Sensory Awareness & Perceptual Certainty

Hakwan Lau
Columbia Psychology

This is a difficult talk. In 35 mins it is hard to focus on both the empirical details and the conceptual issue. So I've chosen to focus on the latter.

In philosophy…

- Access vs Phenomenality



- Easier vs Harder problems



In philosophy we distinguish between two kinds of concepts for consciousness. One more about basic information processing. One more about the subjective experience.

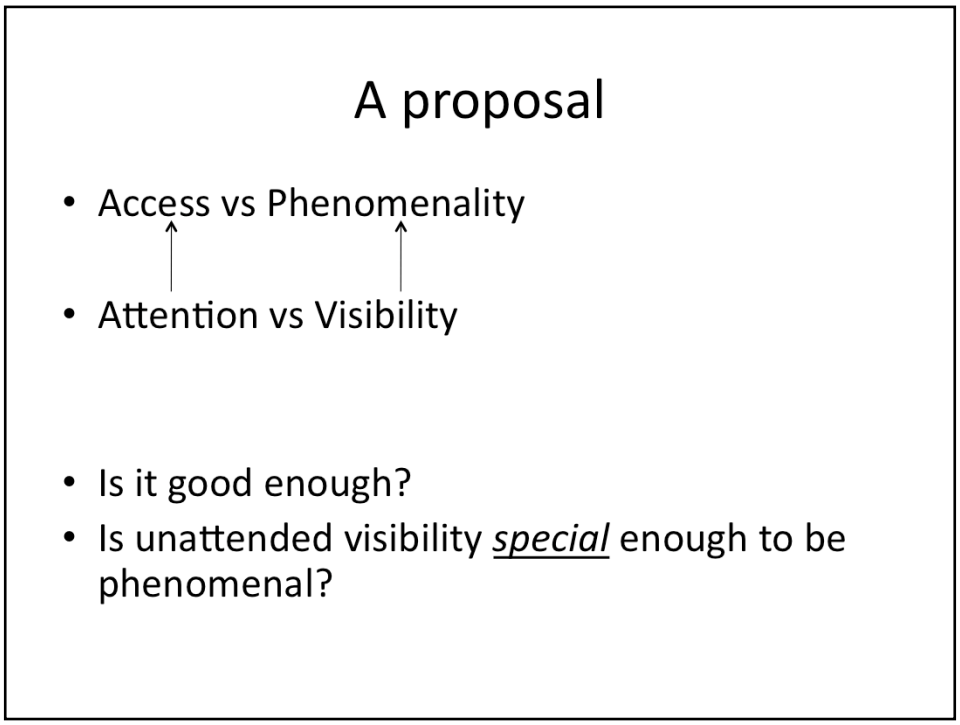
In science…

- Ignore the distinction.. (standard Workspace)
  - consciousness = full attention conscious report
  - we are not conscious of the "gorilla"
  - consciousness ~= whole brain
- Attention vs Visibility
  - prefrontal vs posterior

In science, there are two options.

One is to ignore the distinction. This has been popular up until about 5 years ago actually: All there is to consciousness is conscious access. The "standard" global workspace people basically take this view. So in inattentional blindness, when subjects didn't notice something so salient as a walking gorilla right in front of their eyes because of paying attention to a different task, we would say that the subjects have no conscious perception of the gorilla. For this kind of consciousness, pretty much the whole brain is relevant (visual signal in the back + fronto-parietal)

Another approach has gained popularity in recent years. People realize it is useful to distinguish between attention and visibility. Psychophysically, sometimes attention and visibilty have opposite effects. And the former maps more to the prefrontal-parietal areas, whereas the latter seems to be supported by posterior visual activity.



So Ned Block has proposed that what scientists call “attention” (prefrontal stuff) may correspond to the “access” aspect of consciousness. And that “visibility” may correspond to “phenomenality”

I think this is great. Definitely better than ignoring the philosophical distinction. But is it good enough? I think what scientists call “visibility” isn’t really what we want to equate with “phenomenality”

What is visibility?

- Masked vs Unmasked
- (potential) Hits vs Misses
- …
- Strong vs Weak signal (performance capacity)

- too *Zombie*… not HARD enough…

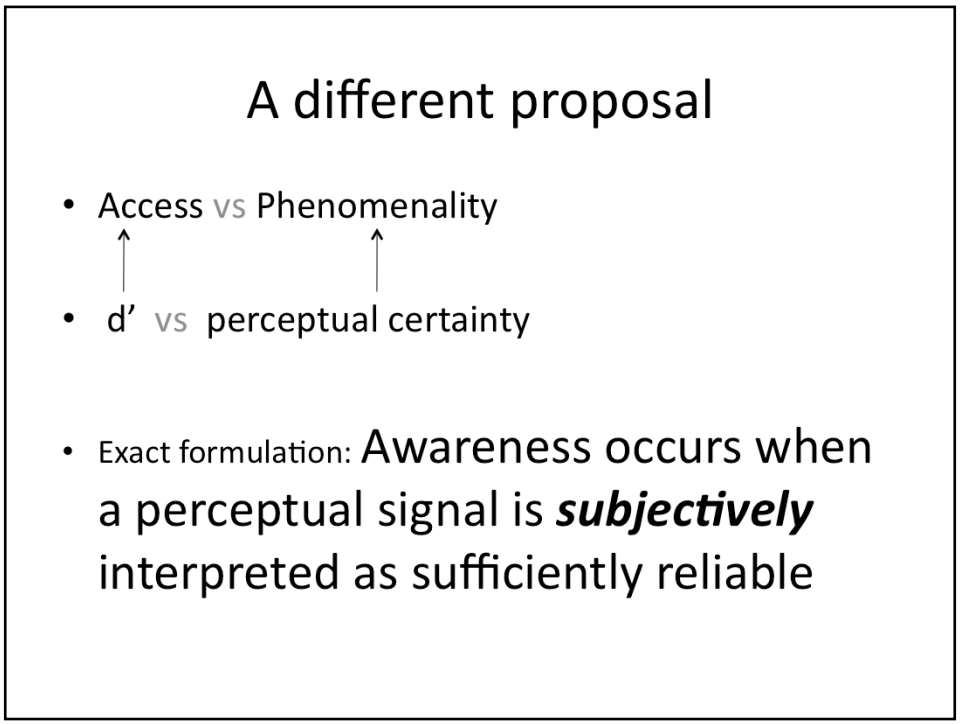
What is visibility anyway?

Typically, in experiments, that means e.g. comparing visually maked vs unmasked stimuli. Or that stimuli that can lead to a potential Hit (i.e. detection of target) compared against Misses (i.e. failure of detection).

Fair enough, if the visual target remained the same, when you sometimes can see it, and sometimes not. It has probably got something to do with your brain. Maybe that's subjective enough?

But really, one way to look at it is that visibility is just the same as signal, or performance capacity. When visibility is high, it just means that the target is yielding a lot of information in your head.

That doesn't sound special enough to warrant the label "phenomenal". Blindsight patients have good signal. Actually, photodiodes too.

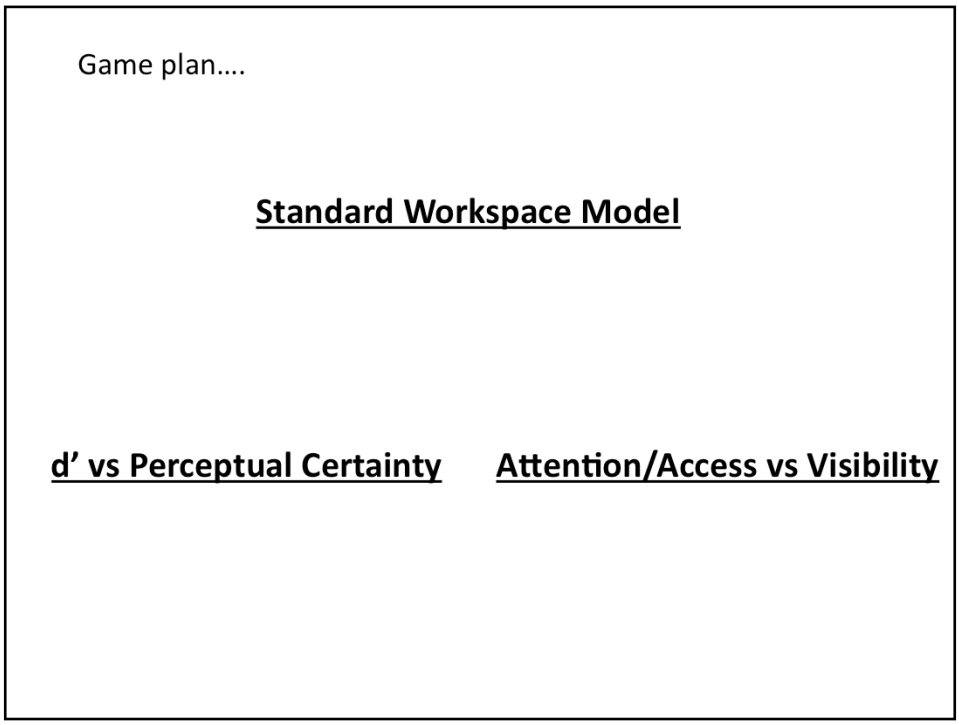


So I want to offer a slightly different proposal

We can distinguish between d' and perceptual certainty. And these two would map to the informational and subjective aspects of consciousness pretty well.

The exact formulation is as follows… (see slide)

D' is just performance capacity (formally it is a measure of the signal-to-noise ratio in processing regarding a stimulus). When you have a good d', you can make a probabilistic judgment of how good the d' really is. All probabilistic judgments are, of course, subjective. I will expand on this idea for the remaining of the talk…



So the talk of the plan really is this: I compare these 3 approaches to empirically studying consciousness.

The standard workspace folks conflate the informational and experiential aspects of consciousness. So they will receive little attention today.

The focus really is on comparing my approach (d' vs perceptual certainty) against the recently popular approach (attention vs visibility) which has people like Ned as proponents.

- **d'** = performance capacity / signal strength
  ~ % correct in forced-choice tasks
  – ***objective***

- **Perceptual Certainty**
  ~ confidence judgments
  ~ visibility ratings
  ~ tendency to say "yes I see it clearly" (detection bias)
  – ***subjective***

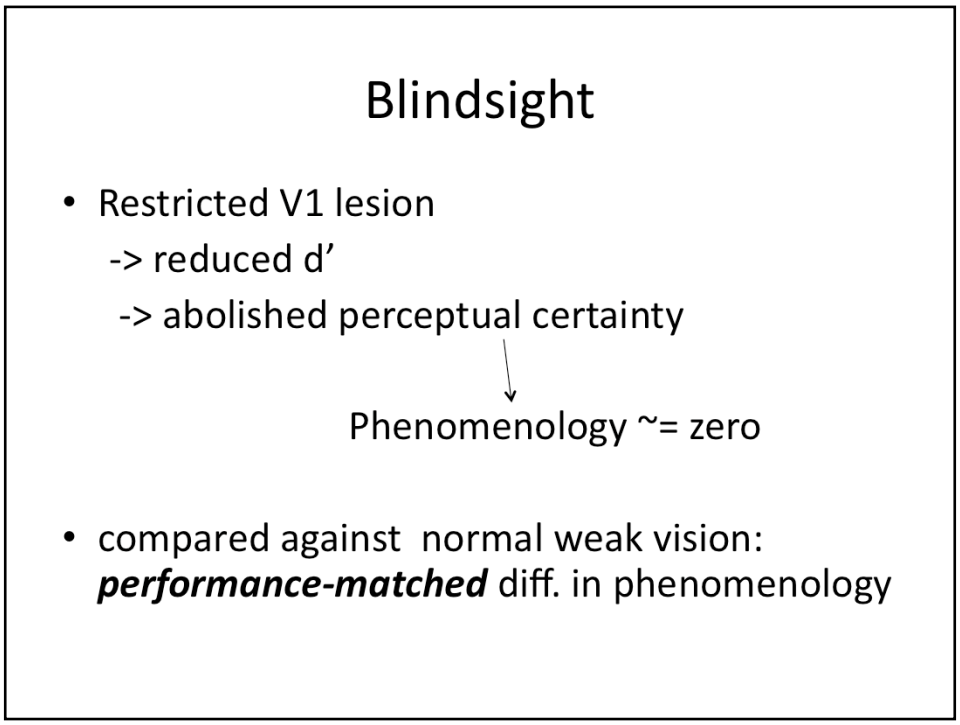
- they correlate, but not always…

In my approach, I distinguish between d' and perceptual certainty.

D' is just a measure of performance capacity. It is similar to the % accuracy in forced-choice tasks. It is what I call an OBJECTIVE measure. It measures how good the signal ***is***.

Perceptual certainty can be reflected by things like confidence judgments, visibility ratings, or the tendency to say "yes I see it clearly", etc. There are subtle differences between these measures, but they are subtle, so I try to lump them together for simplicity in this talk. These are all what I call SUBJECTIVE measures. They measure how good one ***thinks*** the signal is.

Now then, how good the signal is, is often correlated with how good you think the signal is, assuming you're thinking straight. But they don't ALWAYS correlate…



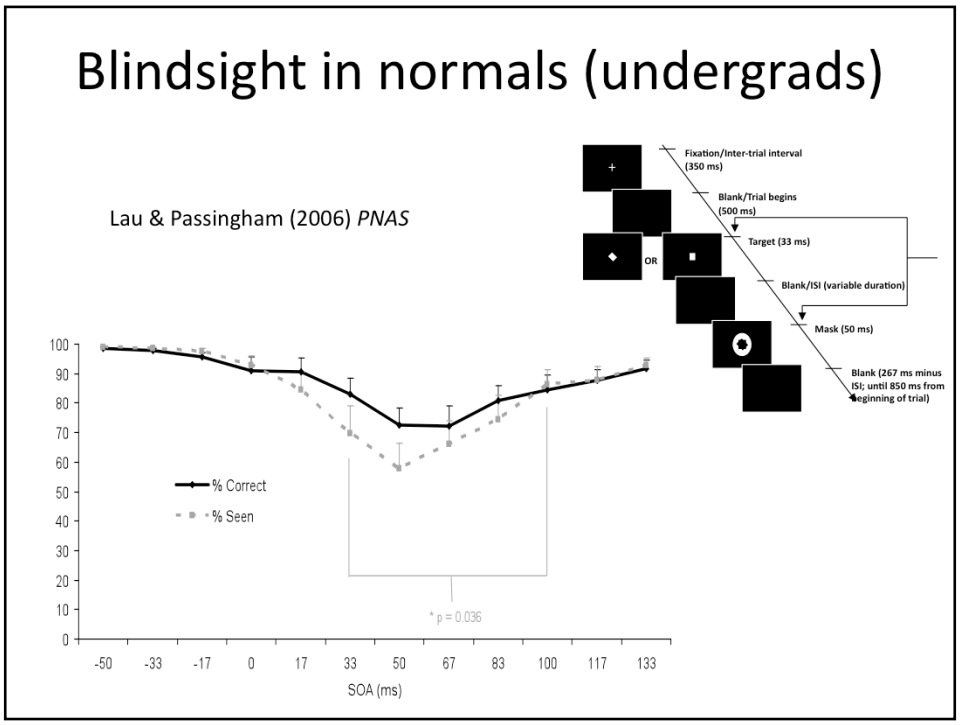
Patients with restricted v1 lesion may develop blindsight

Blindsight subjects have reduced but well above chance d'.

But they are almost completely uncertain about their percepts. They claim they don't see anything

One can send a strong stimulus to the blindfield. And send a weak stimulus to the good field (in some patients only part of the visual field is affected). This way, you can have performance-matched (i.e. d' matched) conditions where the levels of perceptual certainty differ dramatically

So, d' and perceptual certainty can dissociate

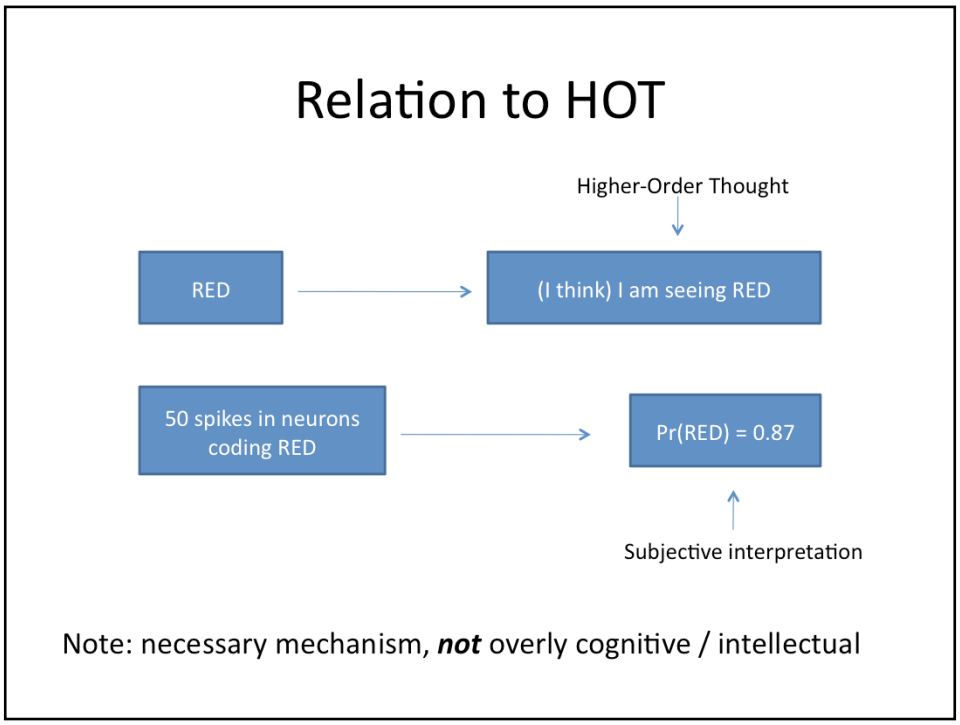


And they dissociate in normal subjects too.

The details are not so important, but I have used metacontrast masking to create conditions where d' (for discriminating between a square and a diamong) is matched. And yet, the frequencies at which people report clearly seeing the target can differ systematically.

So, I can manipulate people's perceptual certainty level while keeping d' constant.

This means the dissociation is not just some weird thing in brain-damaged patients. It happens in all of us.

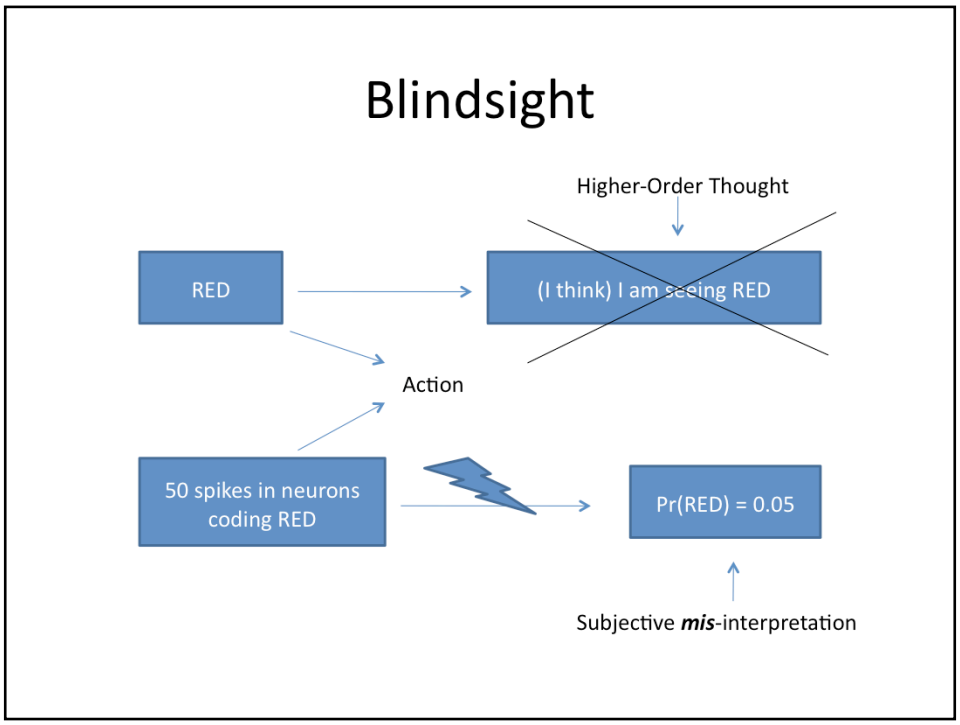


In philosophy, of course the HOT theory may naturally explain such dissociations. In fact, my view is closely related to HOT.

In HOT: e.g. you have a "first-order" representation of RED. And then, you're conscious of seeing red if you have the HO representation that allows you to *think* that you're seeing red.

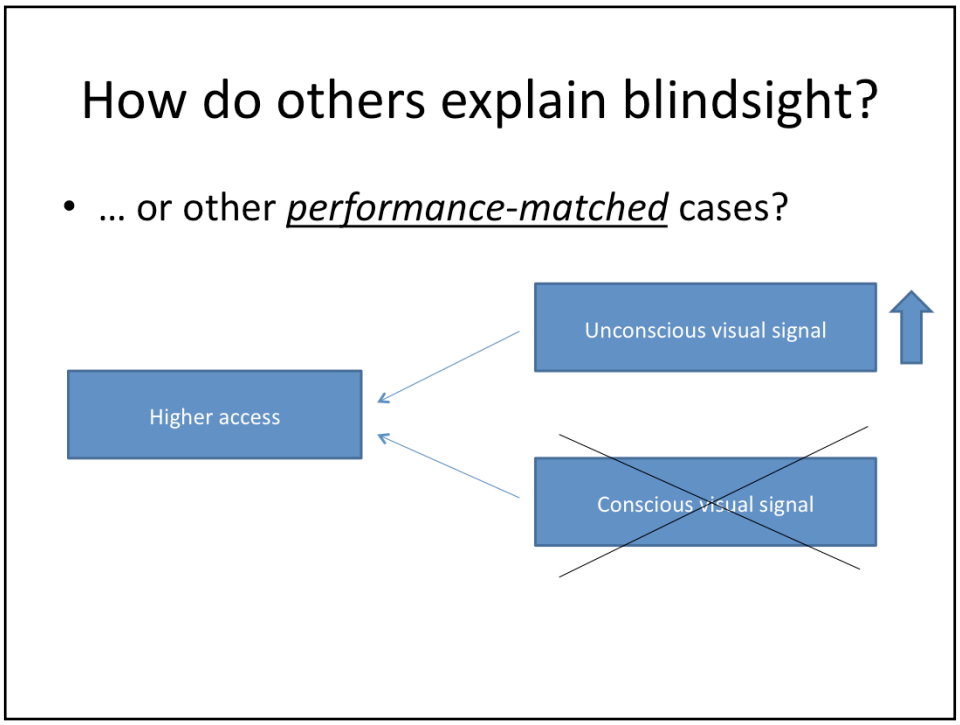
My view is similar, but perhaps more in cognitive neuroscience terms. E.g. You have 50 spikes of firing in total in neurons that code red (first order representation). But you need to have another process to translate or interpret that into probabilistic terms. What does 50 spikes means? It could mean, the probability that the target is red is 0.87. This ***inference*** depends on your baseline firing rate, and other statistical info. This is an active area of research (cf Alex Pouget, probabilistic population code). I argue that to have a phenomenological conscious percept, you need to have this higher-order process to support this probabilistic judgment. When you judge a signal to be sufficiently reliable, say P(something is there)>0.95, you're aware of the relevant stimulus.

One minor note: Sometimes the HOT is criticized on the grounds that the higher-order THOUGHT seems to be rather clumsy / cognitively demanding. Children/animal may not have such thought, but they should be conscious. I think it's not entirely fair because such thoughts needn't be the normal kind of conscious thoughts we usually talk about. But anyway, my higher-order computation is actually rather mechanistic. And you probably need that anyway, to rate your confidence – even for children and animals.



You can see how this kind of higher-order views explain blindsight or similar dissociations of d' and perceptual certainty here.

The first-order representation contributes functionally to action. So if that's intact, d' is the same. But in HOT, if the thought is missing, you don't have the conscious percept. In my view, you can still have the higher-order process. But if the process goes wrong (such as because the cortical connections is broken after the V1 damage), you may mis-interpret the strength of the signal. So you can have a d' of 3, but you make a wrong statistical inference and mis-estimate your own d' to be near zero. That's why you're not certain. And unaware.



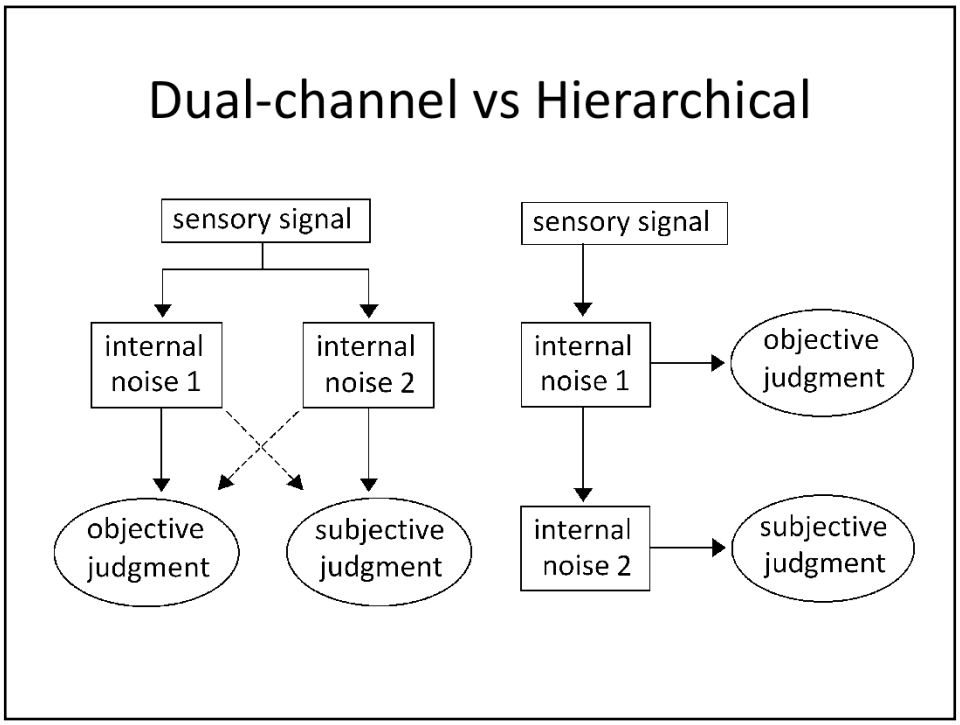
How do other people explain blindsight and other performance-matched cases where awareness differ?

Turns out, both the workspace people and the people like Ned would have to give a similar answer, I think. They can disagree on whether higher-order access is necessary for full-blown consciousness. But for someone to lose the conscious percept, and yet have the same performance capacity, the only obvious way is that: the neural signal supporting a conscious percept is gone, but some other unconscious visual signal compensates for performance capacity.

Ned would like that, I suppose. He would say that the conscious signal is carried by a special biological substrate (something specific to the posterior visual cortex). There may be an unconscious signal to be carried by some other substrate (say in a subcortical pathway).

Let's us call this a dual-channel account of conscious and unconscious perception. In some ways, I think that's what most people assume is the only way to think about the issue. Some areas of the visual brain is conscious. Some is not.

I apologize for failing to discuss Mel Goodale's view here in more details, because of lack of time. But of course, his dorsal-ventral account is one classic example of a dual-channel system. On his view, the ventral system is conscious (with respect to object recognition) and the dorsal isn't.



I want to invite you to think of an alternative view, which I call hierarchical.

(see figure)

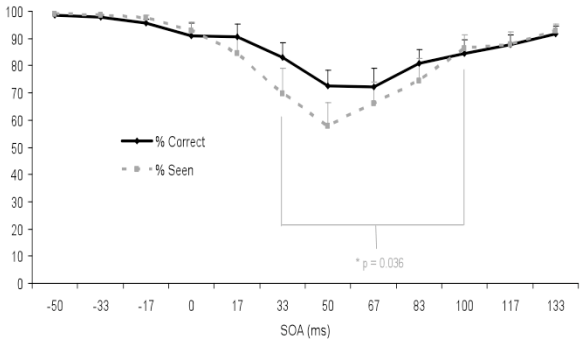
The critical difference is that in the dual channel view there are two largely independent channels, where the noise is largely independent. The hierarchical model has it that the internal noise at an early stage that drives actions and forced-choice discriminations gets inherited to a later stage, where one makes subjective judgments such as awareness and perceptual certainty.

The hierarchical model is supposed to reflect the higher-order views.

Because the noise structure is different. We can distinguish them computationally.

Formal model comparison

- Dual-channel can fit
- BUT hierarchical fits much better



Maniscalco et al, COSYNE 09 abstract & poster

Having spelt out the models in these formal terms, now we can program them into a computer stimulation. Because the noise structure is different. We can distinguish them computationally.

We tried to create these models and see how well they fit with our data (metacontrast masking) where d' and perceptual certainty dissociates.

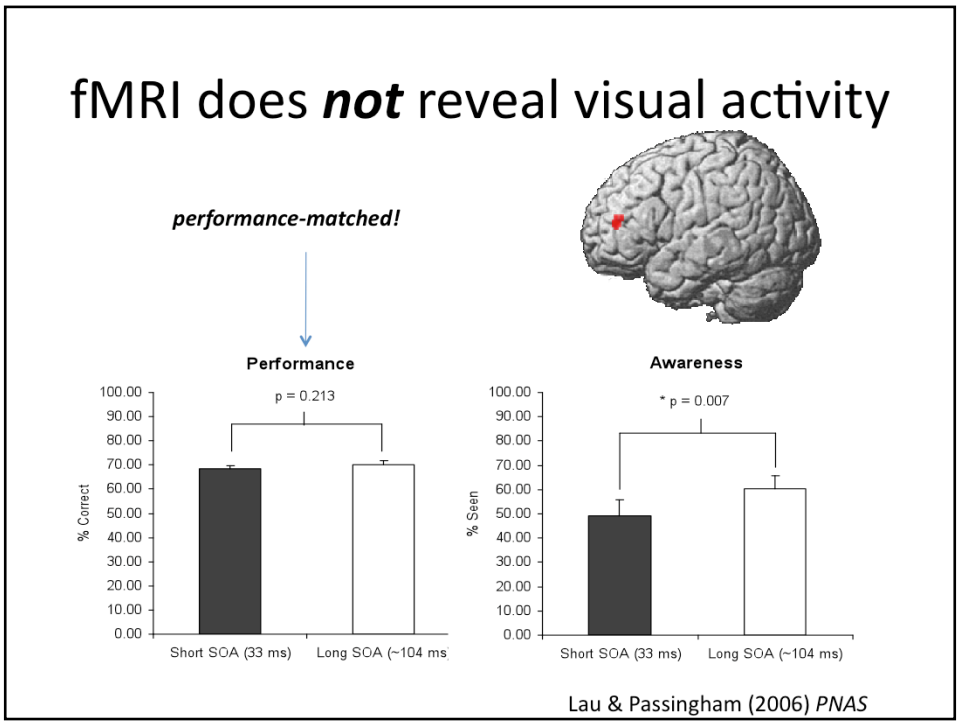
Turns out, the dual-channel can fit the data alright.

This may not be surprising. Stan Dehaene, a proponent of the standard workspace view, recently has discovered that to account for unconscious behavior, he needs to extend his model to include a secondary unconscious channel too. He claims that his dual channel model can fit the data very well, as compared to a single channel model. (appendix of a recent paper in Brain)

What they have not done though, is to compare it against a similarly complex model, such as the hierarchical model. Two channels can almost certainly fit better than one. The more parameters you have in a model, the more you can fit. The gold standard in model comparison is to control for model complexity, using Bayesian model selection techniques. When we do that, we find that the hierarchical model fits MUCH better.

The details are complicated. And I have to acknowledge that this is work in progress. This works for our task, but we haven't tried to see how well it generalize yet. There aren't many tasks where d' and subjective judgment systematically dissociates.

But I can give you more other evidence.



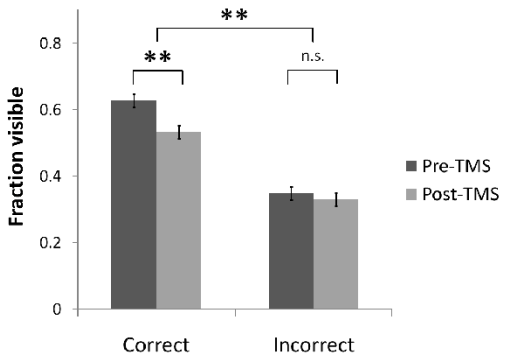
e.g. I have replicated the psychophysical results of metacontrat masking in the fMRI.

On the dual-channel view, I should find differences in 2 visual areas (one conscious and unconscious). It's the change of balance between the conscious and the unconscious visual signal that could account for the data here. This is NOT what I find.

Comparing the two performance matched conditions where reported levels of awareness differ, I find difference in the prefrontal cortex but not the visual cortex.

Assuming the higher-order process for making perceptual certainty inference is in the prefrontal cortex, this fits well.

TMS to prefrontal induces "blindsight"



- Similar recent work by De Cul & Dehaene et al

Similarly, we tried using TMS (theta-burst) to the prefrontal cortex.

One striking thing is that TMS to this area does NOT affect the visual performance. (That's in part why critics have previously argued that prefrontal is not relevant to awareness)

But it changes the perceptual certainty ratings. In a way that is qualitatively similar to blindsight: i.e. TMS lowers the certainty ratings when they are correct.

Recently, De Cul in Dehaene's lab has looked into prefrontal lesioned patients, and found that such lesions affect the certainty ratings more than performance too.

Again, this is hard to be explained by the dual channel model, unless you think the prefrontal is part of the conscious channel.

The prefrontal (DLPFC) is the converging point of the dorsal and ventral stream. If the prefrontal is part of the conscious, then the unconscious must be subcortical…..

Cognitive Unconscious

- **Generation of Spontaneous actions** (Lau et al 2004 Science, 2006 J. Neurosci, 2007 JoCN)
- **Cognitive control** (Lau & Passingham 2007 J.Neurosci)
- **Inhibition** (van Gaal, Lamme & co. 2008 J. Neurosci)
- **Motivation** (collaboration with Pessiglione et al 2007 Science)

- **surprisingly powerful…**

(see review chapter in Gazzaniga's Cognitive Neuroscience IV)

But the unconscious CANNOT be subcortical.

There's a wealth of evidence suggesting that the unconscious is SO surprisingly powerful that it's hard to imagine how the subcortical can do so much

In many cases, one can DEMONSTRATE that unconscious info can activate the cortex, even DLPFC

On HOT, all this is easily explained. The conscious and unconscious share the same channel. All is needed for a percept to change from conscious and unconscious is a change in the HO process. The FO representation can be exactly the same. That may be why you can have unconscious process that is similarly powerful to conscious process.

interim summary

- Hierarchical > dual-channel
  - Computational modeling
  - fMRI
  - TMS
- Consciousness is NOT just strong signal
  - surprising power of Cognitive Unconscious
- but……

Take stock.

So all is good. All is on our side.

Except for one problem….

The Phenomenological Outflow puzzle

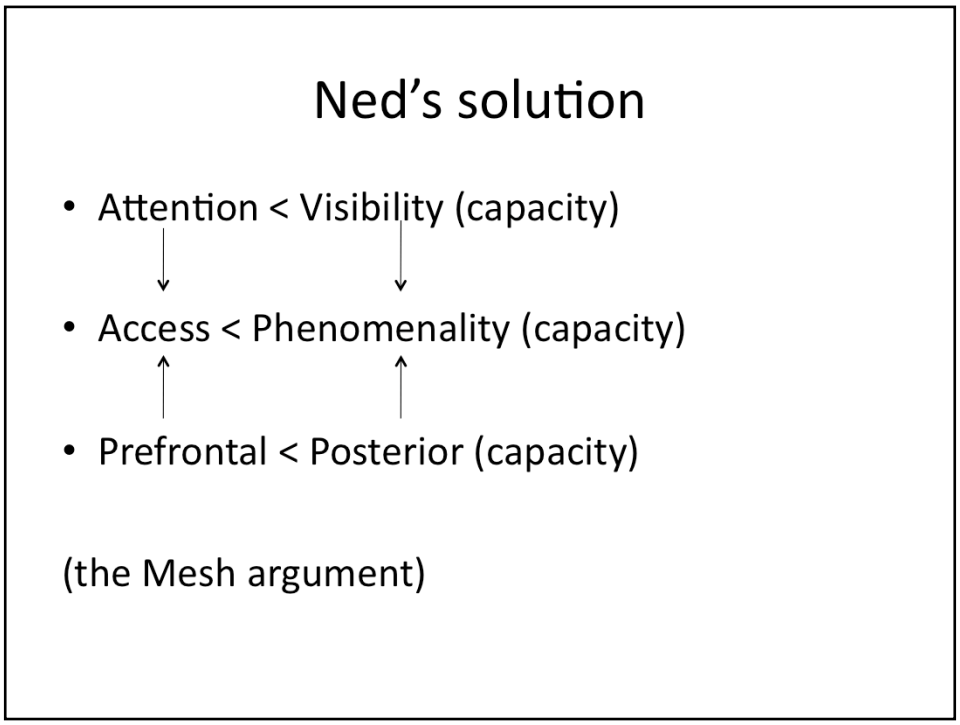


- certainly you think you see more than just the basketball?
- OR - how many notes can you report when you listen to a symphony?

Phenomenology seem to outflow access…

i.e. your phenomenology seems to be more that what you can report/access

i.e. the experiential character of consciousness seems to be what can be measured as the informational capacity



Ned Block's solution is very convincing.

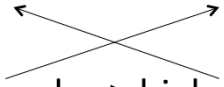
Phenomenality has higher capacity than access. That's because visibility has higher capacity than attention (known in psychophysics). In turn, one would also expect posterior visual system has higher capacity than prefrontal system (more or less expected based on physiology – differences in receptive field, etc).

So, the fact that findings from these different levels "mesh" together seem pretty nice.

Higher-order solution?

- Intuitively….

Access **<** Phenomenality (capacity)



First-order **>** higher-order (capacity)

We need higher-order "fake" phenomenology,
i.e. mis-attribution of perceptual certainty

Developing a similar solution based on a higher-order may run seem difficult….

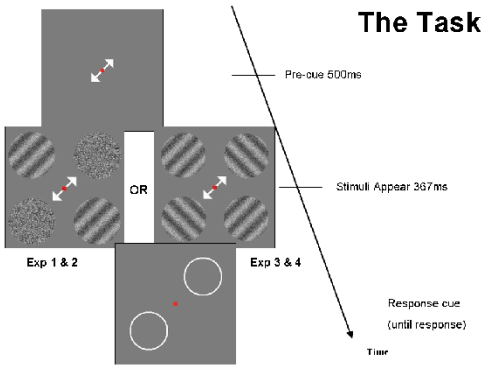
Intuitively, the FO system should have higher capacity than the HO system. So the mapping wouldn't be what we want.

One way that it could work is to not think in terms of capacity, which is an informational term. If it's not information, but ***interpretation***, we could have HO system giving "fake" phenomenology sometimes. I.e. HO process giving perceptual certainty to FO representation that isn't worth the certainty (i.e. it has low d' to begin with).

Inattention increases perceptual certainty

Uncued diagonal
-> liberal detection bias
-> higher certainty rating

(when d' is ***matched***, or
when stimulus is v. weak)



… unattended = more phenomenology?

Rahnev et al (in review)

Turns out this is EMPIRICALLY true. Inattention in general increases perceptual certainty. Yes it sounds crazy. But what it means is: inattention biases the ***interpretation*****.** The brain gives more perceptual certainty than its d' is ***worth,*** under inattention. Of course, attention boosts the signal (d'). So usually, attention still gives more perceptual certainty (because of higher d'). But the rate that it gives perceptual certainty based on d' is more generous under inattention.

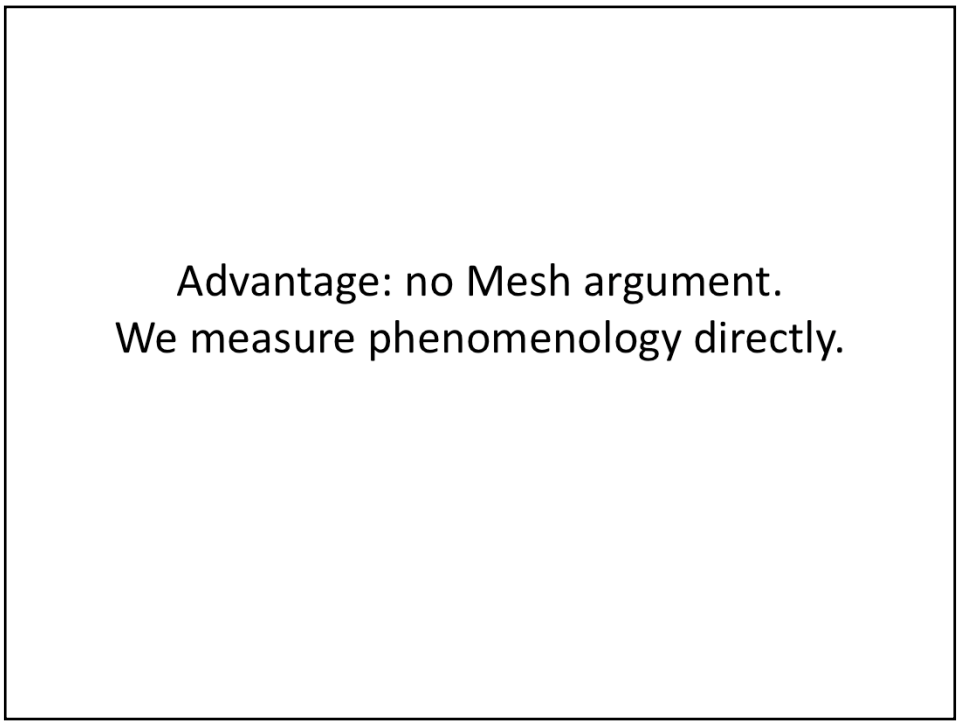
To make it more clear: So we did this task, where subjects were cued to one diagonal in every trial. In some experiments, they DETECT if there is a pattern or not. When we take conditions where d' is matched between the cued and uncued diagonal, the uncued has more liberal detection bias (they say "yes there is a pattern" more frequently). In other experiments, we ask people to DISCRIMINATE between two patterns (left titled or right titled). And then we ask them to rate their perceptual certainty. Again, when we take conditions where d' is matched, people give higher perceptual certainty rating to the uncued. This is quite crazy because they are just as good in both. But they say they are more certain of what they see in the uncued, even though they shouldn't.

This also works when we give them trial-by-trial feedback and training.

This also works when the stimulus is *very* weak, and matched between the cued and uncued. (i.e. no stimulus confound)

So… maybe, in general, inattention gives more phenomenology? We can call it "inattentional filling-in" (you heard it here first)

This fits with the phenomenon of Phenomenological Outflow. Access is limited under inattentional. Phenomenality > acess because the former is inflated by inattention.



So we can have a HO solution to Phenomenological Outflow too. One could even say there's an advantage. Here the solution does not depend on indirectly piecing together differrent evidence. It's based on DIRECT measure. We can do experiments to directly test our predictions – and get them confirmed.

Sometimes people worry that we're measuring only the REPORTS of perceptual certainty, which can be biased, etc. I think, if we're careful, the REPORTS should be taken seriously. And they can make good data. Here it seems to work pretty well.

Standard Workspace Model

- Denies Phenomenological Outflow
- Conflates Harder and Easier problems

d' vs Perceptual Certainty	**Attention/Access vs Visibility**
• phenomenal = subjective (probabilistic interpretation)	• treat SIGNAL as phenomenal
•explains Phenomenological Outflow, blindsight, fMRI, modeling, TMS results, Cognitive Unconscious	•Cognitive Unconscious?
	•committed to ***dual-channel account***, which is… UNSUPPORTED by fMRI, modeling, & TMS

A summary

Major problem with Access Vs Visibility approach: they treat the SIGNAL / performance capacity as phenomenal. That's conceptually bad.

Challenge #1: methodological

- d' and subjective ratings correlate most of the time anyway, so what's the big deal?
- They don't ALWAYS correlate, which means they are not the SAME. Therefore, d' has to be CONTROLLED for. ***Performance-matching!***

Without performance-matching, one can just be looking at performance capacity / the signal, which can be fairly irrelevant to phenomenology per se. Not many experiments involve performance-matching. Hence the bias in discussion of results.

Challenge #2: methodological/conceptual

- Aren't you just studying the ***reports*** of "perceptual certainty"?

- Yes that's how we get at it at the moment. We try to be careful. And there may be other ways to implicitly tap into that…. (new data coming)

Challenge #3: conceptual

- Is subjective probabilistic interpretation ***special*** enough?

- Maybe not. But it may be as good as it gets in cognitive neuroscience. And it's better than everything else we have right now….

Standard Workspace Model

- Denies Phenomenological Outflow
- Conflates Harder and Easier problems

d' vs Perceptual Certainty **Attention/Access vs Visibility**

Bottom line: Of course I'm biased in giving you evidence in supporting my view. But at least, we should realize that the standard workspace doesn't work anymore (for basic conceptual reasons). As a result, a lot of attention has been given to this "new" approach of "attention vs visibility". Maybe we can consider a different approach too, which is at least just as interesting conceptually.

Maybe neither approaches solves the HARD problem. But that's what we have for now. Let's choose between what we have, and we'll see where do we go from here?