Signal detection theory

Filip Děchtěrenko, Jiří Lukavský

Example

- Hygiene calls you: your covid test results were positive

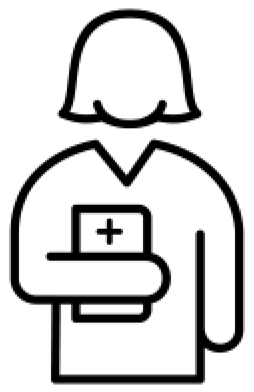

- Ugh... but how good is the test, anyway?

Possible test outcomes

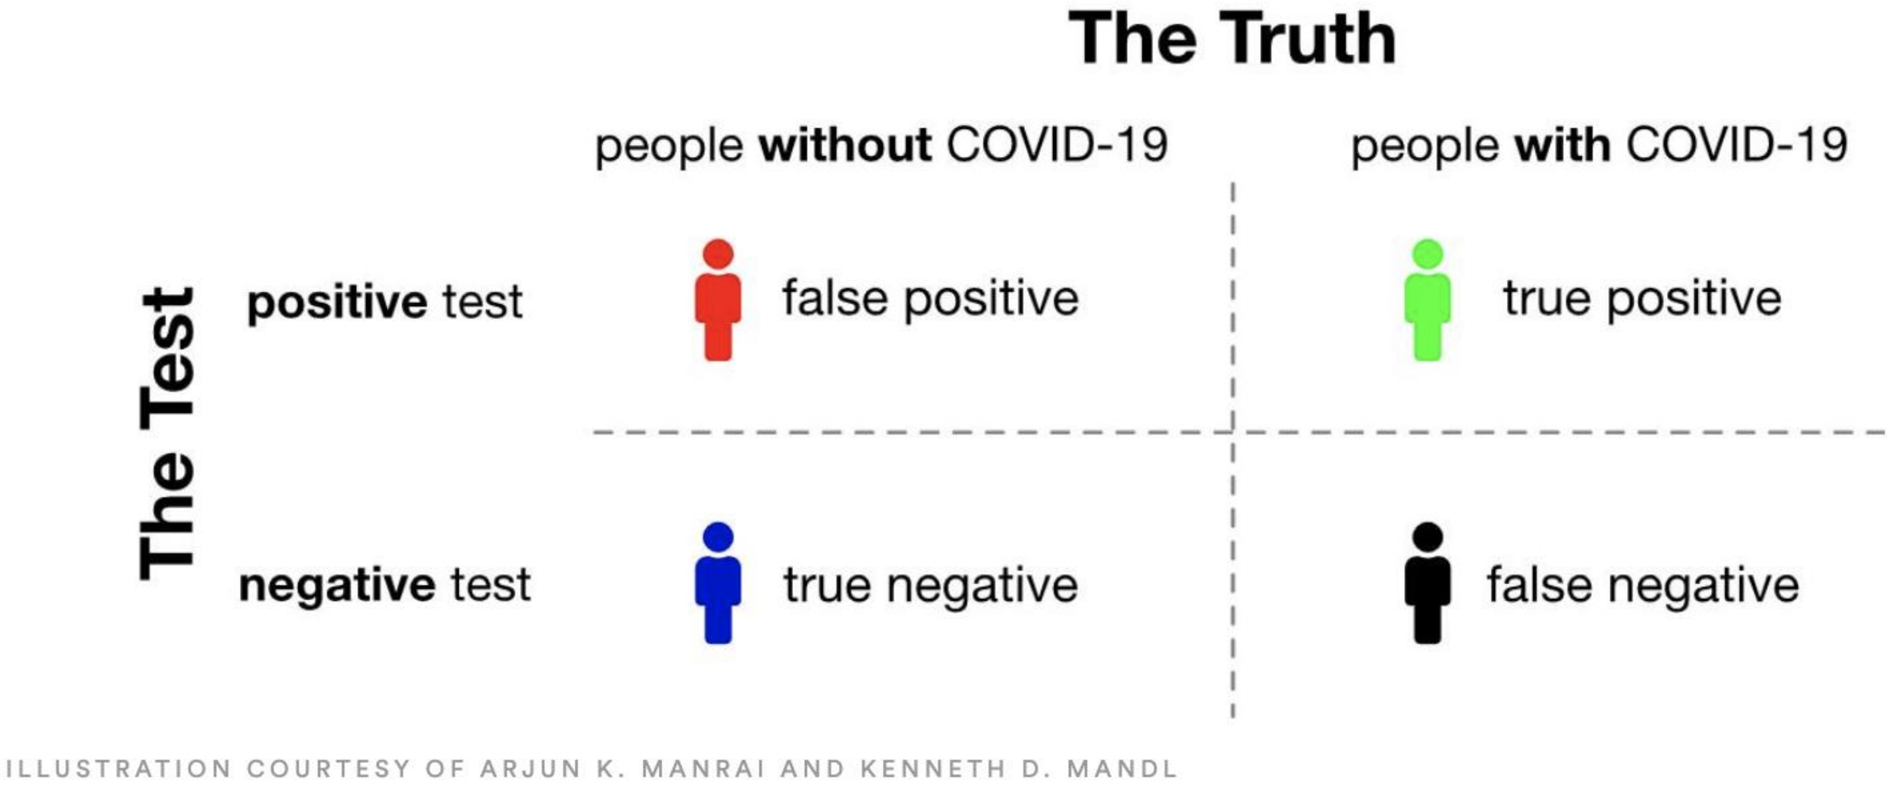

Possible test outcomes

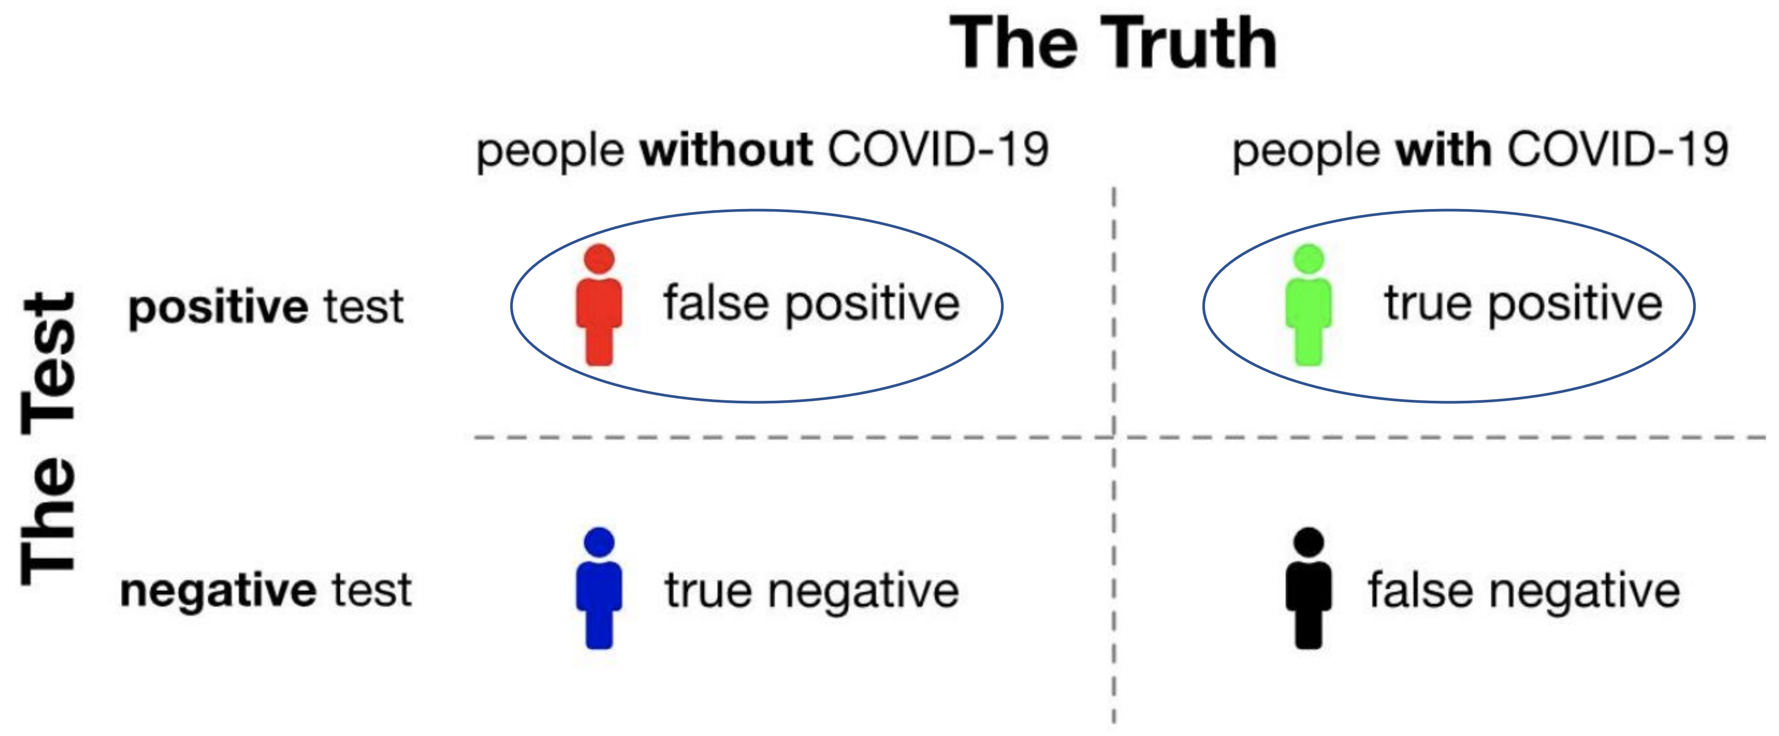

ILLUSTRATION COURTESY OF ARJUN K. MANRAI AND KENNETH D. MANDL

Using false positives and true positives we can describe the whole table. It is important to understand their relationships

- Hygiene calls you: your covid test results were positive

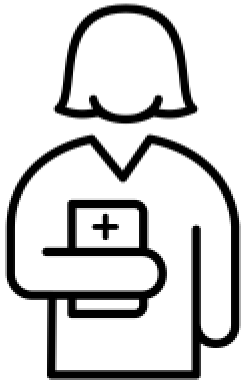

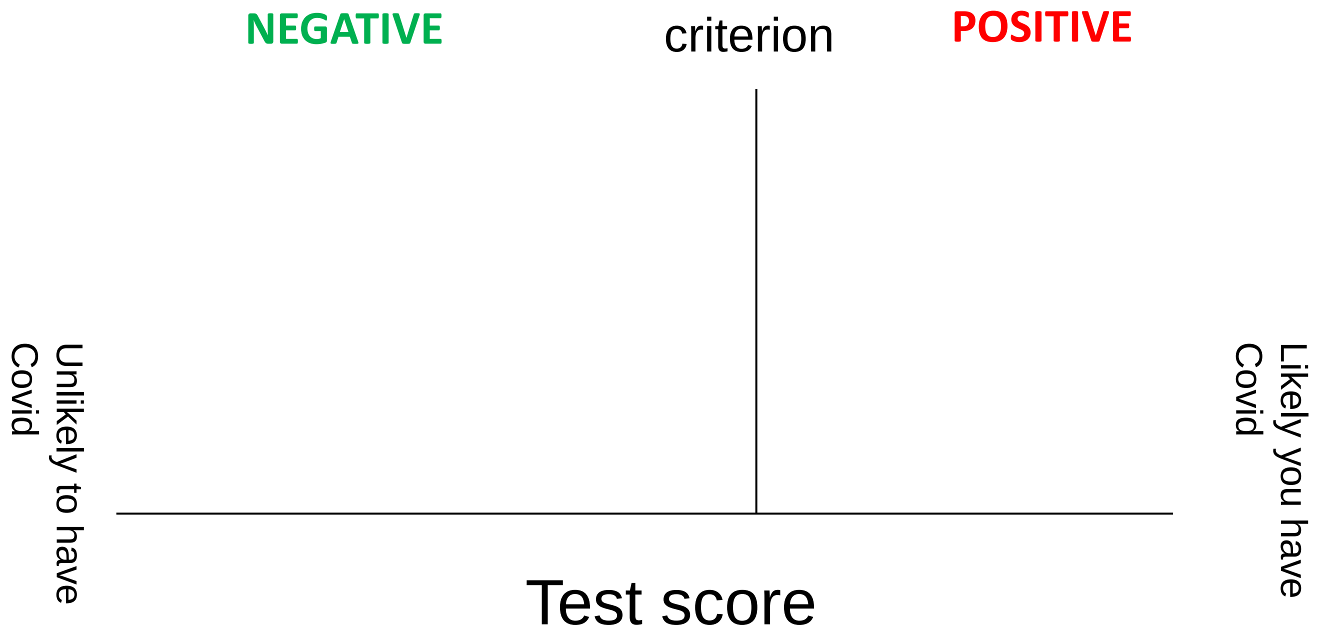

Běžný jazyk

- 🔥 How sensitive should we make this smoke alarm?
- ⚽ He's a great football referee
- How well can we distinguish between cases with and without covid?
- Beyond a reasonable doubt ⚖️ 🧑‍⚖️

Běžný jazyk

- 🔥 How sensitive should we make this smoke alarm?
- ⚽ He's a great football referee
- How well can we distinguish between cases with and without covid?
- Beyond a reasonable doubt ⚖️ 👨‍⚖️

SDT

- criterion, false alarms
- criterion, sensitivity
- sensitivity
- criterion, false alarms

Analysis of data from Yes/No experiment

- Participants are shown a series of faces to remember

Analysis of data from Yes/No experiment

- Participants are shown a series of faces to remember

Data

Stimuli type	Was this image shown?		
	yes	no	total
old	20	5	25
New	10	15	25

Citlivost

- Typically, we are interested in how able a participant is to distinguish between two categories of stimulus
- We're talking about sensitivity

Stimuli type	Was this image shown?	
	yes	no
Old (S_2)	Hit	Miss
New (S_1)	False alarm	Correct Rejection

- H = P(„yes“|S_2)
- F = P(„yes“|S_1)

For data from the experiment

$H = P(„yes“|S_2)$
$F = P(„yes“|S_1)$

We denote by (F,H), so for this experiment (.4, .8)

Stimuli type	Was this image shown?		
	yes	no	total
old (S_2)	20	5	25
new (S_1)	10	15	25

Stimuli type	Was this image shown?		
	yes	no	total
old (S_2)	.8	.2	1.0
new (S_1)	.4	.6	1.0

Simple approach

- Consider only the hits

80%

Stimuli type	Was this image shown?		
	yes	no	total
old (S_2)	.8	.2	1.0
new (S_1)	.4	.6	1.0

Simple approach

- Which is not ideal

80%

Stimuli type	Was this image shown?		
	yes	no	total
old (S_2)	.8	.2	1.0
new (S_1)	.4	.6	1.0

- We need to work with F as well

32%

Stimuli type	Was this image shown?		
	yes	no	total
old (S_2)	8	17	25
new (S_1)	1	24	25

Two alternatives

- Sensitivity

  H-F

.8 - .4 = .4

Stimuli type	Was this image shown?		
	yes	no	total
old (S_2)	.8	.2	1.0
new (S_1)	.4	.6	1.0

.32 - .04 = .28

Stimuli type	Was this image shown?		
	yes	no	total
old (S_2)	.32	.68	1.0
new (S_1)	.04	.96	1.0

Two alternatives

- Sensitivity

  H-F

.5(.8 + .6) = .7

Stimuli type	Was this image shown?		
	yes	no	total
old (S_2)	.8	.2	1.0
new (S_1)	.4	.6	1.0

- Sensitivity

  p(c) = 0.5(H+(1-F))

.5(.32 + .96) = .64

Stimuli type	Was this image shown?		
	yes	no	total
old (S_2)	.32	.68	1.0
new (S_1)	.04	.96	1.0

Two alternatives - Which is better?

H-F vs 0.5(H+(1-F))

Two alternatives - Which is better?

$H-F$ vs $0.5(H+(1-F))$

- They are equal!

$$0.5(H+(1-F)) = 0.5(H-F) + 0.5$$

Signal detection theory approach

- d’ = z(H) – z(F)

Psychological Review
Vol 61, No 6, 1954

A DECISION–MAKING THEORY OF VISUAL DETECTION [1]

WILSON P. TANNER, JR. AND JOHN A. SWETS
University of Michigan

This paper is concerned with the human observer's behavior in detecting light signals in a uniform light background. Detection of these signals depends on information transmitted to cortical centers by way of the visual pathways. An analysis is made of the form of this information, and the types of decisions which can be based on information of this form. Based on this analysis, the expected form of data collected in "yes-no" and "forced-choice" psychophysical experiments is defined, and experi-

where p' is the observed proportion of positive responses, p is the corrected proportion of positive responses, and c is the intercept of the dotted curve at $\Delta I = 0$.

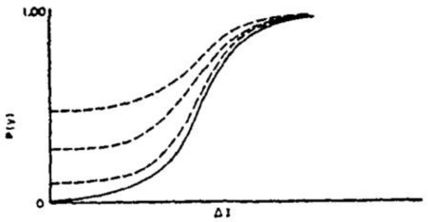

The parameter d' is the square root of Peterson and Birdsall's d. The square root of d is more convenient here; d' is the difference between the means of N and $S + N$ in terms of the standard deviation of N. The cri-

Signal detection theory approach

- d' = z(H) – z(F)

d' = .842-(-.253) = 1.095

Stimuli type	Was this image shown?		
	yes	no	total
old (S_2)	.8	.2	1.0
new (S_1)	.4	.6	1.0

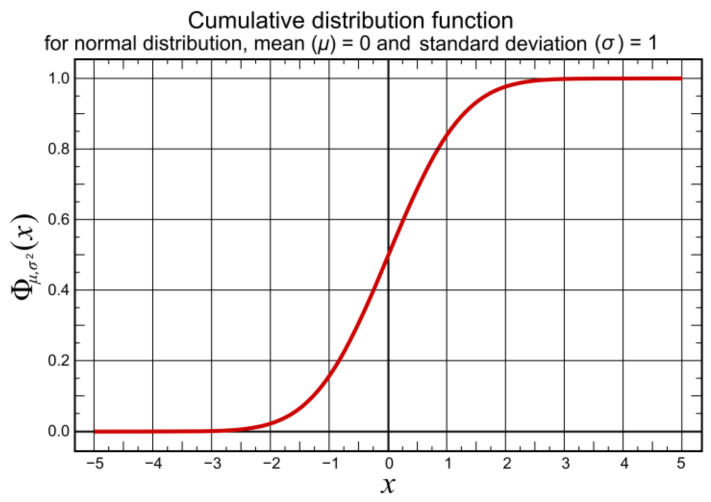

d' = -.468-(-1.751) = 1.283

Stimuli type	Was this image shown?		
	yes	no	total
old (S_2)	.32	.68	1.0
new (S_1)	.04	.96	1.0

Signal detection theory approach

- $d' = z(H) - z(F)$

- if H=F, i.e. the observer cannot distinguish between signal and noise, d' = 0
- if H is .99 and F is .01, d' = 4.65
- perfect accuracy is a problem (z(0) = -Inf , z(1) = Inf)

Stimuli type	Was this image shown?	
	yes	no
old	20	5
new	0	15

Stimuli type	Was this image shown?	
	yes	no
Old	20.5	5.5
new	0.5	15.5

Conversion between p(c) and d’

- If we assume that H = 1-F

$$d' = 2z(p(c))$$

$$p(c) = .9 \rightarrow d' = 2.56$$

- But if they are significantly different, we get a big difference
- For example, if H = .99 and F = .19, then p(c) = .9, but d' = 3.20

ROC curve

- The sensitivity is constant for the participant for a given task, i.e. only their preference for yes/no can change
- Since d' is a combination of H and F, we get the same d' for different values of (.4, .8) ~ (.2, .6) ~ (.07, .35)

ROC curve

- The sensitivity is constant for the participant for a given task, i.e. only their preference for yes/no can change
- Since d' is a combination of H and F, we get the same d' for different values of (.4, .8) ~ (.2, .6) ~ (.07, .35)
- The line connecting the same values is the isosensitivity curve (also sometimes Receiver operating characteristic, or Relative operating characteristic)
- If the curves pass through 0 and 1, they are regular

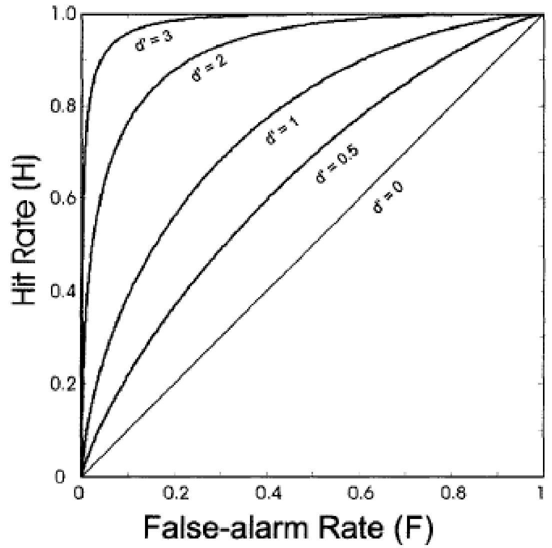

FIG. 1.1. ROCs for SDT on linear coordinates. Curves connect locations with constant d'.

ROC in transformed coordinates

- because $z(H) = z(F) + d'$

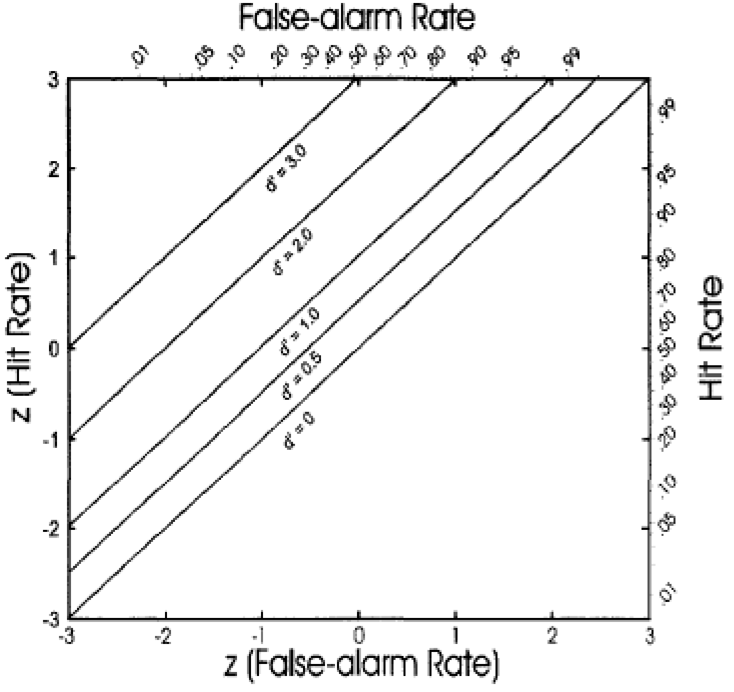

FIG. 1.2. ROCs for SDT on *z* coordinates.

ROC for p(c)

- H = F + 2p(c) – 1

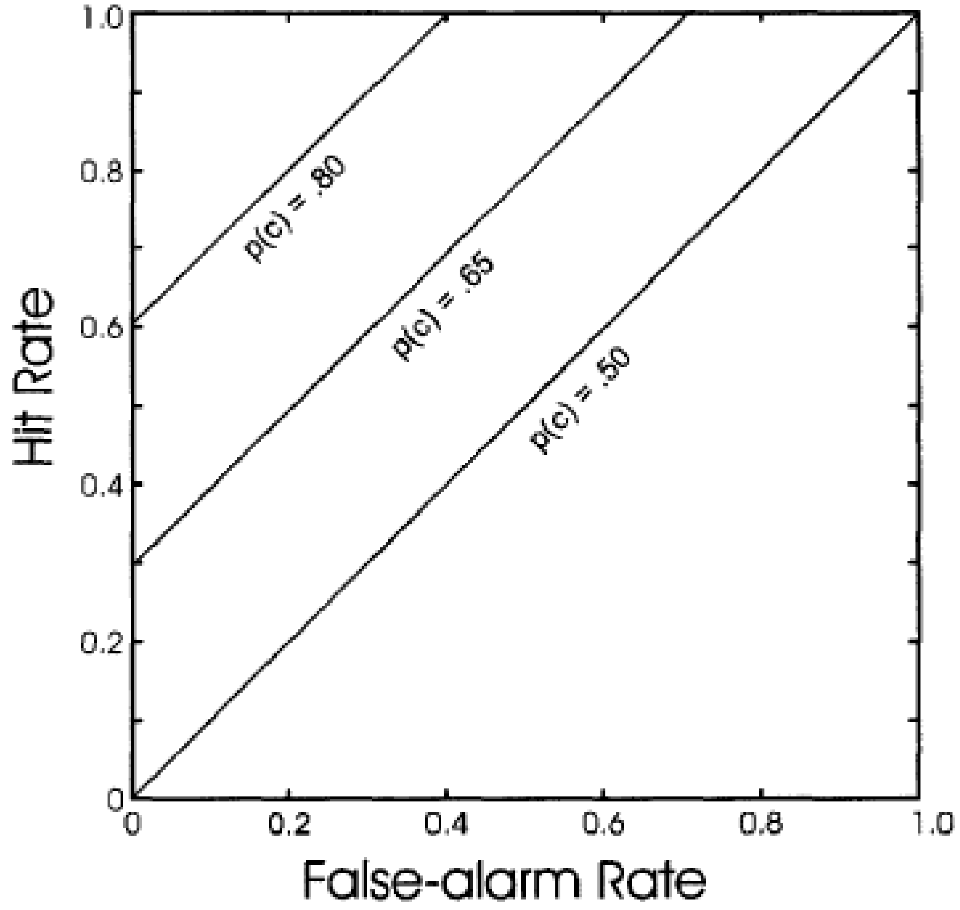

FIG. 1.3. ROCs implied by *p*(*c*) on linear coordinates.

Which is better?

- So why use d' and not p(c)?

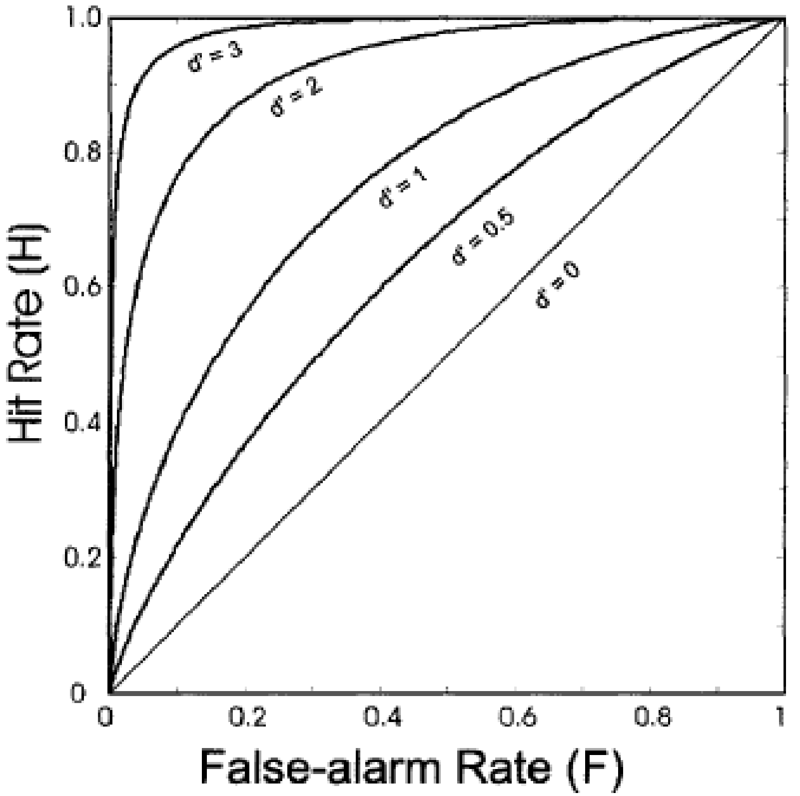

FIG. 1.1. ROCs for SDT on linear coordinates. Curves connect locations with constant d'.

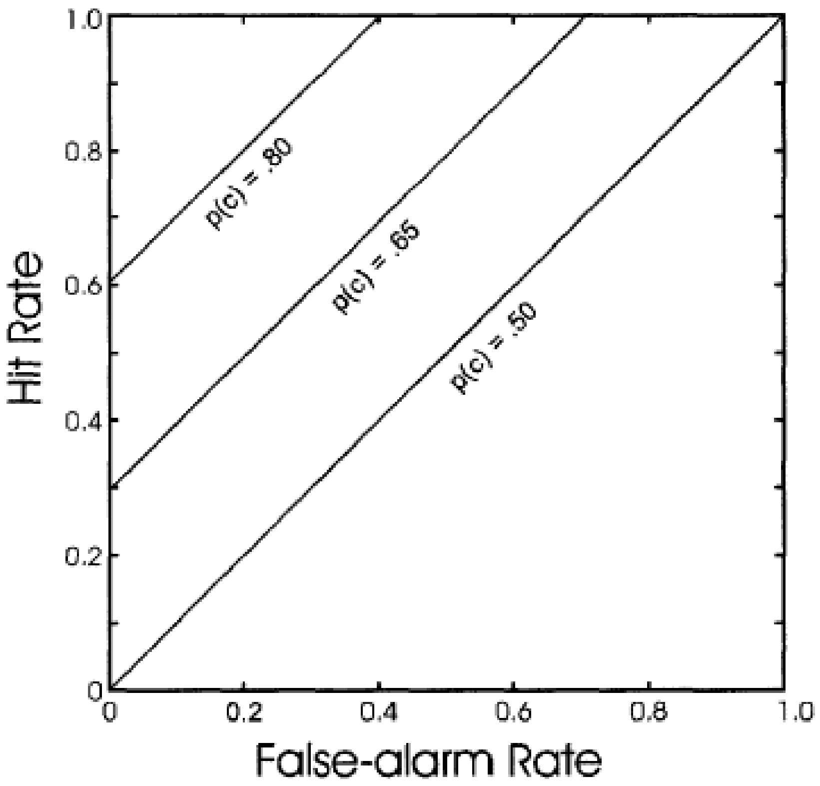

FIG. 1.3. ROCs implied by $p(c)$ on linear coordinates.

Which is better?

- So why use d' and not p(c)?
- d' fits the real data better
- d' can be used as perceptual distance
  $d'(x, w) < d'(x, y) + d'(y, w)$.

Which is better?

- So why use d' and not p(c)?
- d' fits the real data better
- d' can be used as perceptual distance $d'(x, w) < d'(x, y) + d'(y, w)$.
- d' is of ratio scale , so the stimuli can be said to be twice as perceptually distant from each other (which is not the case for p(c))

Which is better?

- So why use d' and not p(c)?
- d' fits the real data better
- d' can be used as perceptual distance d'(x, w) < d'(x, y) + d'(y, w).
- d' is of ratio scale , so the stimuli can be said to be twice as perceptually distant from each other (which is not the case for p(c))
- At unit slope, symmetry applies

  d’(1-H,1-F) = d’(F,H)

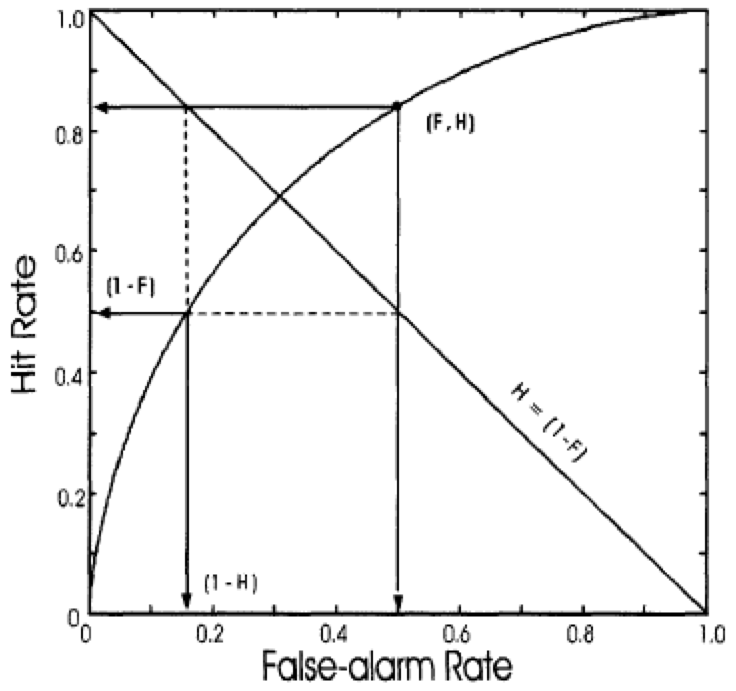

FIG. 1.4. The points (*F*, *H*) and (1 – *H*, 1 – *F*) lie on the same symmetric ROC curve.

Comparison between tasks

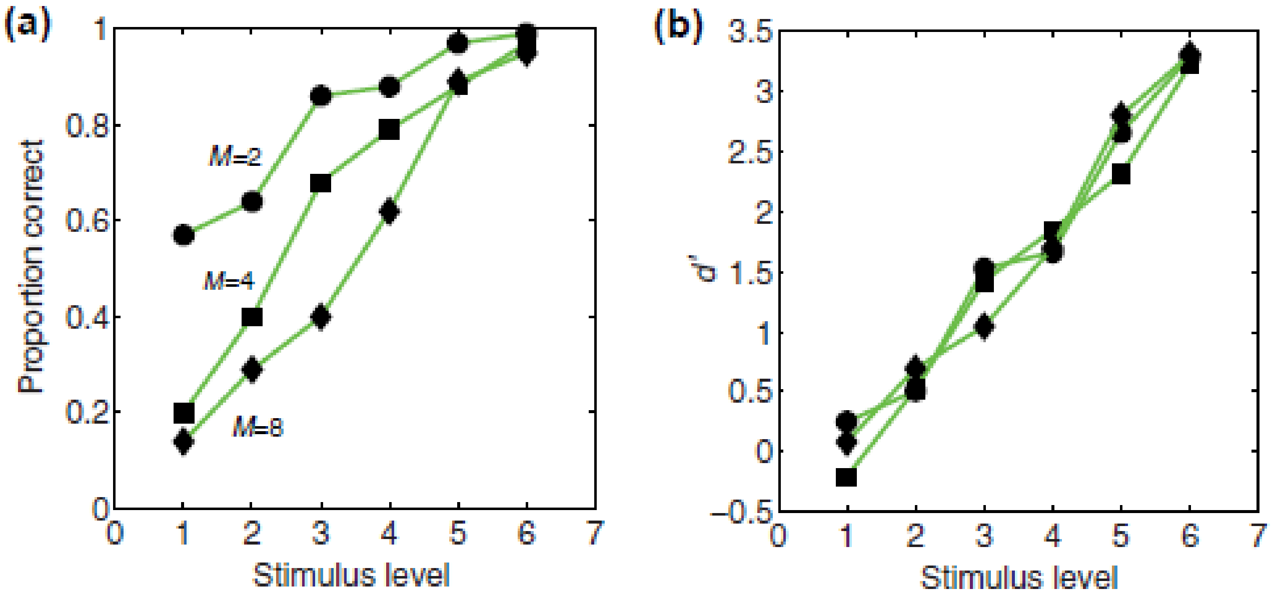

FIGURE 6.1 (a) Hypothetical Pc (proportion correct) data for an $M = 2$, $M = 4$, and $M = 8$ forced-choice task. (b) The same data plotted as d's.

ROC in another tasks

- In diagnostic setting, we often base our diagnosis on some cutoff value. We can obtain similar false alarms and hit rates, however it is often denoted as sensitivity (=H) and specificity (=1-F)
- We often use AUC (area under curve) to describe ROC curve

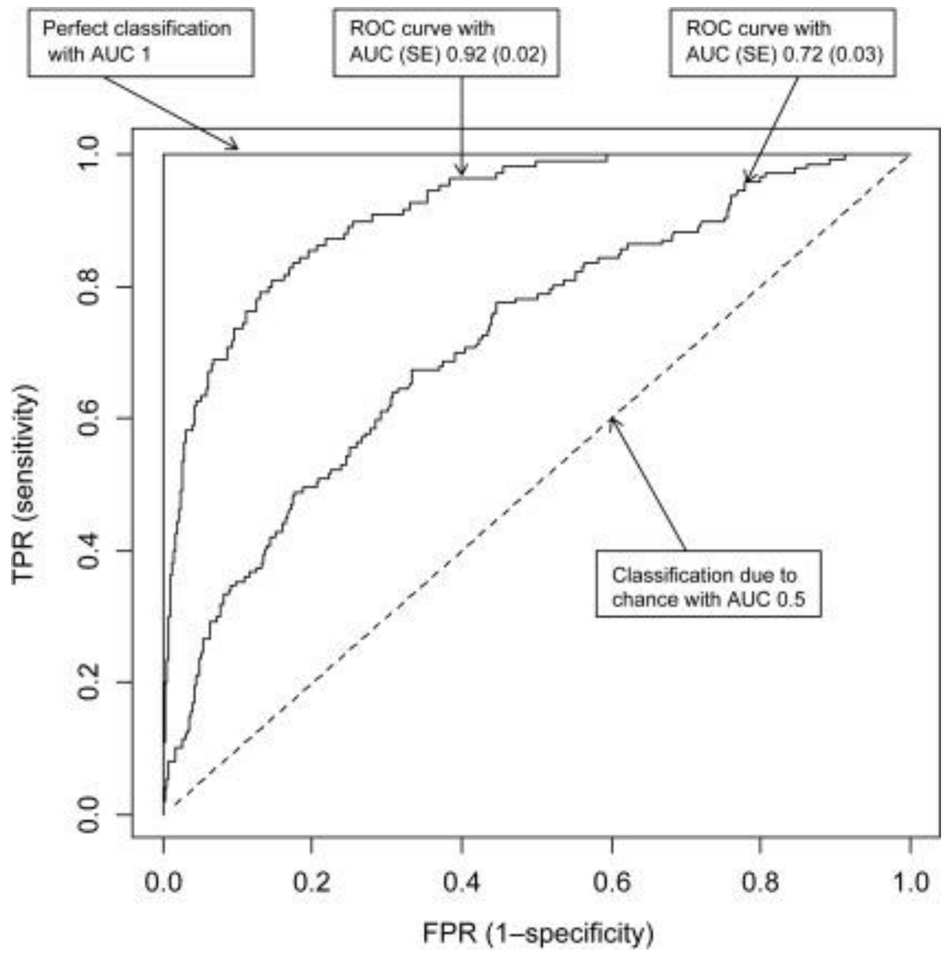

Selection from the normal distribution

- We'll use visualizations with normal distributions that look like this

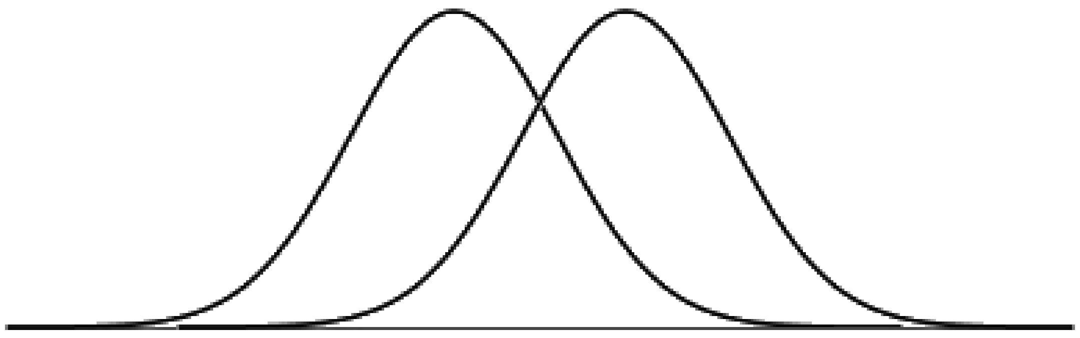

Selection from the normal distribution

- It's a density chart

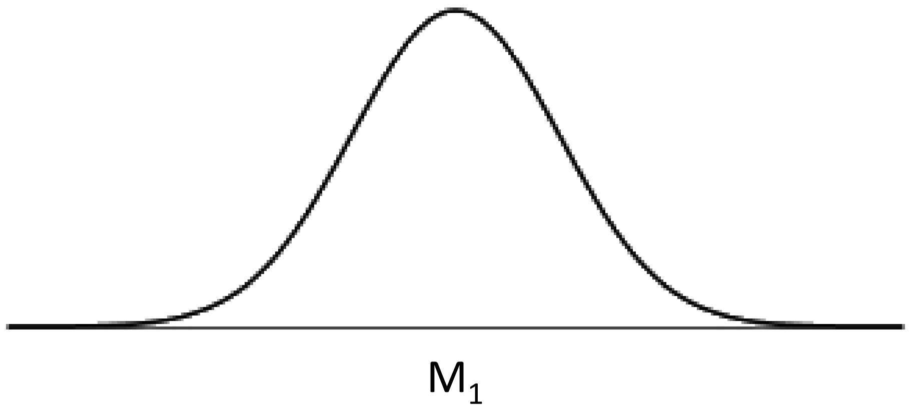

Selection from the normal distribution

- It is a density graph
- If we pick a random number from this distribution, we will have most of the numbers from the range

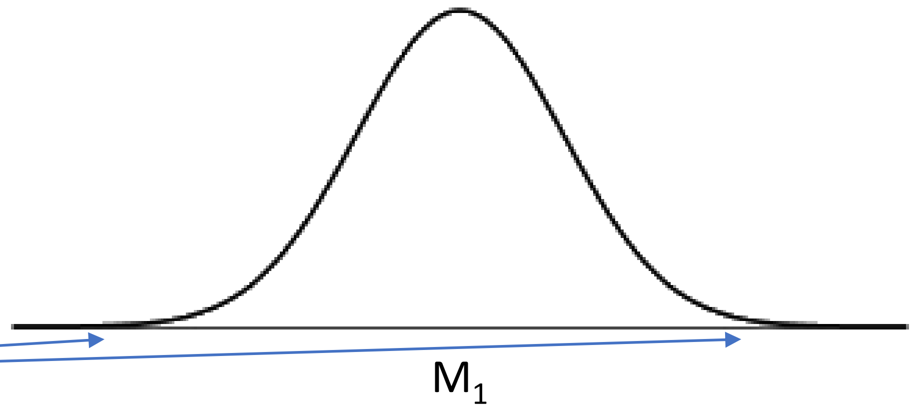

Selection from the normal distribution

- It is a density graph
- If we pick a random number from this distribution, we will have most of the numbers from the range
- Most of them will be around the mean

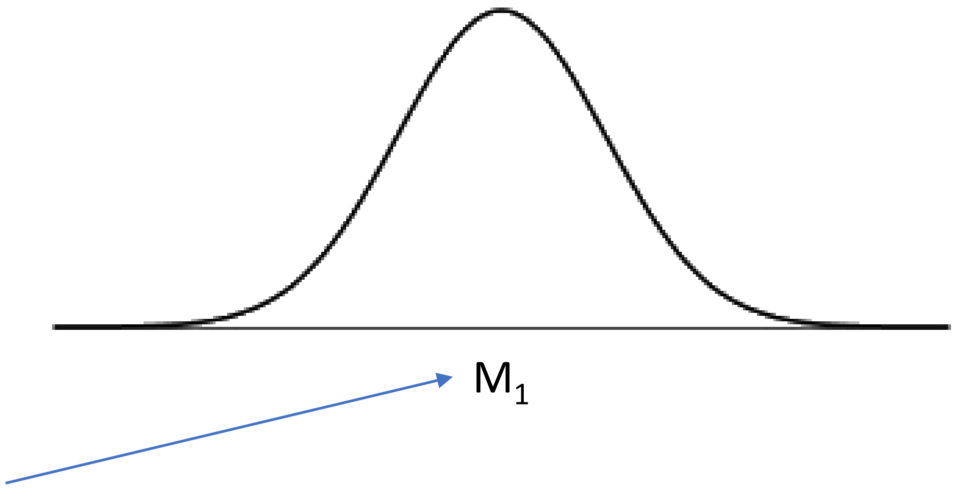

Selection from the normal distribution

- It is a density graph
- If we pick a random number from this distribution, we will have most of the numbers from the range
- Most of them will be around the mean
- If we theoretically had thousands of them and made a histogram, it would have the same shape as the density plot (distribution function)

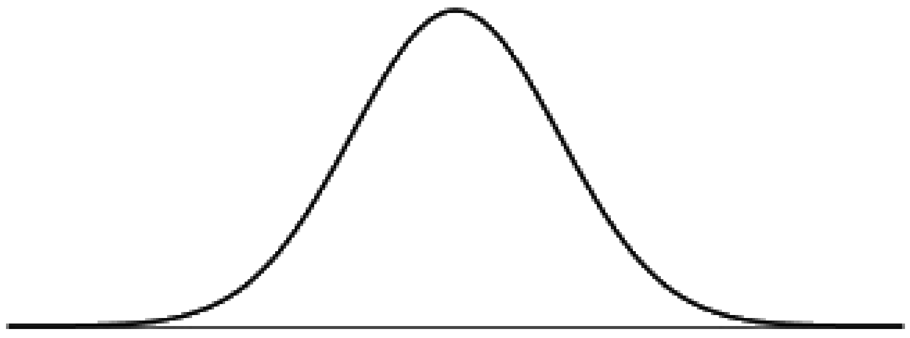

SDT model

- Sensitivity M_2- M_1

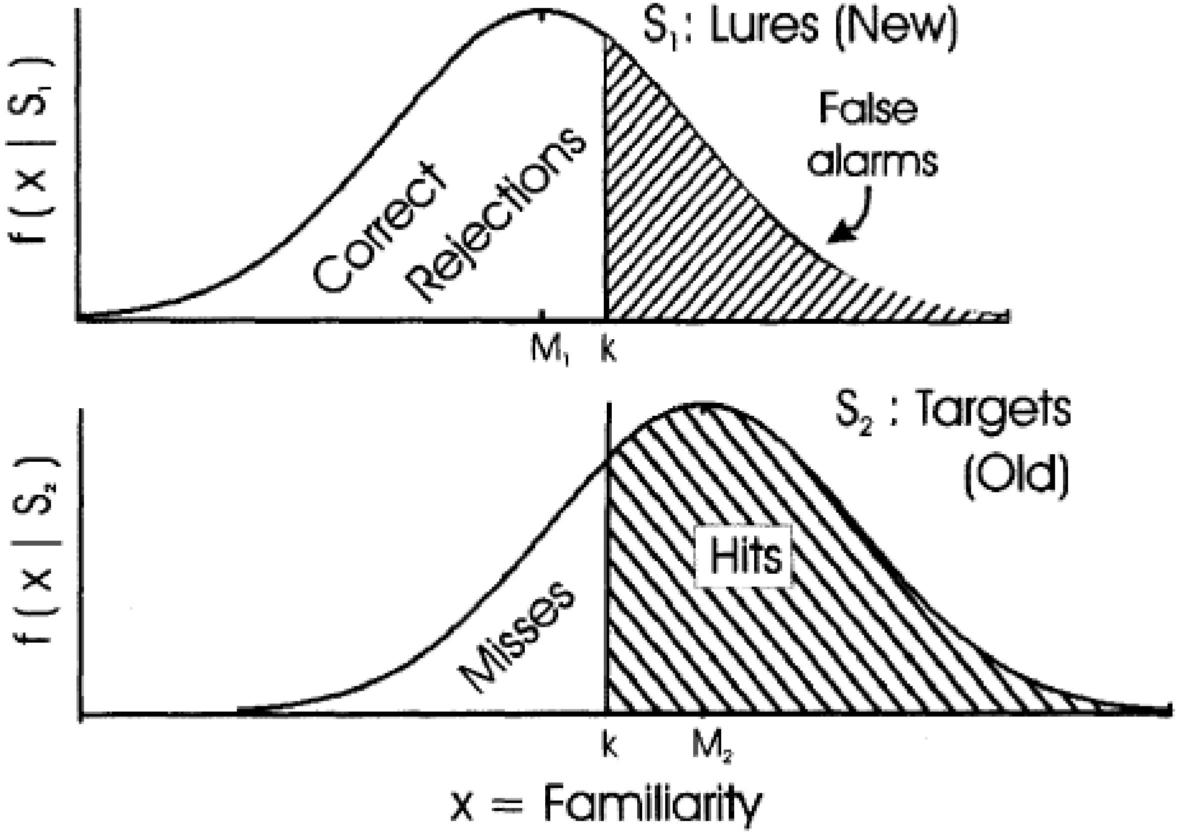

SDT model

- Sensitivity M_2- M_1

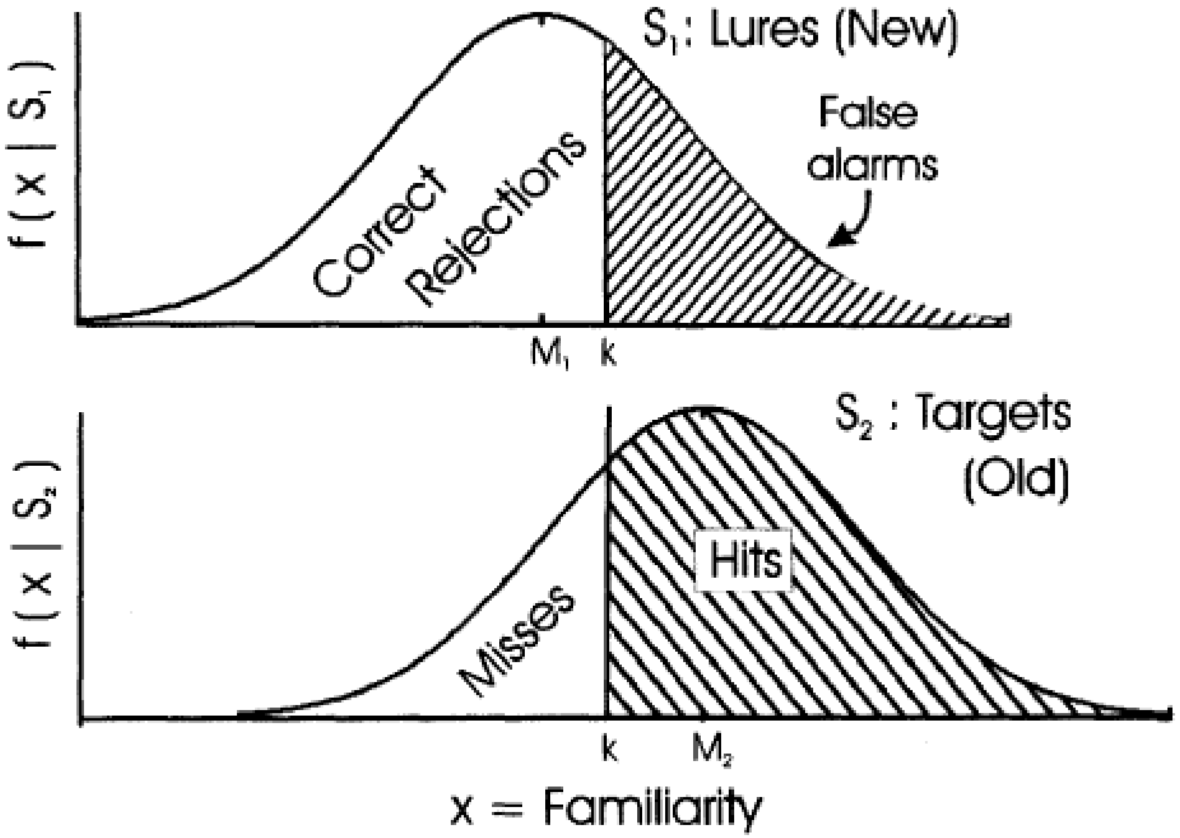

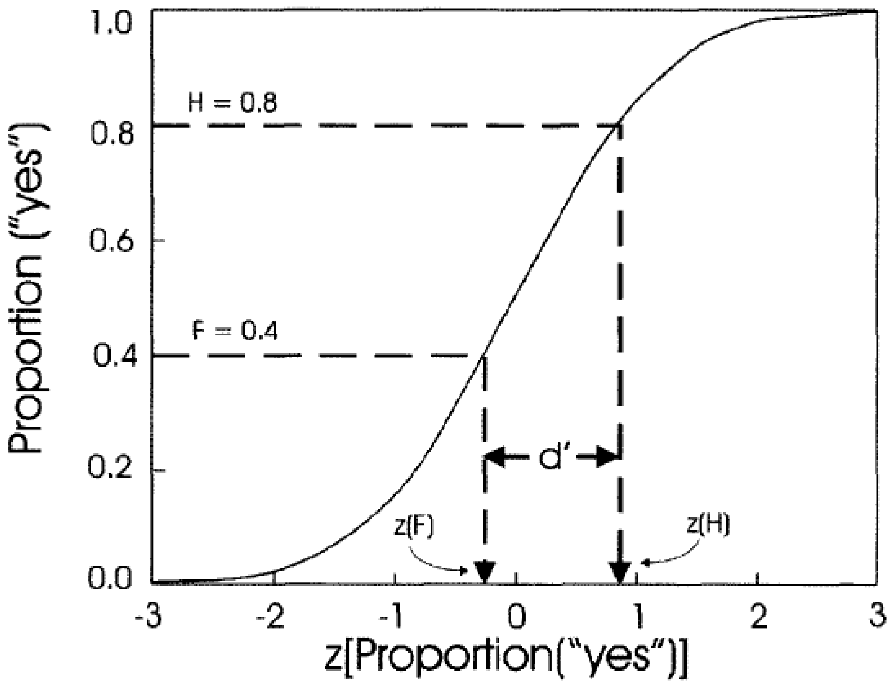

Covid again

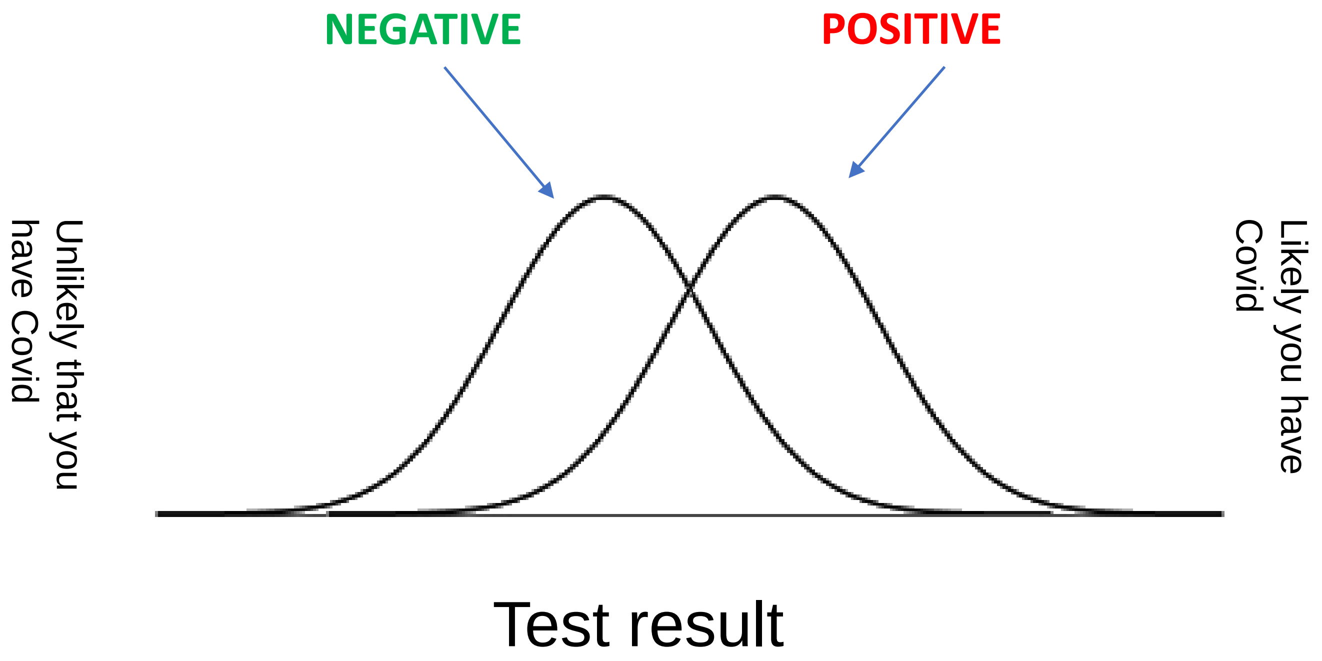
NEGATIVE
POSITIVE
Unlikely that you have Covid
Likely you have Covid
Test result

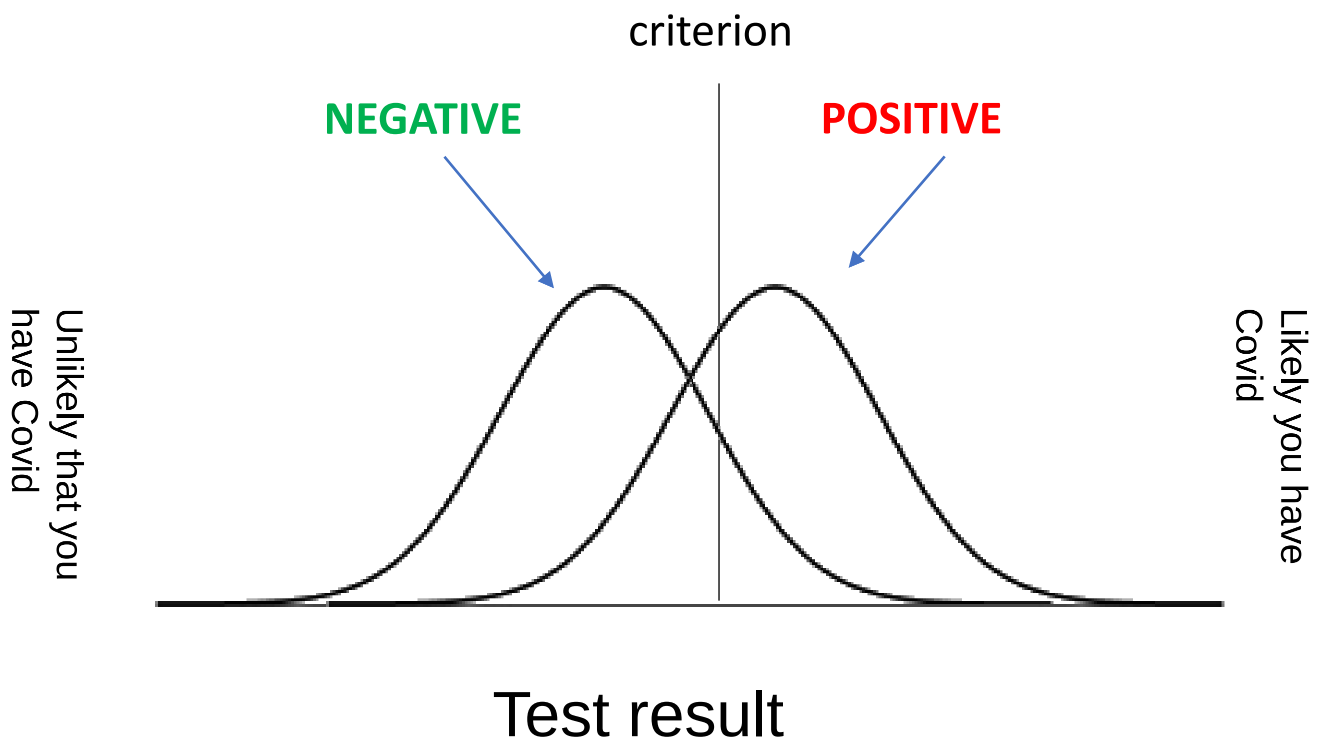
criterion
NEGATIVE
POSITIVE
Unlikely that you have Covid
Likely you have Covid
Test result

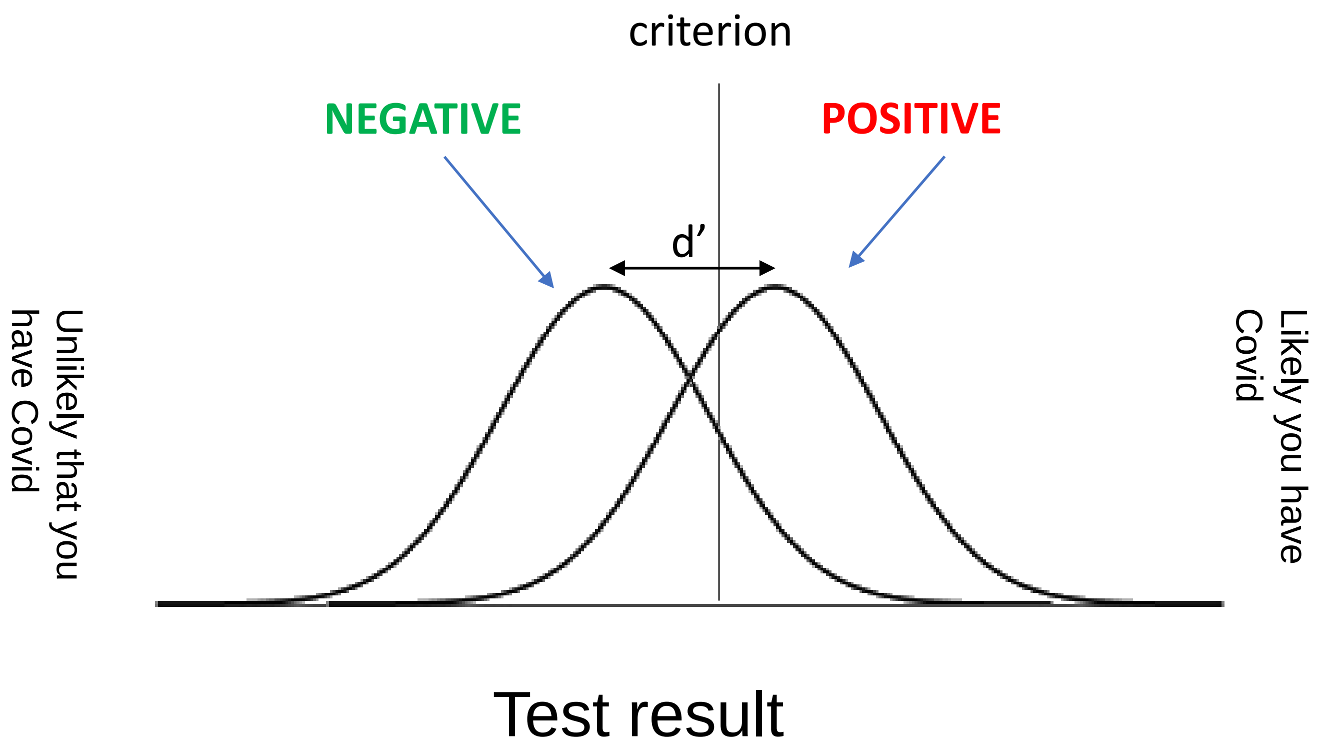
criterion
NEGATIVE
POSITIVE
d’
Unlikely that you have Covid
Likely you have Covid
Test result

Example with umbrellas

- https://r.tquant.eu/GrazApps/Group7_SignalDetection/

Bias

- Memory of faces normally and in hypnosis

	Normal		*Hypnotized*	
	"Yes"	"No"	"Yes"	"No"
Old	69	31	89	11
New	31	69	59	41

  - In both cases d' = 1, so only the bias changes
- Training for finding tumours in images
  - Before training d' is the same, after training both d' and bias are different

	Before Training	*After Training*
Trainee 1	$H = .89$	$H = .96$
	$F = .59$	$F = .39$
Trainee 2	$H = .89$	$H = .993$
	$F = .59$	$F = .68$
Trainee 3	$H = .89$	$H = .915$
	$F = .59$	$F = .265$

Three types of bias

- Bias should be some form of the sum of H and F
- Frequently used types:
  - c
  - c’
  - ln(β)

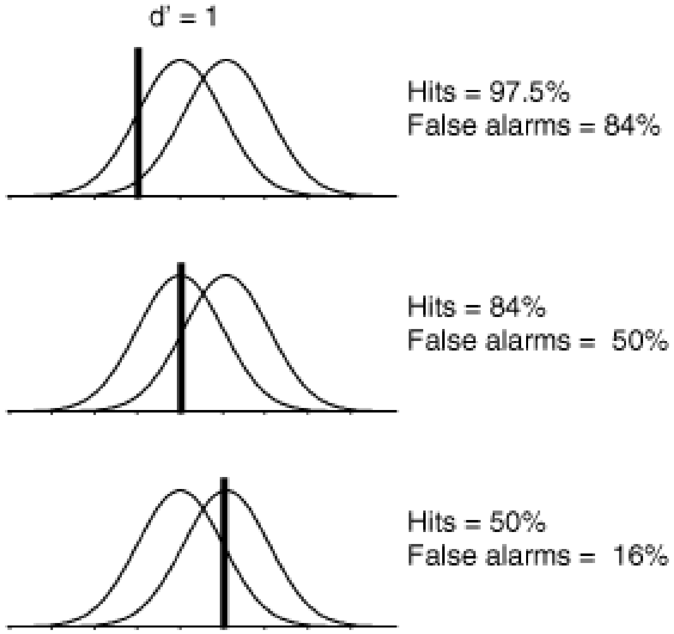

Criterion c

- c = -0.5(z(H) + z(F))
- At the extremes of H and F it is large, only reversed signs

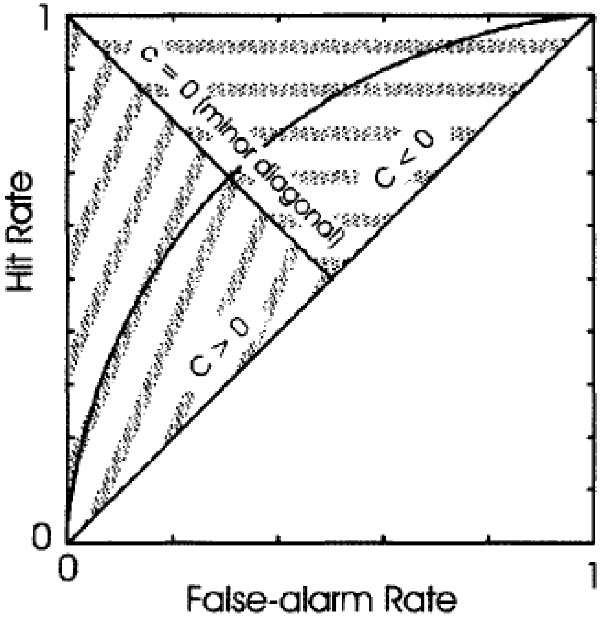

FIG. 2.1. The representation of criterion location in ROC space. Points in the shaded regions arise from criteria that are positive (below the minor diagonal) and negative (above the minor diagonal). Points in the unshaded region below the major diagonal result from negative sensitivity.

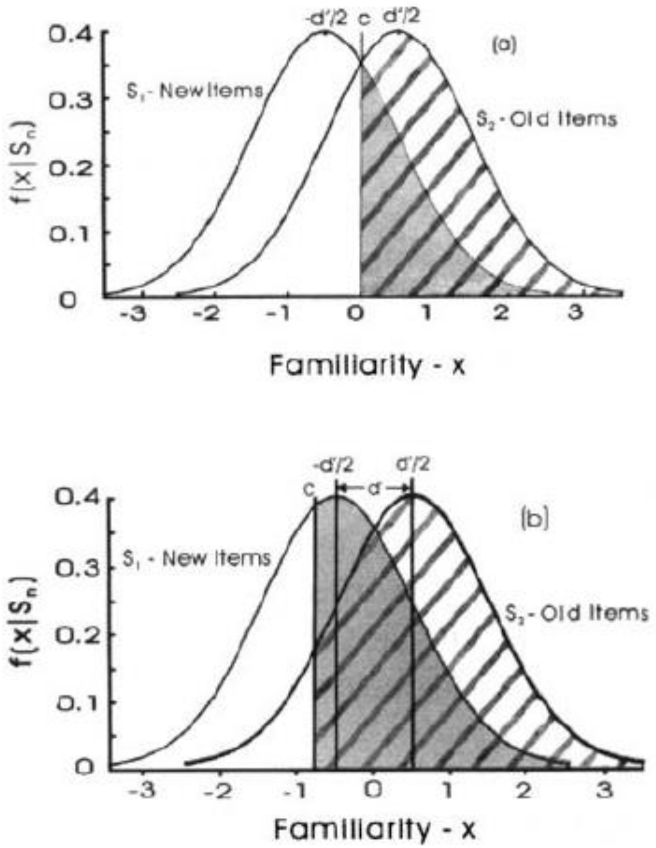

FIG. 2.2. Decision spaces for the Normal and Hypnotized conditions of Example 2a, according to SDT. Shaded area corresponds to *F*, diagonally striped area to *H*. (a) Normal controls have a symmetric criterion, $d' = 1.0$. (b) Hypnotized participants display identical sensitivity but a lower criterion, and thus have higher hit and false-alarm rates.

Bias - c’

- c is an absolute number, but there are other states with d'=0.5 and c=2 and d'=4 and c=2
- c’ = c/d’ = $\frac{(z(H)+z(F))}{2(z(H)-z(F))}$

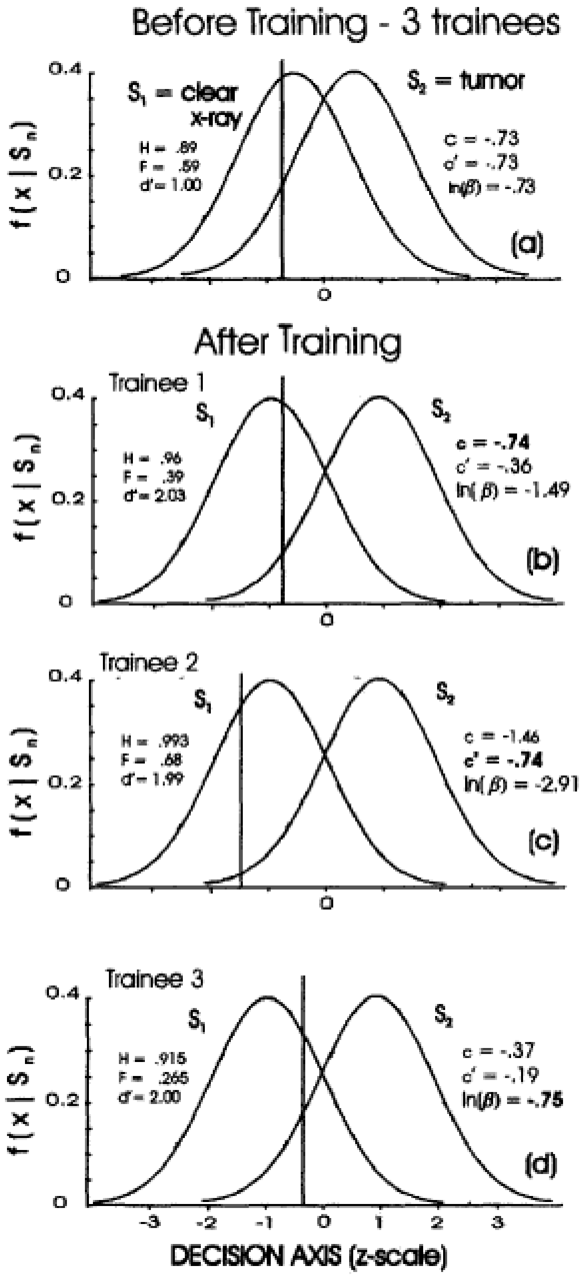

Likelihood ratio ln(β)

- The alternative is the likelihood ratio of the two distributions
- $LR(x) = f(x|S_2)/f(x|S_1)$

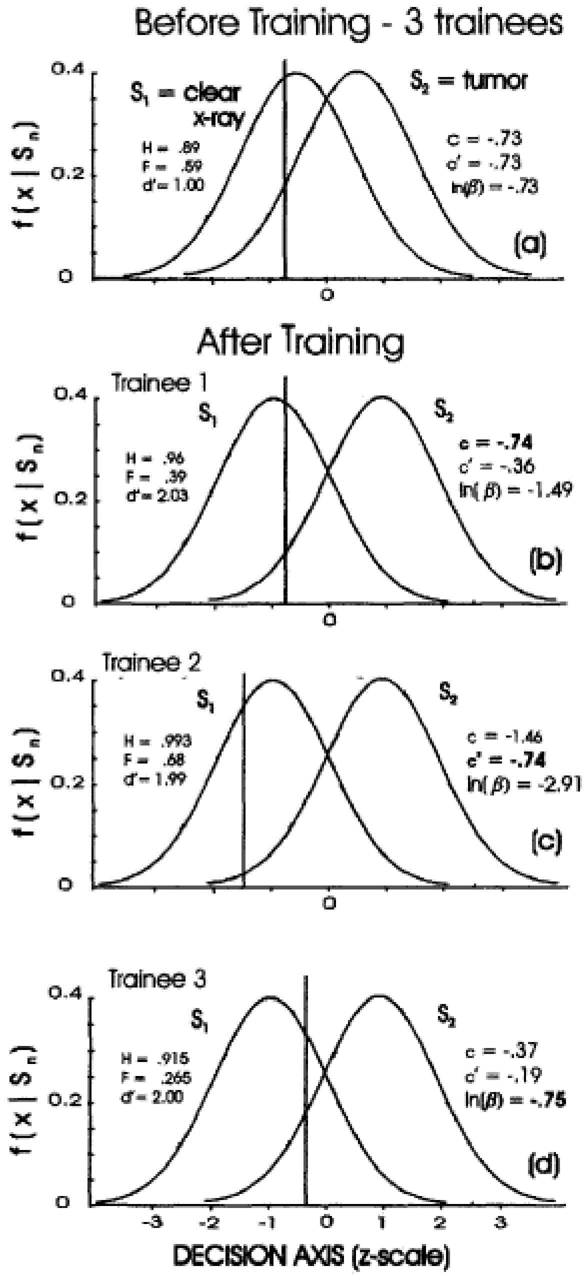

Likelihood ratio ln(β)

- ln(β) = cd’

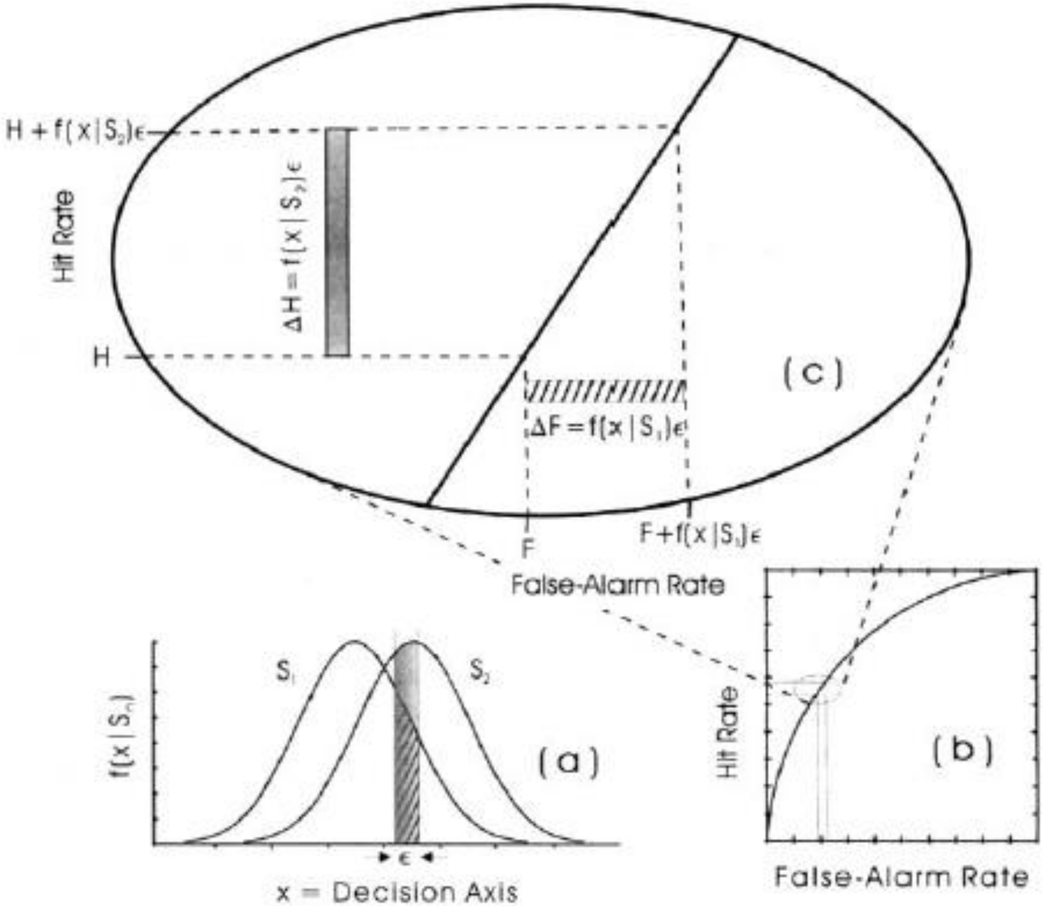

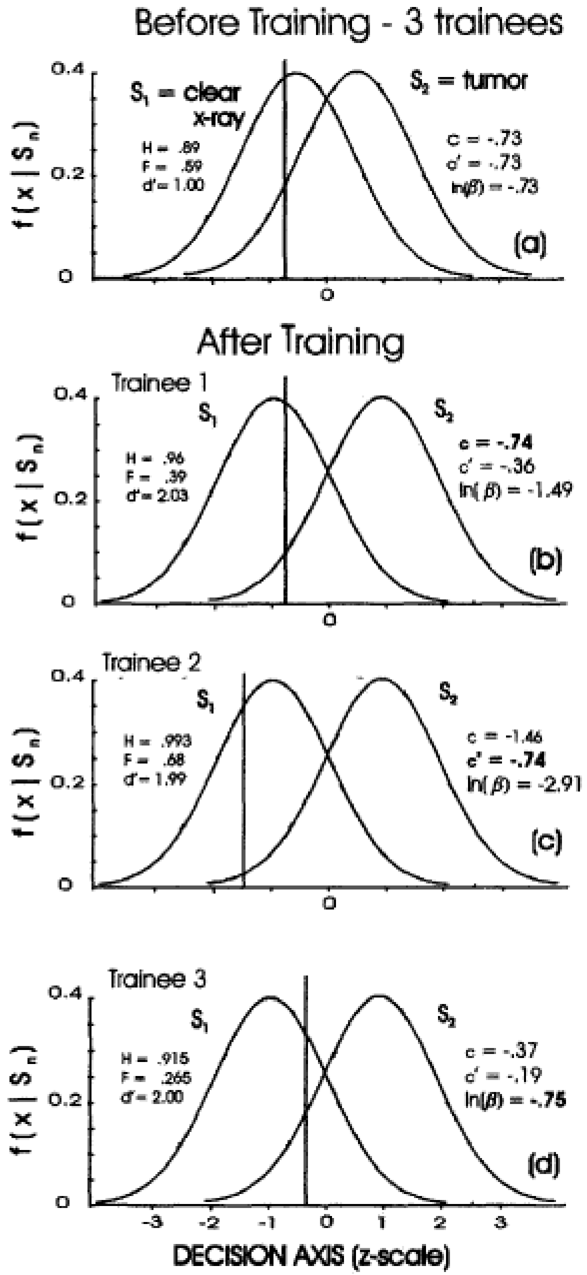

Isobias curves

- Again, we can express the relationship between z(H) and z(F): z(H) = -2c-z(F)

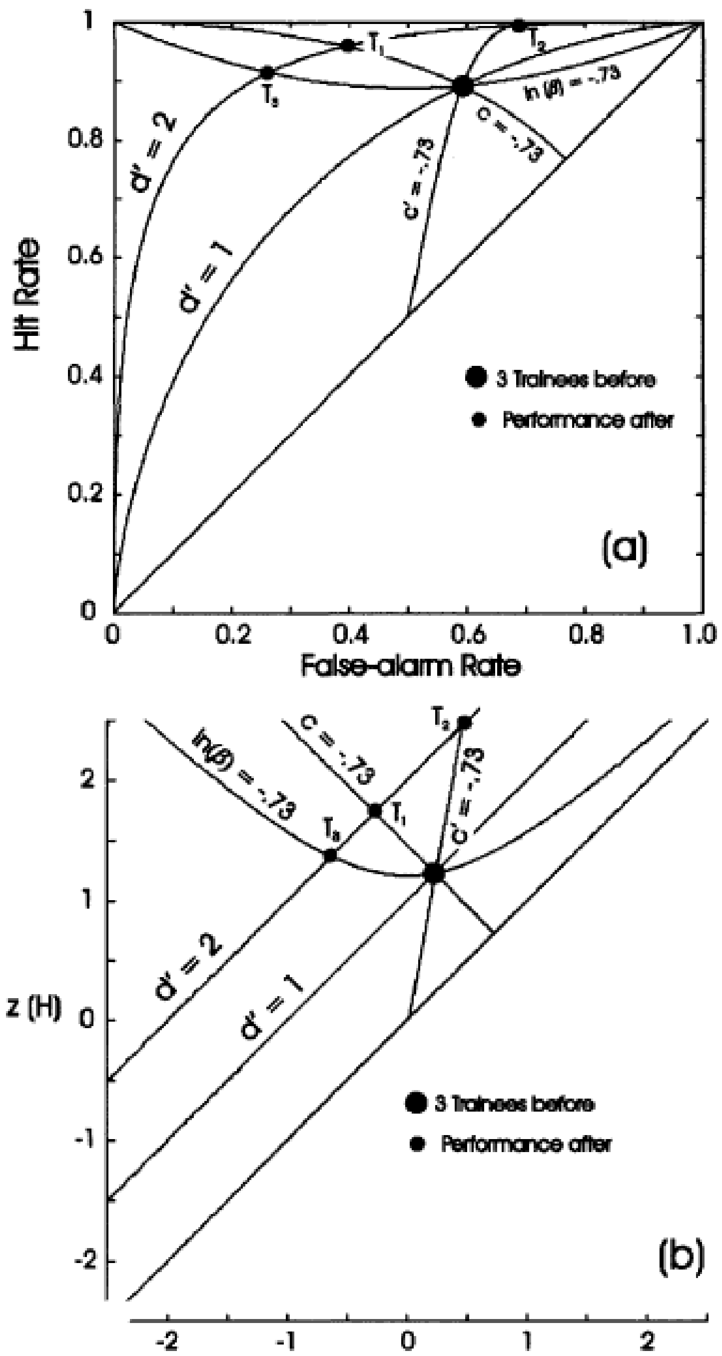

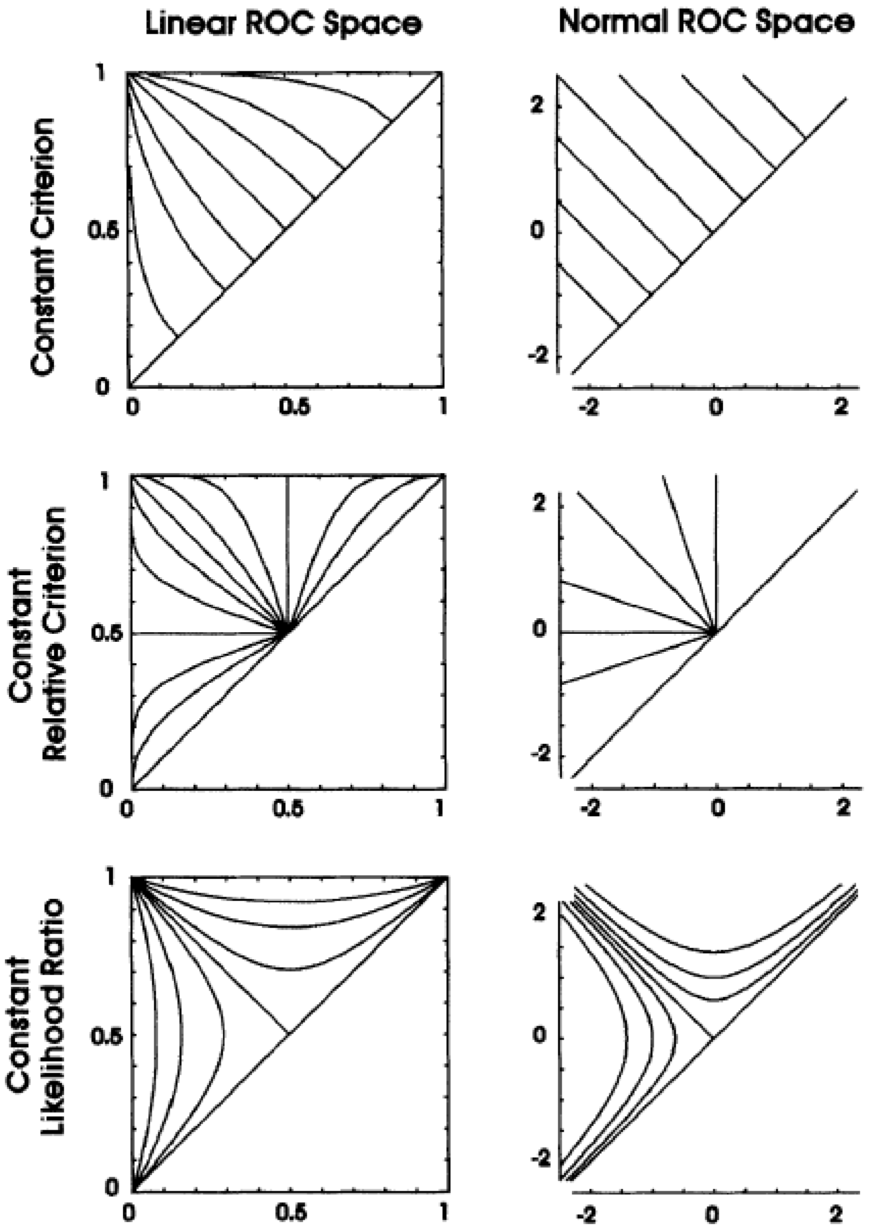

FIG. 2.6. Families of isobias curves (hit rate vs. false-alarm rate, d' varying) for constant criterion c, relative criterion c', and likelihood ratio β, on linear and z axes.

SDT for 2AFC

- Is 2AFC more difficult or easier than yes/no?

Data for 2AFC

- Two faces appear, one at the top, one at the bottom, one of which is old

	Responses		
Stimulus Sequences	*"Old on Top"*	*"Old on Bottom"*	*N*
<Old, New>	16	9	25
<New, Old>	7	18	25

- H = P(„old is above“|<new,new>)
- F = P(„old is above“|<new,new>)

Bias for 2AFC

- d’ = $\frac{1}{\sqrt{2}}$(z(H)-z(F))

- c = 0.5(z(H)+z(F))
  bias in one does not affect bias in the other

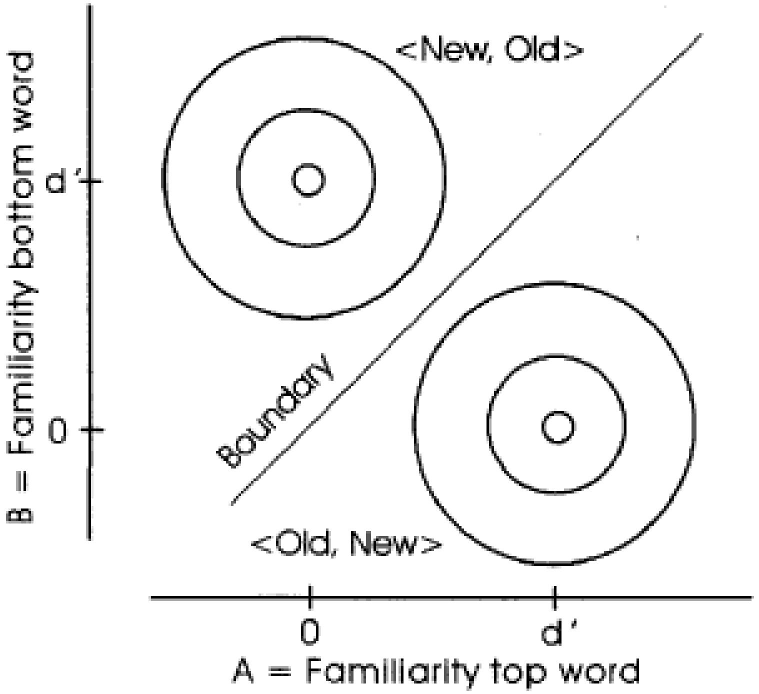

Maximum proportion correct

- Given d', what is the maximum p(c) at zero bias?

$$p(c)_{max,\ yes/no} = \phi(d'/2)$$

$$p(c)_{max,\ 2AFC} = \phi(d'/\sqrt{2})$$

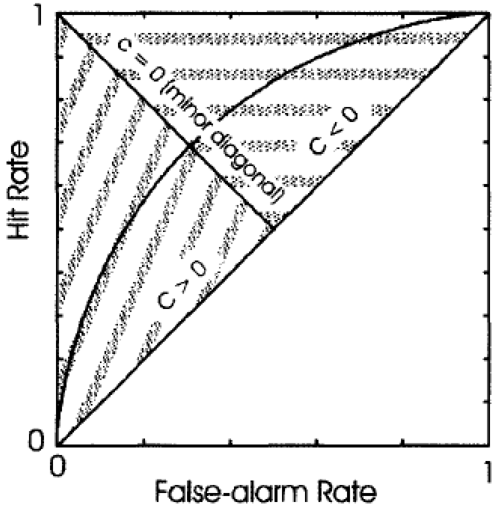

FIG. 2.1. The representation of criterion location in ROC space. Points in the shaded regions arise from criteria that are positive (below the minor diagonal) and negative (above the minor diagonal). Points in the unshaded region below the major diagonal result from negative sensitivity.

Allows us to use d' with psychometric curves

Correction for guessing

- For Yes/No

$$q = \frac{H - F}{1 - F}$$

- For 2AFC (H is p(c) and F is ratio of trials, in which guessing would lead to correct answer, so 0.5)

$$q_{2AFC} = \frac{p(c)_{2AFC} - 0.5}{1 - 0.5} = 2p(c)_{2AFC} - 1$$

Non-parametric sensitivity

- When we do not want to assume shape of ROC curve

$$A' = \frac{1}{2} + \frac{(H-F)(1+H-F)}{4H(1-F)} \quad \text{if } H \geq F$$

- Area under curve

$$A' = \frac{1}{2} - \frac{(F-H)(1+F-H)}{4F(1-H)} \quad \text{if } H \leq F$$

$$B'' = \frac{H(1-H) - F(1-F)}{H(1-H) + F(1-F)} \quad \text{if } H \geq F$$

- Bias

$$B'' = \frac{F(1-F) - H(1-H)}{H(1-H) + F(1-F)} \quad \text{if } H \leq F\ .$$

- But it misses theoretical grounds..