

Human Information Processing

AN INTRODUCTION TO PSYCHOLOGY

An Introduction to Psychology

PETER H. LINDSAY and DONALD A. NORMAN

University of California, San Diego

ACADEMIC PRESS

New York and London

Copyright © 1972, by Academic Press, Inc.
All Rights Reserved

No part of this book may be reproduced in any form,
by photostat, microfilm, retrieval system, or any
other means, without written permission from
the publishers.

ACADEMIC PRESS, INC.
111 Fifth Avenue, New York, New York 10003

United Kingdom Edition published by
ACADEMIC PRESS, INC. (LONDON) LTD.
24/28 Oval Road, London NW1

Library of Congress Catalog Card Number: 70-182657

Fourth Printing, 1973

PRINTED IN THE UNITED STATES OF AMERICA

Cover design by Leanne Hinton
Text design by Wladislaw Finne

APPENDIX B

Operating characteristics

THE DECISION PROBLEM

THE DICE GAME

The criterion rule

Hits and misses

False alarms

Moving the criterion

The operating characteristic

Confidence ratings

The normal distribution

PROBLEMS

The fire sprinkler problem

Memory

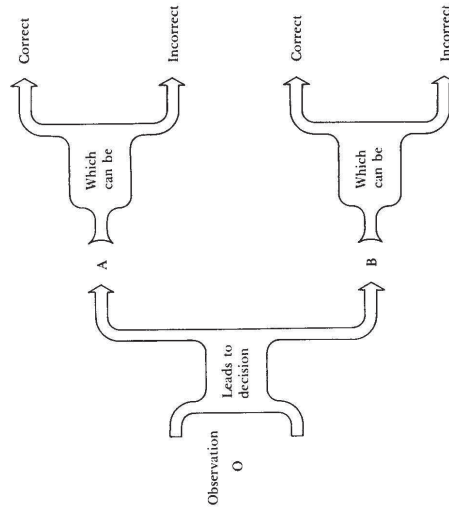
The dice game revisited

SUGGESTED READINGS

THE DECISION PROBLEM

In most real decisions that we must make, there is no answer that is guaranteed to be correct. For most situations, it is necessary to choose among virtues and evils, hoping to minimize the chance of misfortune and maximize the chance of good.

In this section, we analyze one common, simple form of decision making. The situation can be described by the following description. First, there is some information that the decision maker uses to help reach his decision. This information comes from the observations of the decision maker. We represent the results of the observations by the letter **O**. Second, the only choice open to the decision maker is to decide which of two acts, **A** or **B**, he will choose to do. Finally, the choice of decision can be correct or incorrect. Thus, we have the simple chain of events with four possible outcomes shown in Figure B-1. The prototypical example is that of the following, a dice game.



THE DICE GAME

You are gambling while playing a guessing game. Your partner throws three dice. Two of the dice are normal, one die is very special in that three of its sides are marked with the number “3” and the other three sides with the number “0.” Your job is to guess which side the special die came up on; you are told only the total scores from all three dice. (Obviously, if the total is 2, 3, or 4 the special die must be “0”;

the total is 13, 14, or 15, the special die must be “3.”) You are told that the score is 8. What should you respond?

Your observation **O** is 8.

Alternative A: Decide that it was a 3.

Possible results: 1. It was a 3. You win the bet.
2. It was a 0. You lose your money.

Alternative B: Decide that the 3 did not turn up.

Possible results: 1. It was a 3. You lose your money.
2. It was a 0. You win.

The important thing to notice about this situation is that, on the average, you cannot help but make mistakes. There is no possible way of guaranteeing perfect performance.

To analyze the dice game, consider all the possible ways that it can come out. First, how many possible results are there? Well, the lowest number that is possible comes if the two regular dice both turn up “1” and the special die comes up “0”: that gives a total of 2. The highest number comes if the two regular dice both come up 6 and the special die comes up 3: that gives a total of 15. Thus, there are 14 possible outcomes, ranging from 2 through 15.

Now consider the chance that we can get any one of those 14 scores, given that the special die was a 3 or a 0. To do this, we have to figure out how many ways the dice can turn up to give any particular score. Here is how we do that.

Suppose the total were 8: this can happen in different ways, depending upon whether the special die is 3 or 0.

If the special die is a 3:

The two regular dice must total 5: There are four ways for that to happen. The two regular dice can come up 1 and 4, 2 and 3, 3 and 2, or 4 and 1.

If the special die is 0:

The two regular dice must total 8: There are five ways for that to happen. Two regular dice can come up 2 and 6, 3 and 5, 4 and 4, 5 and 3, or 6 and 2.

In fact, for all the possible scores of the dice, we get the number of possibilities shown in Table B-1.

Now, suppose that we observed a score of 10. How many ways can that happen? Looking at Table B-1, we see that this can happen in 3 ways if the special die is 0 and 6 ways if the special die is 3: There

THE DICE GAME

is a total of 9 ways that the 3 dice can combine to give a 10. Thus, because we know that we got a score of 10, we also know that, on the average, $\frac{9}{36}$ of the time this will be a result of the special die coming up a 3 and $\frac{27}{36}$ of the time from the special die coming up 0. So, if we guess that a score of 10 means that the special die is a 3, we will be correct 6 out of every 9 trials and incorrect 3 out of every 9 trials, on the average. If you are a gambler, you would say that a score of 10 means that the odds of the special die being a 3 are 2 to 1 (six to three).

It would seem to be sensible to say the special die is 3 whenever the total is 10, because the odds favor it. In fact, look at this:

Total	Proportion of Times Special Die Is 3
7	$\frac{8}{36} = 33\%$
8	$\frac{4}{9} = 44\%$
9	$\frac{5}{9} = 56\%$
10	$\frac{6}{9} = 67\%$
11	$\frac{7}{9} = 71\%$
12	$\frac{4}{5} = 80\%$
13	$\frac{3}{3} = 100\%$

On the average, the percentage of times that calling the special die 3 will be correct rises steadily as the total score rises, with the chance being more favorable than not as soon as the total is 9 or greater. A good decision rule, thus, would appear to be "Say that the special die is 3 whenever the total score is 9 or greater." Let us see what this would cause to happen.

Hits and misses. Suppose we tossed the dice 100 times, and on each toss had to decide whether the special die was a 3. We use the decision rule of responding "Yes" every time a total of 9 or more occurs. Now go ahead to Table B-1. We will answer yes (the die is a 3) for a score of 9, 10, 11, 12, 13, 14, and 15. Otherwise we will say no. But there are 36 possible combinations of the regular two dice, and if the special die is a 3, only 26 of them give totals of 9 or greater. (We can get scores less than 9—scores of 8, 7, 6, and 5—in 10 ways.) Thus, we will be correct by saying yes 26 out of every 36 trials on which the special die really is a 3. We miss 10 out of every 36 trials that the special die really was 3.

The proportion of times that we get a hit by correctly deciding yes, the 3 has turned up, is represented as $p(\text{yes}|3)$. The vertical bar (|)

Table B-1

Number of Ways This Can Happen if the Special Die Is a

Total	0	3
0	0	0
1	0	0
2	1	0
3	2	0
4	3	0
5	4	1
6	5	2
7	6	3
8	5	4
9	4	5
10	3	6
11	2	5
12	1	4
13	0	3
14	0	2
15	0	1
16	0	0
Total Combinations:	36	36

means "conditional" or "given." Thus the terms read "the proportion of hits is equal to $p(\text{yes}|3)$, which is the proportion of yes, given that a 3 actually was rolled on the special die." In this example the hit rate or $p(\text{yes}|3) = \frac{26}{36} = 72\%$. Similarly, the miss rate $p(\text{no}|3)$, is $\frac{10}{36}$ or 28%.

False alarms. What about when the special die really was a 0? We respond yes anytime the total score is 9, 10, 11, or 12. Thus, out of the 36 combinations of the two regular dice when the special die is a 0, exactly 10 of them lead to a total score of 9 or more, the other 26 combinations lead to a total score of 8 or less. Saying yes for the wrong event is called a false alarm. In this example, the false-alarm rate is $\frac{10}{36}$: $p(\text{yes}|0) = \frac{10}{36} = 28\%$.

Moving the criterion. We can adjust how often we correctly guess that the special die turned up "3" by adjusting the critical score at which we change from an answer of "yes" to "no." But, as the critical score varies, so do the hits and the false alarms. The hit- and false-alarm rates are related; increasing one always increases the other. In fact,

the exact relationship between the hit and false-alarm rate is very important in decision theory. Call the critical score on which we base our decisions the *criterion*. Whenever the dice total equals or exceeds the criterion, we say that the special die is most likely to be a "3"; otherwise we say that it is probably a "0."


Criterion	False-Alarm-Rate		Hit Rate	
	Fraction	Percentage	Fraction	Percentage
1	$\frac{3}{8}$	100	$\frac{3}{8}$	100
2	$\frac{3}{8}$	100	$\frac{3}{8}$	100
3	$\frac{3}{8}$	97	$\frac{3}{8}$	100
4	$\frac{3}{8}$	92	$\frac{3}{8}$	100
5	$\frac{3}{8}$	83	$\frac{3}{8}$	100
6	$\frac{3}{8}$	72	$\frac{3}{8}$	97
7	$\frac{15}{32}$	58	$\frac{3}{8}$	92
8	$\frac{15}{32}$	42	$\frac{3}{8}$	83
9	$\frac{1}{8}$	28	$\frac{3}{8}$	72
10	$\frac{3}{8}$	17	$\frac{3}{8}$	58
11	$\frac{3}{8}$	8	$\frac{15}{32}$	42
12	$\frac{3}{8}$	3	$\frac{3}{8}$	28
13	$\frac{0}{8}$	0	$\frac{0}{8}$	17
14	$\frac{0}{8}$	0	$\frac{3}{8}$	8
15	$\frac{3}{8}$	0	$\frac{3}{8}$	0

The *operating characteristic*. It is easier to see the relationship between false alarms and hits if we plot them together, as shown in Figure B-2. This relationship is called an *operating characteristic*.¹ This curve shows explicitly how changing the criterion (the numbers beside the points) changes both the percentage of hits and the percentage of false alarms.

Another way of seeing how the decision rule must always be related to the trade off between hits and false alarms is to look again at the distribution of total scores shown in Table B-1. This time, draw a diagram of the distributions (Figure B-3). This is the same information originally presented in the table, but now it is clear why there must always be errors. The distribution of dice scores when the special die


¹ Originally, this relationship came from the study of radar receivers attempting to determine whether the signal seen was a real one or simply noise. Hence, the curve was called a *Receiver Operating Characteristic* or *ROC curve*. The term *ROC curve* is still widely used in the psychological literature.

FIGURE B-2



THE OPERATING CHARACTERISTIC

FIGURE B-3



is a 0 (the distribution on the left) overlaps considerably with the distribution of scores when the special die is a 3 (the distribution on the right). There is nothing that can be done about that overlap: If the dice total is 8, it could be a result of either outcome of the special dice. In the figure, a criterion of 11 is drawn in. For this criterion value, we decide to say that the special die is a 3 if we get a dice total of 11 or more, so the chance that we are correct is the chance that we get an observation of 11 or more from the distribution shown on

the right. The chance of a false alarm is the chance of an observation of 11 or more from the distribution shown on the left. Thus, simply by examining how much of each distribution lies to the right of the criterion, we can see the way the relative hit- and false-alarm rates vary as we move the criterion back and forth. This, of course, is exactly what we did in drawing the operating characteristic.

The operating characteristic shows how performance varies as we vary the decision rule. Now, what happens if we make the task easier? Suppose we change the dice game so that the special die has a 6 on three sides and 0 on the other three sides. What then? We leave this as a problem for the reader. Draw the new distribution of observations of the dice scores for the special die coming up a 6. (You already have the distribution for the case when the special die is 0.) Now draw the operating characteristic. It should include the point that has a hit rate of 83% and a false-alarm rate of 8%. If it does not, you had better review this section on operating characteristics.

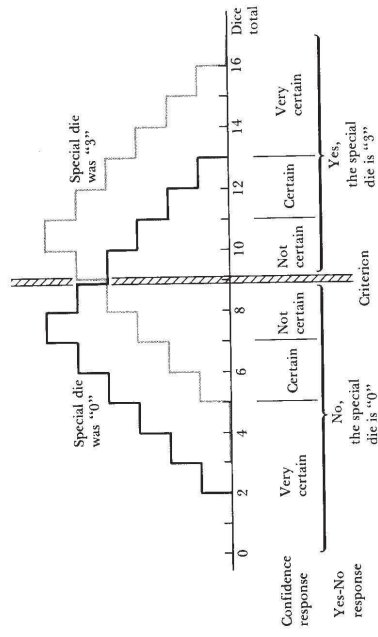
Confidence ratings

The diagram of the distributions points out something else about the decision rule: If we simply adopt a strategy of saying "3" whenever the dice total exceeds the criterion, we are wasting information. There are times when we have absolutely no doubt about the accuracy of our response, and there are times when we know that we are simply guessing: How does the decision rule describe this? The answer is simple. Whenever we get a low total on the dice, say between 2 and 4, we are certain that the special die was a 0, whenever we get a high score, say between 13 and 15, we are certain that the special die was a 3. With a value of 8 or 9 for the total score, we are guessing. Thus, we can say more than simply yes or no whether the special die is likely to be a 0 or a 3: We can also assign a statement of how confident we are in that response. We can easily qualify our answers by adding a statement like "I am very certain," or "I am pretty certain," or "I am really just guessing" to our statement of yes or no. When this is done, we see that there are really six responses:

- Yes, the special die is a 3 and I am very certain
- very certain
- not certain
- not certain
- No, the special die is a 0 and I am not certain
- certain
- very certain

These six responses can be ordered according to the dice score, with a response of "very certain that it is a 3" always coming from the highest total and "very certain that it is a 0" coming from the lowest. If we draw the way responses come from the distributions of dice totals, we might get something like that shown in Figure B-4.

FIGURE B-4



These confidence ratings are extremely useful. Note that we can treat the six different responses somewhat as if we had six different criteria for responding. Thus, the operating characteristic can be drawn to reflect confidence ratings, rather than the criteria illustrated previously. To do this, simply note that the chance of responding yes with a confidence of **certain** or greater is given by the chance that the dice total is 11 or greater. Thus, in the illustration shown in Figure B-4, the translation between criteria and confidence ratings looks like this:

<i>To Simulate a Criterion of</i>	<i>Combine These Responses</i>
13	Yes—very certain
11	Yes—very certain and certain
9	Any response of yes
7	No—Not certain and any response of yes
5	No—Not certain, certain, and any response of yes
2	Any response whatsoever

Note that in order to plot the operating characteristic we do not really need to know what the criteria are. All we need to know is what the hit- and false-alarm rates are for the various responses.

Suppose we did the dice game for 200 times. Furthermore, suppose that on 100 trials the special die came up 0 and on 100 trials it came up 3. After the experiment, we sort out the responses according to whether they resulted from a 3 or 0 on the special die. Suppose that this is what we found.

Responses	Number of Occurrences When Special Die Was	
	0	3
A. Yes—very certain	0	17
B. Yes—certain	17	41
C. Yes—not certain	11	14
D. No—not certain	30	20
E. No—certain	25	8
F. No—very certain	17	0
TOTAL:	100	100


Now, without bothering to figure out what criterion each response represents, we simply realize that we can treat these responses as if each came from a criterion, if we lump together all responses of a certain confidence or greater:

Response	(Special Die Was 0) False-Alarm Rate	(Special Die Was 3) Hit Rate
A	0	17
B or A	17	58
C, B, or A	28	72
D, C, B, or A	58	92
E, D, C, B, or A	83	100
F, E, D, C, B, or A	100	100

If we plot the hit- and false-alarm rates, we get the operating characteristic (Figure B-5)—the same curve shown in Figure B-2.

This is exactly how we analyze real data, the only exception being that in a real experiment the numbers would not come out quite so cleanly. People are inconsistent in where they place their criteria. These inconsistencies are relatively small, however.

FIGURE B-5



The normal distribution

We ask a human subject to listen to a very weak signal that is presented to him periodically over a pair of earphones. We want to find out whether or not he can hear the signal. The question is actually much more complex, however, because the subject is always hearing something: He must decide whether what he has heard resulted from the signal presented, or whether it was simply a result of the normal fluctuations in hearing that occur. These fluctuations come about for many reasons. In fact, in many experiments, we add noise to the earphones in order to see how well the subject can pick out the signal from the noise.

The situation for the subject is very much like the situation described for the dice game. He listens during the interval when the signal could be presented and ends up with some observation, much like our rolling the dice and ending up with some total score. The question is, did that observation come from the signal or just from noise. The analogous question for the dice game, of course, is, "Did that total result from the special die being a 0 or a 3?" We assume the subject who tries to detect the signal chooses some criterion: If his observation exceeds that criterion he says "signal." Otherwise he says, "no signal." From his hit- and false-alarm rates, we try to determine the separation of the distributions that he must be using to make his decision. Then,

from our determination of the distributions, we try to decide how the auditory system must be converting the signals.
 Let us now work through some examples. Before we do, however, we need to introduce a special type of distribution of observations, the *normal distribution*.

When we played the dice game, we developed the distribution of outcomes of the dice (Figure B-3). In general, however, a different type of distribution is frequently encountered. This distribution is called the *normal distribution*, and is an extremely useful one to know about. It is widely used in many fields of study, including psychology, and it usually turns out that even if the actual distributions under study are not normal, the normal is an excellent approximation to the true one. A drawing of the normal distribution is shown in Figure B-6. Notice





FIGURE B-6

that it looks very much like the distribution for the dice game, except it is drawn smoothly, rather than with steps. This is because the total score from a dice game can only take on an integral value—it must be a number like 6 or 7, it cannot lie between. The normal distribution, however, can take on any real number, positive or negative. The normal distribution shown here is characterized by one number—the mean or average value. As it is drawn, it has an average value of zero. If it were to have an average value of, say, 1.5, the distribution would simply be shifted to the right, so that its peak was at 1.5: This is shown in Figure B-7.

The values of the normal are shown in Table B-2. Here we see the height of the curve for different values along the horizontal axis. In addition, we also show the percentage of the curve that lies to the right of any criterion. It is this latter figure that we use for computing the operating characteristics.

What we usually care about is how far apart the mean values of two distributions are from one another. Suppose we do the experiment

Table B-2 The Normal Distribution Height and Percentage of the Curve (Area) to the Right of Any Criterion



Crite- rion	Y	Per- centage	Crite- rion	Y	Per- centage
-3.0	.004	99.9	-1.0	.242	84.1
-2.9	.006	99.8	-.9	.266	81.6
-2.8	.008	99.7	-.8	.290	78.8
-2.7	.010	99.7	-.7	.312	75.8
-2.6	.014	99.5	-.6	.333	72.6
-2.5	.018	99.4	-.5	.352	69.2
-2.4	.022	99.2	-.4	.368	65.5
-2.3	.028	98.9	-.3	.381	61.8
-2.2	.035	98.6	-.2	.391	57.9
-2.1	.044	98.2	-.1	.397	54.0
-2.0	.054	97.7	0	.399	50.0
-1.9	.066	97.1	.1	.397	46.0
-1.8	.079	96.4	.2	.391	42.1
-1.7	.094	95.5	.3	.381	38.2
-1.6	.111	94.5	.4	.368	34.5
-1.5	.130	93.3	.5	.352	30.1
-1.4	.150	91.9	.6	.333	27.4
-1.3	.171	90.3	.7	.312	24.2
-1.2	.194	88.5	.8	.290	21.2
-1.1	.218	86.4	.9	.266	18.4
			2.0	.054	2.3
			2.1	.044	1.8
			2.2	.035	1.4
			2.3	.028	1.1
			2.4	.022	.8
			2.5	.018	.6
			2.6	.014	.5
			2.7	.010	.4
			2.8	.008	.3
			2.9	.006	.2




FIGURE B-7




FIGURE B-8

located relative to the noise distribution and also where the criterion is. To start, call the mean value of the noise distribution 0. There is good reason for doing this, in the absence of signal the average observation ought to be around zero. Moreover, since we only care about the relative separation of the two distributions, it doesn't really matter what number we call the mean value of the noise (our measurement will be on an *interval scale*; see Appendix A). We call the distance from the mean of the noise distribution to the criterion A, the distance from the mean of the signal distribution to the criterion B, and the distance from the mean of the noise distribution to the mean of the signal distribution d' . The symbol d' is used for historical reasons: That is what it has been called in the psychological literature. Both A and B are distances from the mean value of distribution. If the criterion is to the right of the mean, A and B are positive. If the criterion is exactly at the mean, then the distance value is 0. If the criterion is to the left of the mean, the distance is negative. Thus $d' = A - B$.

Suppose our subject gives us a false-alarm rate of 14% and a hit rate of 95%. We can immediately determine A: If we look up 14% in Table B-2, we see that the criterion must be located a distance of 1.1 units to the right of the noise distribution. Thus, $A = 1.1$. Similarly, we see that a hit rate of 95% requires that the criterion be 1.6 units to the left of the mean of the signal distribution (the criterion value is at -1.6). Thus, $B = -1.6$. Now we know that $d' = 2.7$. And that is all there is to it.

PROBLEMS

The fire-sprinkler problem

At this point, we can probably learn most about the use of operating characteristics and about the normal distribution by working a few problems.

We are installing a sprinkler system as part of a fire-alarm system for a building. Now we wish to install the temperature control that will turn on all the sprinklers whenever a fire occurs. The control is located near the ceiling of a large store room. The roof is made of tin, and there are no windows in the room. Questions: To what temperature should we set the control? If the temperature is set too low, (say 130°), then on very hot days, when the outside temperature goes as high as 110°, it is quite likely that the hot air will rise to the ceiling of the store room and be heated even more by the sun warming up the tin roof. Thus, it would not take long for the air temperature to reach 130° and set off the system: a *false alarm*.

If, however, the temperature is set higher, say 180°, it is quite likely that a fire could develop and destroy a good deal of the items in the

we mentioned in which we ask the subject to try and detect a signal that we present to him. We want to find out how far the distribution of observations that results from noise is from the distribution which results from the signal. We think that the situation is characterized by Figure B-8.

We want to discover both exactly where the signal distribution is

PROBLEMS

storeroom before the flames got high enough to heat the air at the ceiling to a temperature of 180°. Thus, we would fail to report many fires, at least while they were still small enough that the sprinkler system could put them out. This would be a *miss*. Where do we set the temperature?

To solve this problem, we need information about hits and false alarms. We need to know the probabilities with which these occur. Ideally, we would set up a test situation and watch what happens over, say, a 3-month period, carefully counting the occurrences of hits (correct triggering of the system to a fire), misses (failure to respond within, say, 5 min of a fire), false alarms (triggering of the system in the absence of a fire), and correct rejections (no response from the system in normal conditions). Then we could plot an operating characteristic.

The way we plot the operating characteristics is to vary the temperature setting of the control, collecting information about the hit- and false-alarm rate at each temperature setting. Thus, if we set the control at 140°, we might observe that the actual room temperature reaches that value on one day out of every five—giving a false-alarm figure of 20%—and we might also note that 88% of the fires that we set caused the room temperature to reach that value within the 5 min we require—a hit rate of 88%. This, then, is the first point on our curve: $p(\text{alarm}/\text{fire}) = 88\%$; $p(\text{alarm}/\text{no fire}) = 20\%$.

This point is actually sufficient, if we believe that everything is normally distributed. We can now compute the value of d' and then compute what the rest of the curve should look like.

If we go back to the table of the normal distribution, Table B-2, we see that if we have a false alarm rate of 20%, the criterion must be to the right of the highest point on the distribution, at about .8: That is, the value of A is 0.8. A hit rate of 88% means that the criterion must be located to the left of the highest point on the distribution, at a point around -1.2 . Thus, $B = -1.2$. Now d' is simply the distance that the two distributions are apart, and that is given by $0.8 + 1.2 = 2.0$. Our fire-alarm system has a d' of 2.0. The entire curve, therefore, looks like that shown in Figure B-9.

Now, to complete our information about the setting of the temperature limit we can simplify our procedure: All we do is find out what the false-alarm rate would be at different temperatures. To get this information, we can install an automatic temperature recorder in the building for a few months. Then, we look at the distribution of temperatures reached throughout that period. We might find that at a temperature of 150°, there was a false alarm only 10% of the time, at a temperature




FIGURE B-9

of 160° only 1% of the time, and at a temperature of 130°, 60% of the time. These values then determine points on the operating characteristic, as shown.

At this point, it is obvious that we can never survive with a d' as low as 2.0. If we set the false-alarm value at a level acceptable to the fire department, say 1%—a temperature setting of 160°—then our insurance company will complain that we will only detect a fire with a chance of $\frac{.37}{1.0}$. If we try to detect the fire with a chance as high as $\frac{.37}{.05}$, we will have a false alarm rate of close to 40%—clearly unacceptable to the fire department. It is quite clear that we can never solve the problem by trying to adjust the temperature setting of the sprinklers and alarms. We have to raise the d' value.

Suppose that both the fire department and the insurance company agree that an acceptable hit and false alarm rate would be 99% and 1% respectively. What value of d' would we have to have?

Memory

From experiments in memory, we know that if a list of 30 names is presented to you once (for about 2 sec per name), an hour later the retention of that list will be very low. In fact, for any individual name, $d' = 0.8$.

Suppose you were a member of a receiving line at a formal party and in a 60-sec time period, 30 people had been introduced to you. An hour later you try to recall their names. Assuming that you adjust your false-alarm rate to be 8%, what percentage of the names do you remember?

The dice game revisited

Consider a version of the three-dice game in which the special die has **O** on three sides and **S** on the other three sides. Using the normal distribution as a good approximation of the dice distribution, what is the relationship between d' and the value of **S**?

Assume that there is a fixed criterion at 11. This means there will be a false-alarm rate of 8%. Thus, if **S** is 3, we see from our dice-game table that the hit rate is $\frac{4}{8}$ or 42%. Going to the normal distribution tables, $A = 1.4$, $B = 0.2$, so $d' = A - B = 1.2$. The relationship between d' , hit rate, and false-alarm rate (assuming a fixed criterion of 11) is shown in Table B-3. The relationship between d' and **S** is plotted in Figure B-10. Now, you should try to complete both the table and the figure.

Table B-3

Value of S	False-Alarm Rate	Hit Rate	A	B	A - B = d'
0	8%	8%	1.4	1.4	0
1	8%	—	1.4	—	—
2	8%	—	1.4	—	—
3	8%	42%	1.4	0.2	1.2
4	8%	—	1.4	—	—
5	8%	—	1.4	—	—
6	8%	83%	1.4	-1.0	2.4
7	8%	—	1.4	—	—
8	8%	—	1.4	—	—




FIGURE B-10

SUGGESTED READINGS

The decision theory discussed here grew out of the engineering literature and it has mostly been applied to the study of sensory processes: to psychophysics. Because it was first applied to the analysis of detecting signals in noise, it usually goes under the name of *Signal Detection Theory*, or sometimes simply *SDT*. Thus, to find this topic in book indices, one must usually look for "signal detection theory" or sometimes for d' .

The best overall introduction to the many uses of the decision theory discussed here is the book by David Green and John Swets: *Signal Detection Theory and Psychophysics* (1966). This book does get very technical, but much of the material in the early chapters can be followed without too much difficulty even by those whose mathematics is weak. Some of the latter chapters review the various uses of the decision theory to other areas of psychology.

The chapter by Egan and Clarke (1966) in Sidowski's book on experimental methods offers another very good introduction to the technique. The book of collected readings edited by Swets (1964) gives a collection of uses, but this is very technical material.

This decision theory has been widely used in other areas. A good (and easy to follow) description of its application to the study of the retrieval of material from libraries by various automatic systems is given in the *Science* article by Swets (1963). But perhaps the most widely encountered use of this method of analysis has been in the study of memory. Two typical studies, both introducing the technique and illustrating what can be done with it are the ones by Norman and Wickelgren (1965) and Wickelgren and Norman (1966). A much simpler introduction to these studies is the short description given in Norman's *Memory and Attention* (1969a, pp. 148-161). (This is the same theory discussed in Chapter 9 of this book.) Many of the advanced theories presented in Norman's *Models of Human Memory* (1970) rely heavily on detection theory analysis.