)}}{\text{value(yes and signal)} - \text{cost(no and signal)}}$$

7. In Quattro Pro, mathematical functions are usually preceded by the @ symbol. So log(x) is written @log(x), and so on. The IF function in Quattro Pro is slightly different from the IF function in Excel, which is, @IF(d2=0 #AND# e2=“NO”, 1, 0). The NORMDIST and NORMSINV functions are essentially identical.

8. A simpler way to perform the hit and false-alarm calculations is to use the Excel pivot table (cross-tab) tool. In about six mouse clicks, one can set up a pivot table with the stimulus condition as the columns, the response as the rows, and the number of trials as the cell entries. The spreadsheet will then produce the confusion matrix for that condition. For the example given, the resulting table will be as follows:

Count of resp	stim		
resp	noi	sig	Grand Total
no	6	3	8
yes	3	9	12
Grand Total	9	11	20

The hit and false-alarm probabilities and SDT parameters can then be calculated as shown in the example.

(Manuscript received October 17, 1997;
revision accepted for publication April 10, 1998.)