UCL

AUDL 4007

Auditory perception

(with a healthy dose of psychoacoustics …)

Course structure

- 8 sessions, a mixture of lectures, demonstrations, laboratory sessions and tutorials
- BSc assessment
  - 2 pieces of coursework, each worth 15% of the final mark (max 1000 words each)
    - Written essays presenting a published psychoacoustic study in a journalistic format, appropriate for lay readers.
    - You will also critique another student's paper, and rewrite from comments. Details to follow.
  - 2-hour written paper (70%)
    You must pass the final exam to pass the course.
- MSc assessment
  - 1 piece of coursework (max 1000 words)
    - A write-up of an empirical very small-scale psychoacoustic study
  - Written paper TBA
- Note: You will also be responsible for the information on cochlear implants that I present in the master class.

Readings

- Main text: Plack C. (2005) *The Sense of Hearing*. Erlbaum.
- Supplementary Reading
  - Yost, W.A. (2007) Fundamentals of Hearing: An Introduction, 5th ed.. Academic Press. A more elementary exposition. Particularly good on the anatomy & physiology.
  - Moore, B.C.J. (1997). An Introduction to the Psychology of Hearing, latest edition., Academic Press. A very complete guide to the literature, but at an advanced level.
- Papers on the web site
  - http://www.phon.ucl.ac.uk/courses/spsci/audper/index.html
  - http://tinyurl.com/y9xcpad

How to succeed in this course

- Attend the lectures
- Do the reading
- Check the web site
- Laboratory sessions should help to clarify the material presented
- Bring questions to the sessions
- Keep up with the work
- If you have problems, *ask for help!*

What is psychoacoustics?

- Psychophysics
  - Mapping the relationship between the physical/objective and perceptual/subjective world.
- Psychoacoustics — psychophysics of sounds.
- How does the loudness of a sound relates to its intensity?
  - loudness depends not only on intensity but also on frequency content
- Changing the fundamental frequency of a periodic sound from 100 to 200 Hz will not lead to the same perceived musical interval as a change from 800 Hz to 900 Hz.

100-200 Hz 800-900 Hz 800-1600Hz

What is psychoacoustics?

- Terminology: Objective vs. subjective
  - intensity (W/m^2, Pa, dB SPL) vs. loudness
  - periodic/aperiodic vs. buzziness/noisiness
  - fundamental frequency (Hz) vs. pitch
  - spectral envelope/shape vs. timbre/quality/colour
- Much of psychoacoustics concerns abilities to …
  - detect
    - many HI people and CI users need higher levels to detect sounds
  - discriminate
    - many HI people and CI users need greater differences between stimuli to hear a difference between them
  - but limits on detectability and discriminability can also provide crucial data for developing models of auditory perception even in normal listeners

Gap detection
A fairly typical psychoacoustic task

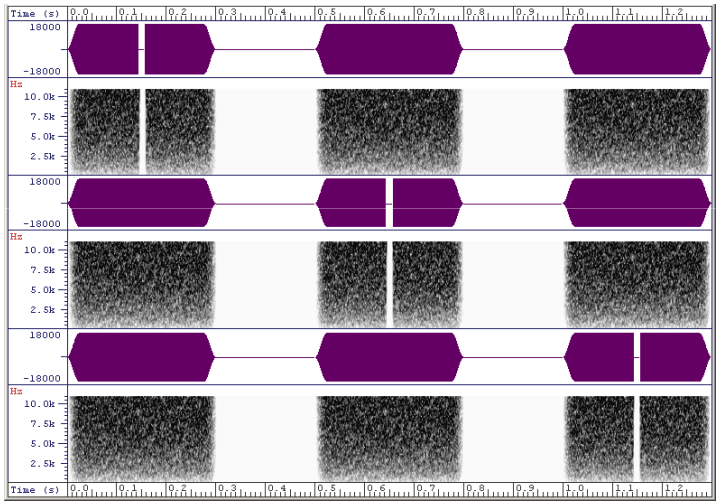

time →

Gap detection

- Pick the sound with the gap – vary the gap duration to find threshold
  - when a listener is 'doing well', make it harder
  - when a listener is 'doing poorly', make it easier
  - What does this remind you of?
    - adaptive procedure
- Thresholds for wide-band noise are around 3 ms

Basic procedures: Psychoacoustic procedures vary in …

- Number of intervals (or stimuli presented) per trial
  - Typically 1, 2, 3 or 4
- What the listener is asked to do
  - detect the presence of a stimulus (absolute threshold)
  - detect a change in a stimulus (discrimination)
  - label a stimulus
  - label the direction of change in a pair of stimuli
- How the stimulus levels are controlled
  - depending on, or independent of, the listener's responses

Number of intervals/stimuli

- 1 interval
  - detection (yes/no) - Is it there?
  - identification/categorization - What is it?
  - scaling - How much of it is there?
- 2 interval
  - detection - When is it?
  - discrimination - Which one is _____-er? (louder, higher, buzzier, smoother, etc)
  - discrimination – are the two sounds same or different (AX – AA vs AB)
  - scaling - How are the two signals related (e.g., what musical ratio are they in?)

Number of intervals/stimuli

- 3 interval
  - detection - When is it?
  - discrimination - Which one is the same? (ABX – ABA vs ABB) Which one is different? (BAA vs ABA vs AAB).
  - triadic comparisons - Which two are most similar? (or different)
- 4 interval
  - especially useful when the perceptual difference is hard to describe
  - discrimination of anything – which of two pairs has different sounds? (4IAX: AB-AA vs AA-AB; or AABA vs ABAA – is the 2nd or 3rd sound the odd one out?)
- variable-interval
  - method of adjustment - adjust one sound to be as *loud, high, etc.* as another

Role of bias

- Perhaps especially strong in 1-interval yes/no experiments
  - some people are much more likely to say yes for a given level of perceptual evidence; others need a lot of convincing!
- Less strong in 2-interval forced choice
  - no obvious reason to prefer 1st or 2nd interval

Psychometric functions : 1 interval
How to define a threshold

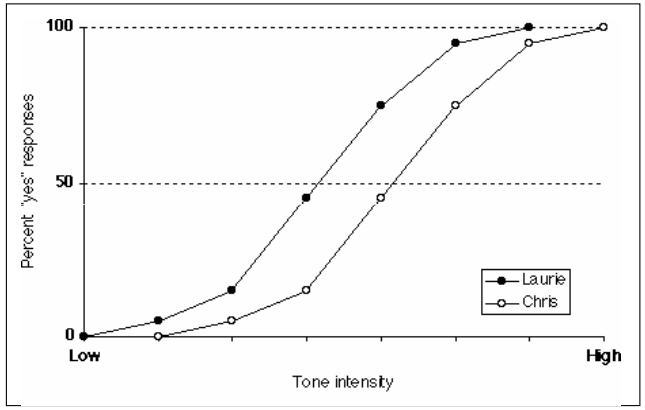

Reiz (stimulus) *Limen (RL),* the absolute threshold (the weakest stimulus that is just detectable)

Speech categorisation function

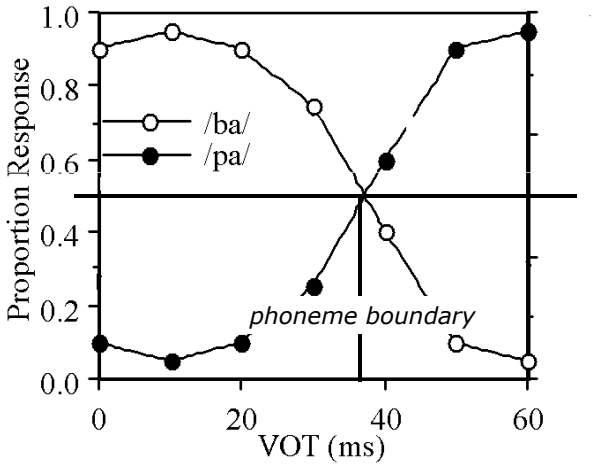

Psychometric functions : 2 intervals
How to define a *jnd*

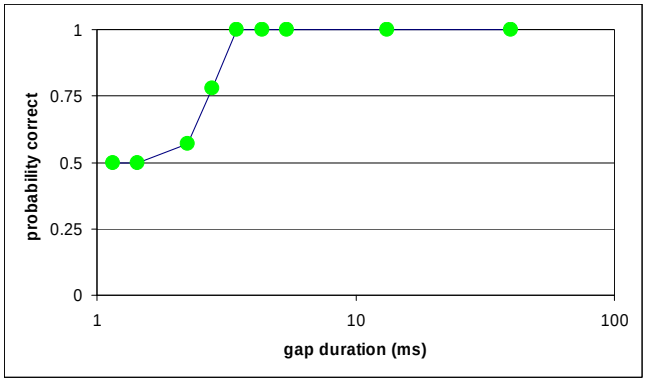

just-noticeable difference (jnd) or DL (difference limen), the smallest stimulus increment that is just detectable (from the German *Differenz Limen*)

Control of stimulus levels

- Method of constant stimuli
  - stimuli presented do not depend on listener responses
  - traditional, and still commonly used
- Adaptive procedures
  - stimuli depend on the responses of the listener
  - try to concentrate stimulus presentations in region of most interest (*e.g.*, near 75% correct in 2I-2AFC)
  - most commonly used nowadays

An adaptive track

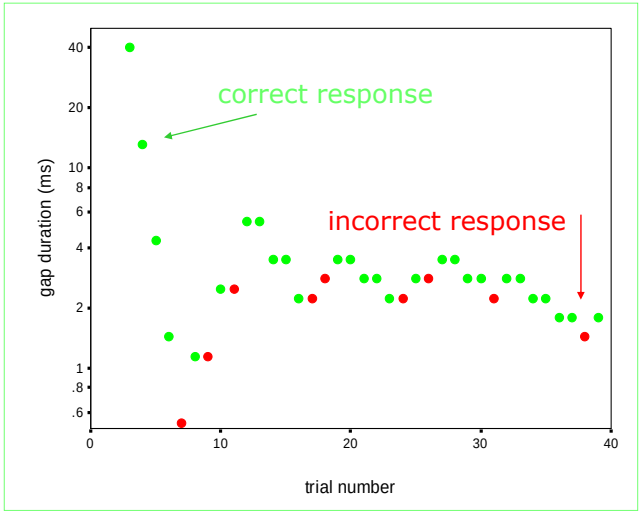

Adaptive procedures

- Staircase
  - 1-up/1-down, especially common in speech intelligibility testing
  - modified into Up-Down Transformed Methods
    - 1/1 tracks 50% correct
    - 1/2 tracks 71% correct
    - 1/3 tracks 79% correct
  - Parameter Estimation by Sequential Testing (PEST)
    - complex rules for changing step-size and deciding when to stop, based on explicit statistical calculations
  - Robust to lapses of attention, but can converge slowly
- Maximum Likelihood
  - use statistical methods to place stimuli *exactly* at the best estimate of the desired point on the psychometric function
  - Move fast but highly sensitive to lapses of attention

Signal Detection Theory

all judgments are made in the presence of uncertainty

The simplest case: 1-interval yes/no

- To determine an absolute threshold.
- During a pre-determined time interval, either present a 125 Hz tone burst, or not.
- Listener responds either …
  - yes, I heard something
  - no, I didn't hear anything
- Sound familiar?

A psychometric function only for trials when a signal is presented, but at various levels

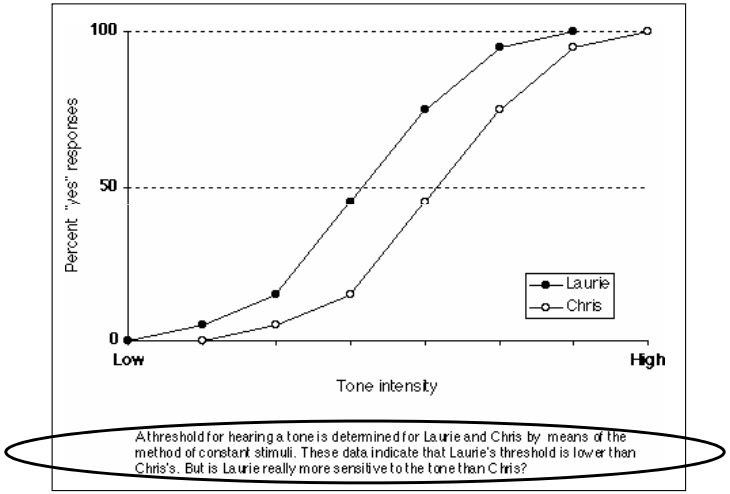

Or just more willing to *say* she hears one?

Terminology

		signal	
		present	absent
decision	present	hit	false alarm
	absent	miss	correct rejection

Knowing the proportion of 'hits' and 'false alarms is sufficient. Why?

Mapping the physical stimulus into the psychological world

- Each interval has 'noise' in it, internal or external so …
- the psychological 'effect' of the presented stimulus varies randomly from trial to trial.
- Sometimes, you think you hear a tone when none was presented, and …
- you don't hear it when one was.

A simple model of a tone detection experiment

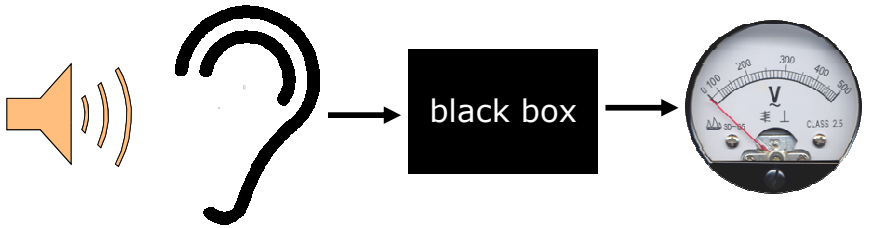

Enter the homunculus

apologies to Jules Feiffer

Present 125 Hz sinusoid at 30 dB SPL

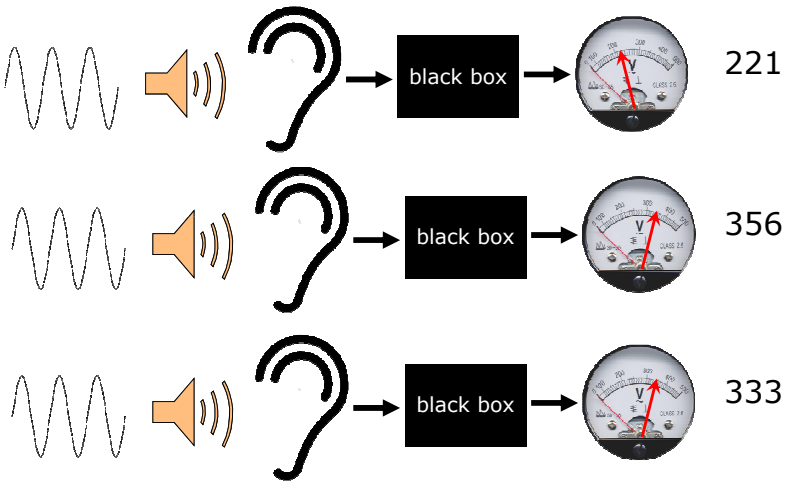

But what happens with no input?

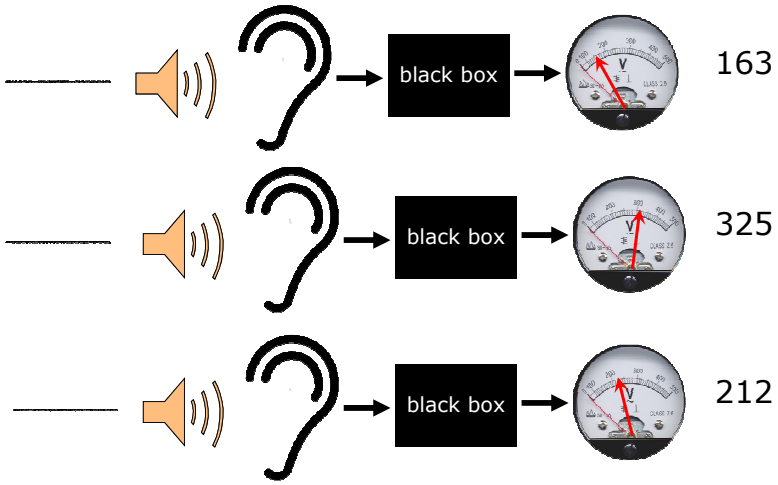

Why do the values vary?

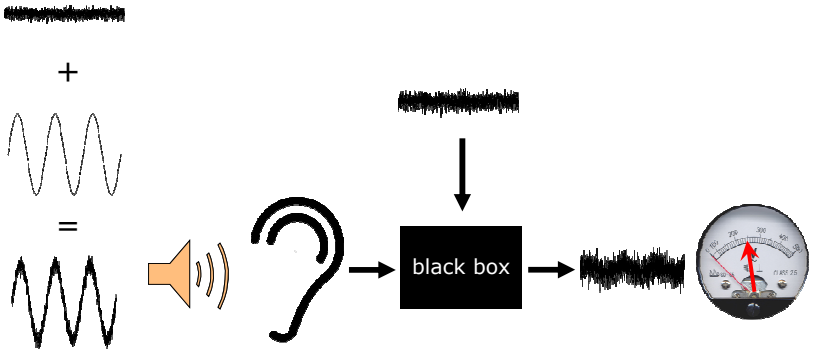

Must consider a *distribution* of values

- The meter reading varies from trial to trial, even when the experimenter does exactly the same thing …
- So the meter reading can be thought of as a *random variable*, which can be described by its *distribution*.
  - normally assumed to be Gaussian (*i.e.*, bell-shaped)

Internal response distributions

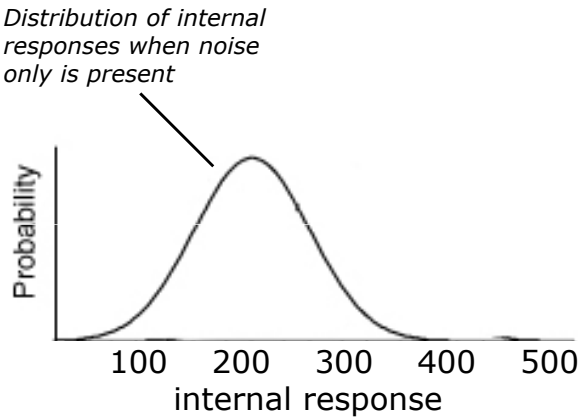

Internal response distributions

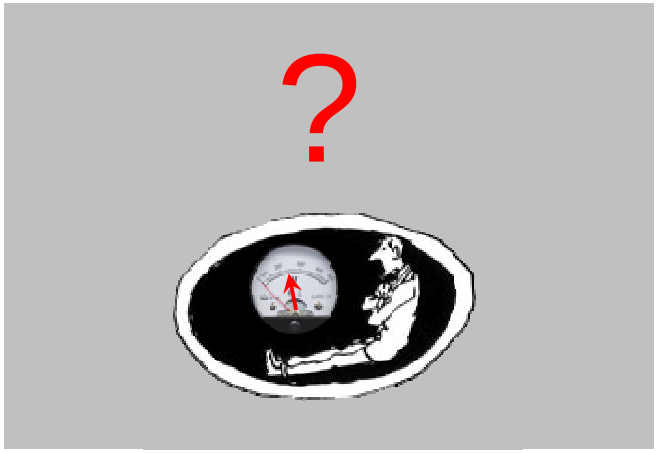

Generally speaking, the bigger the meter reading, the more likely there is to be a signal but …
Given any particular value of the internal response, how do you decide to say 'yes' or 'no'?

Need to set a criterion

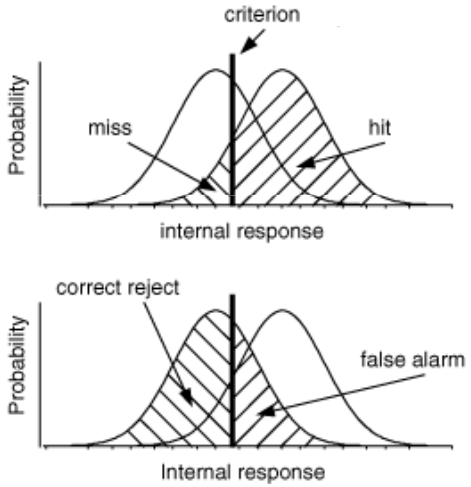

Shifting the criterion: You can't improve your hit rate without also increasing the number of false alarms

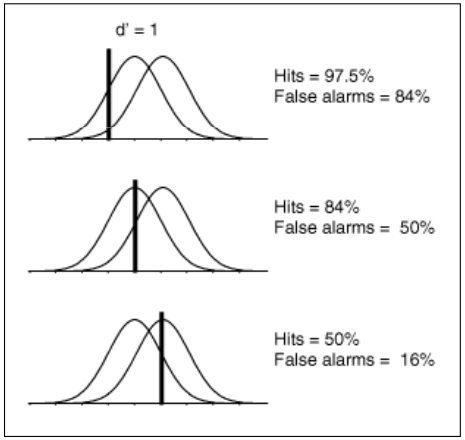

Performance determined by the distance between the distributions, and …

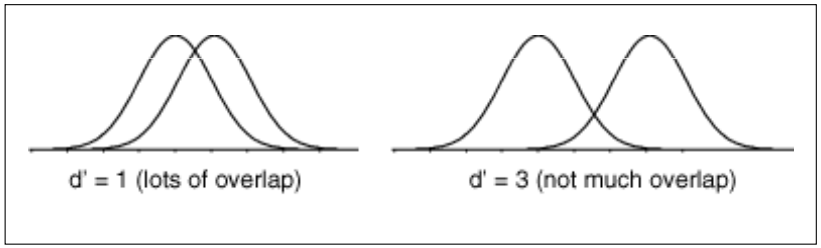

… by the variance of each

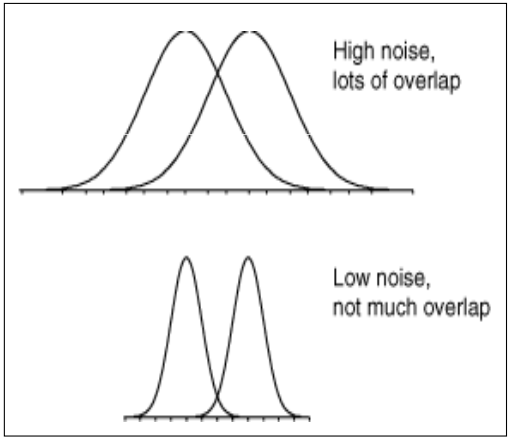

Discriminability index (*d'*):

d' = separation of the distributions / spread of the distributions

$d' = (\mu_2 - \mu_1)/$ s.d.

Related to the statistic known as ***effect size:***
Unlike significance tests,
effect size is independent of sample size, so does not suffer the 'problem' that a tiny difference between samples can be highly significant if sample sizes are sufficiently large.

ROC curves show how shifting the criterion affects judgements

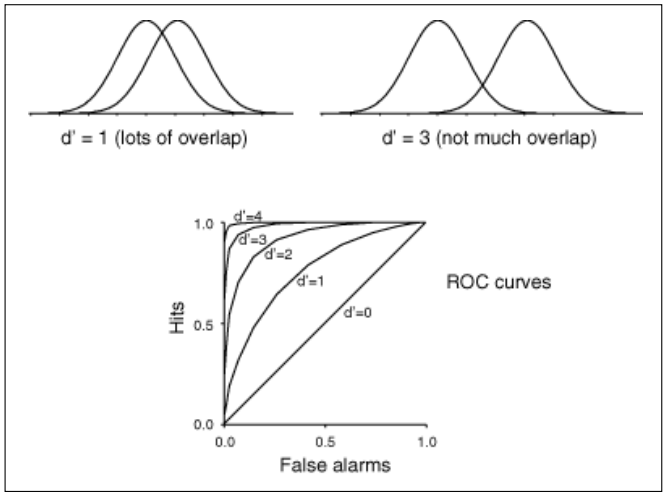

This is what we know: ROC curve

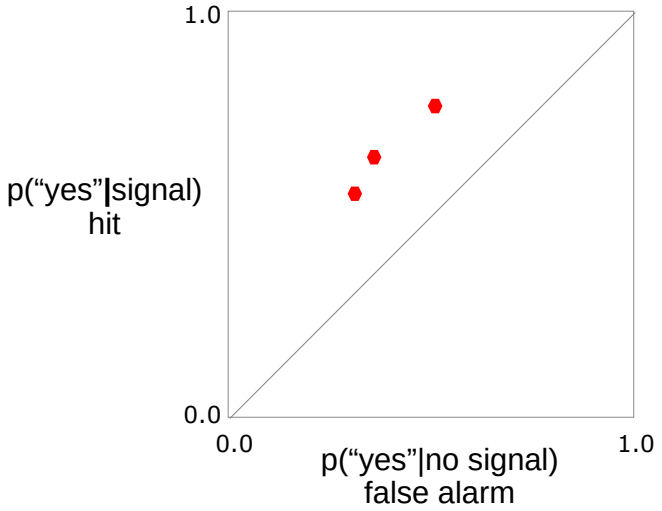

Internal response distributions

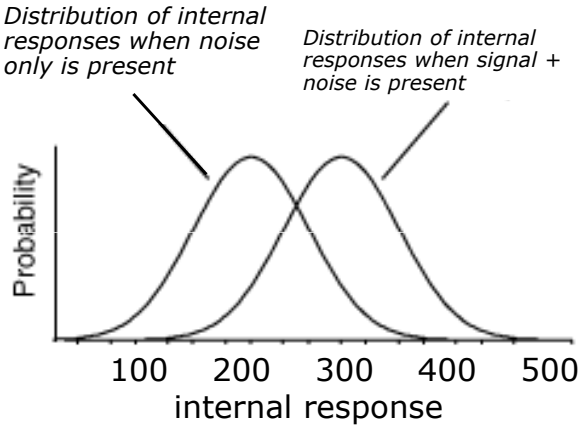

This is what we don't know

We use what we know to infer what we don't

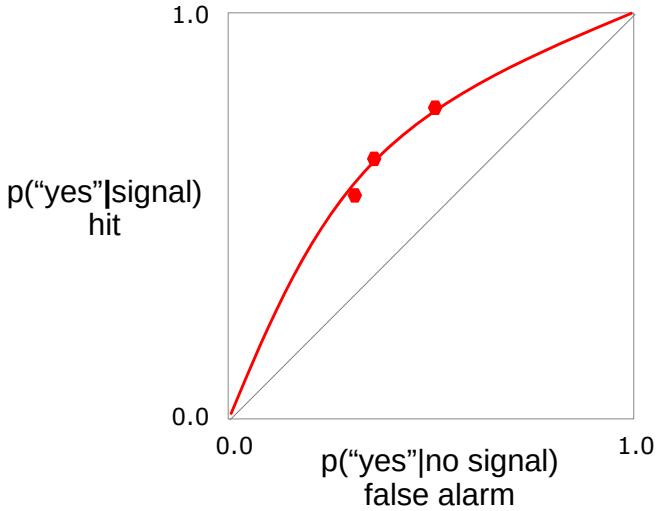

Signal Detection Theory has often been used to account for the variation in thresholds found in different paradigms.

Most research now uses a relatively small subset of procedures to minimise bias.