Old news and new news about single-photon sensitivity

Image of chick retina by Andy Fischer, Ohio State Wexner Medical Center.

These slides will appear at
www.physics.upenn.edu/~pcn
(or just google me)

Plan

1. Indoctrination
2. Light
3. Vision: psychophysics
4. Vision: single-cell
5. Synthesis

Just four ideas

- Understanding your own body sometimes requires top-drawer physics ideas.

- Sometimes a simple physical measurement can give you insight decades earlier than "ought" to be possible.

- But sometimes that measurement needs to be coupled with some mathematical analysis.

- Once you understand, even partially, how Nature has implemented one of its impressive tricks (e.g. vision), often you also gain practical benefits.

[And anyway, it's beautiful.]

Plan

1. Indoctrination
2. Light
3. Vision: psychophysics
4. Vision: single-cell
5. Synthesis

Eyes

Eyes are an ancient invention: Here is Trilobite, half a billion years old.

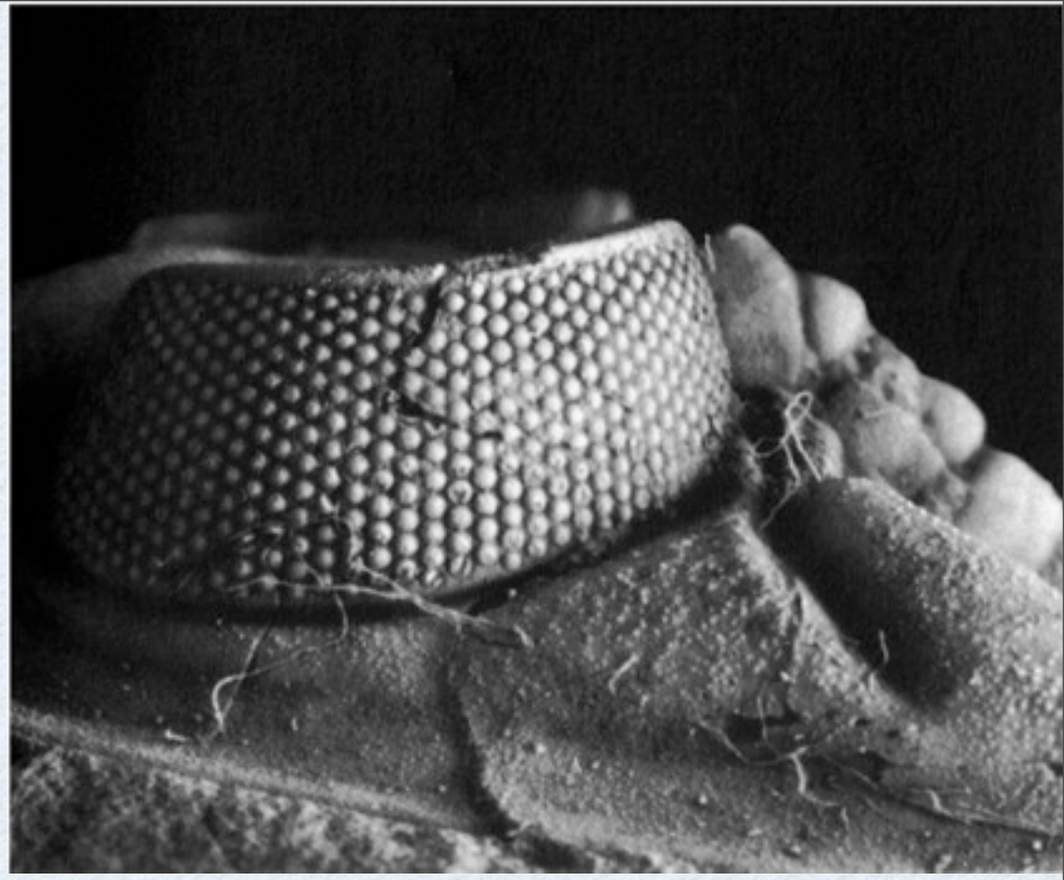

That design was successful: Here is a modern aphid.

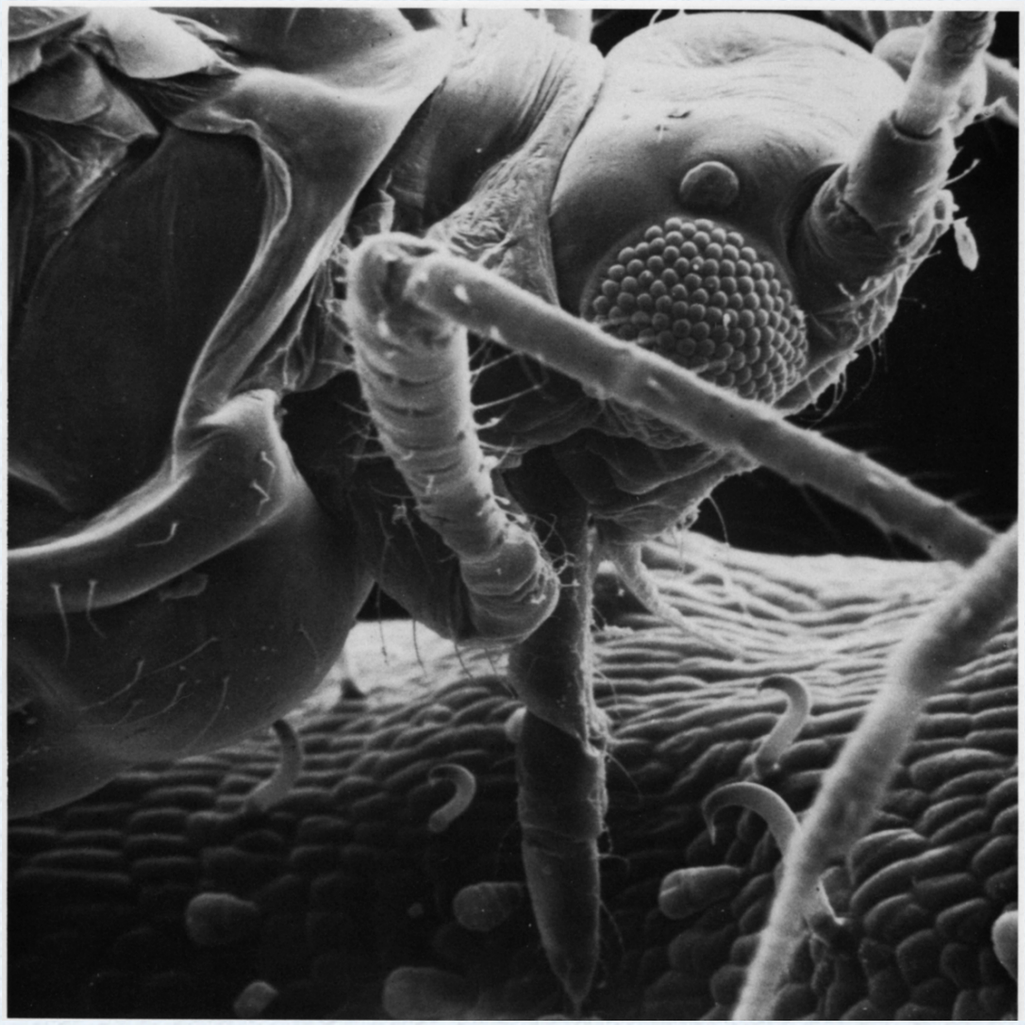

But if you can afford to carry more weight around, here is a better design:

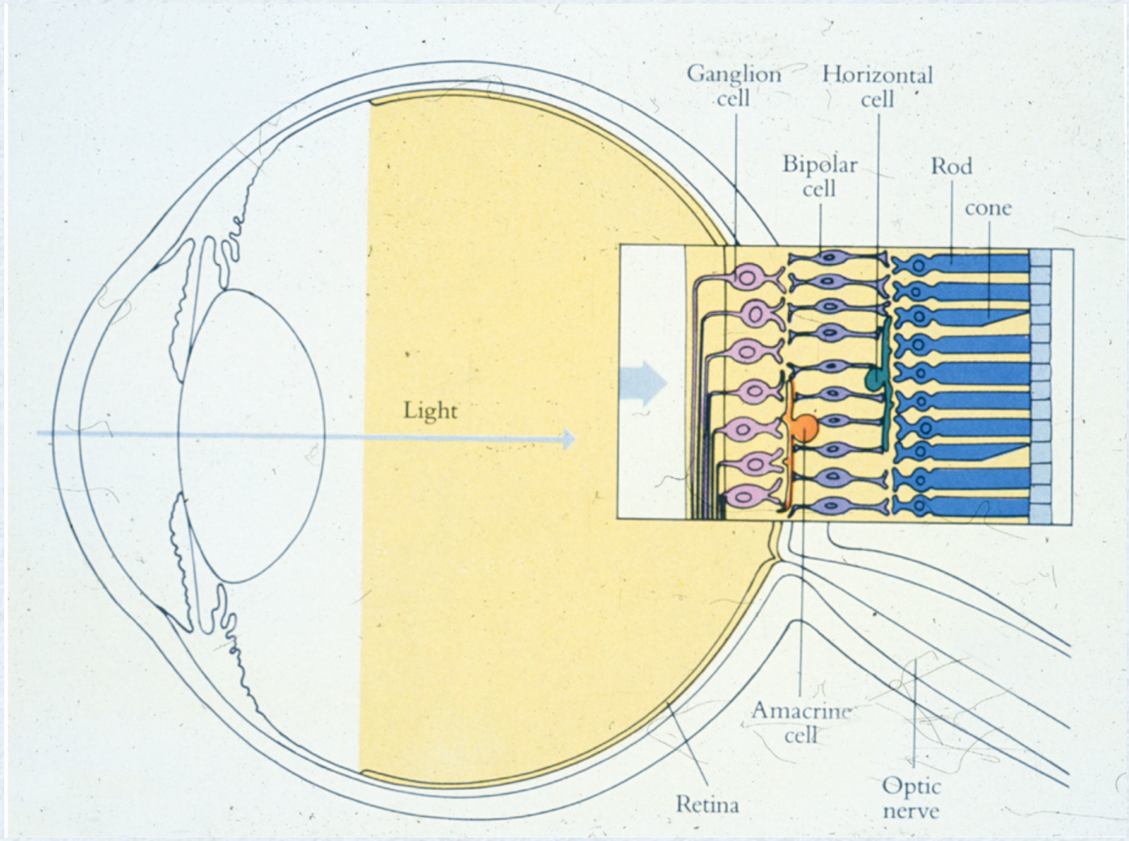

False sense of security

The human eye has a lens-based focusing system, in some ways like a camera. It seems to make sense in terms of Snell's Law, a consequence of the wave theory of light.

The finite size of our pupils limits the resolution we can get at the retina, due to diffraction. There's no point having a pixel size smaller than this resolution limit, and remarkably, *our photoreceptor cells really are about this size*.

Looks like the wave theory of light explains everything.

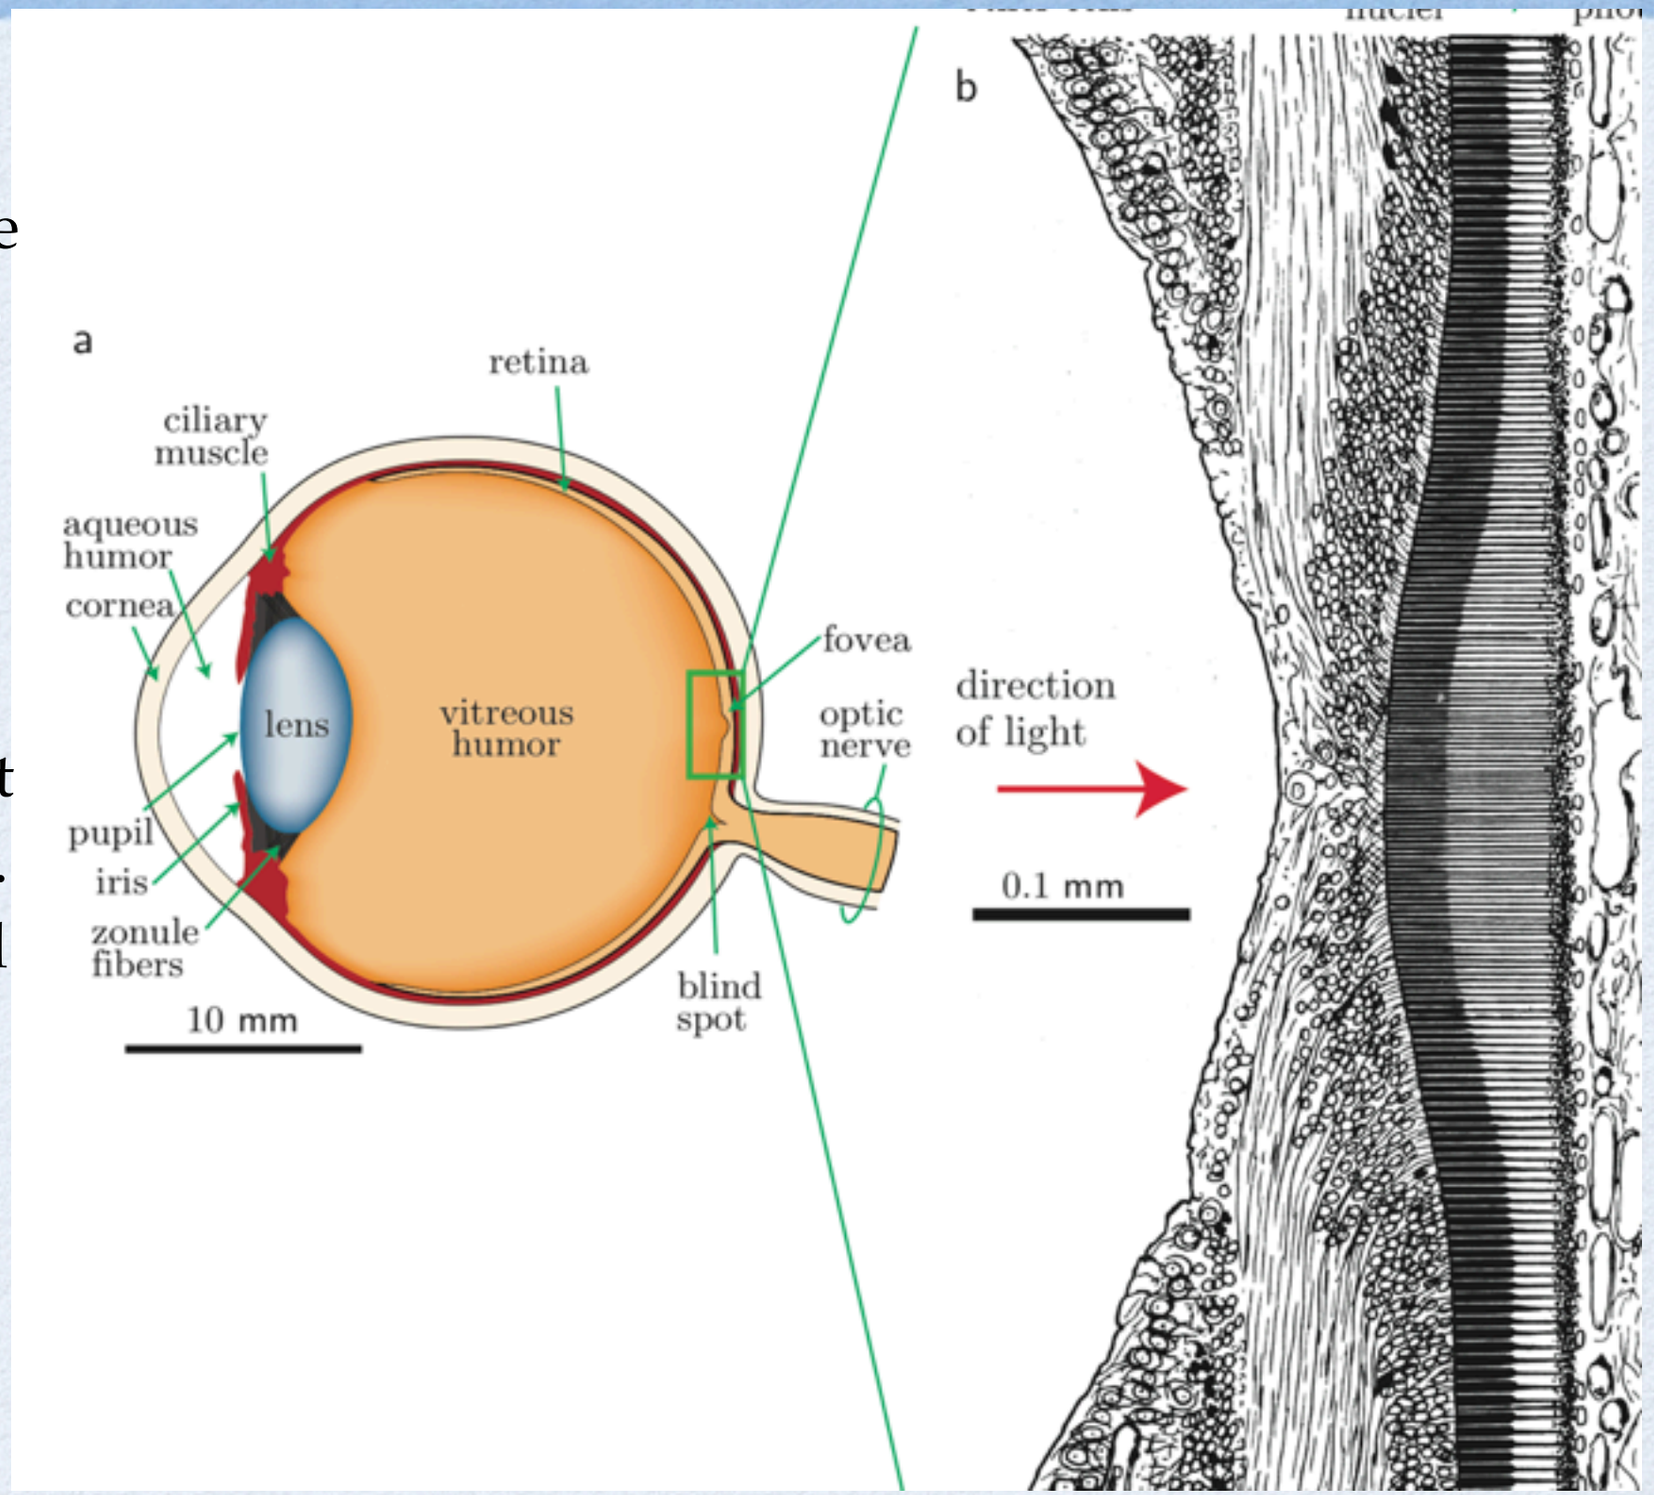

From P. Nelson, ***From Photon to Neuron: Light, Imaging, Vision*** (in preparation).

Uh-oh

But what happens *next?* What happens in those photoreceptor cells that translates light into nerve impulses?

We can detect very dim light with a photomultiplier tube, or a photodiode. Either way, light causes discrete clicks in the detector. *Dimmer light gives equally big clicks, just less frequent.*

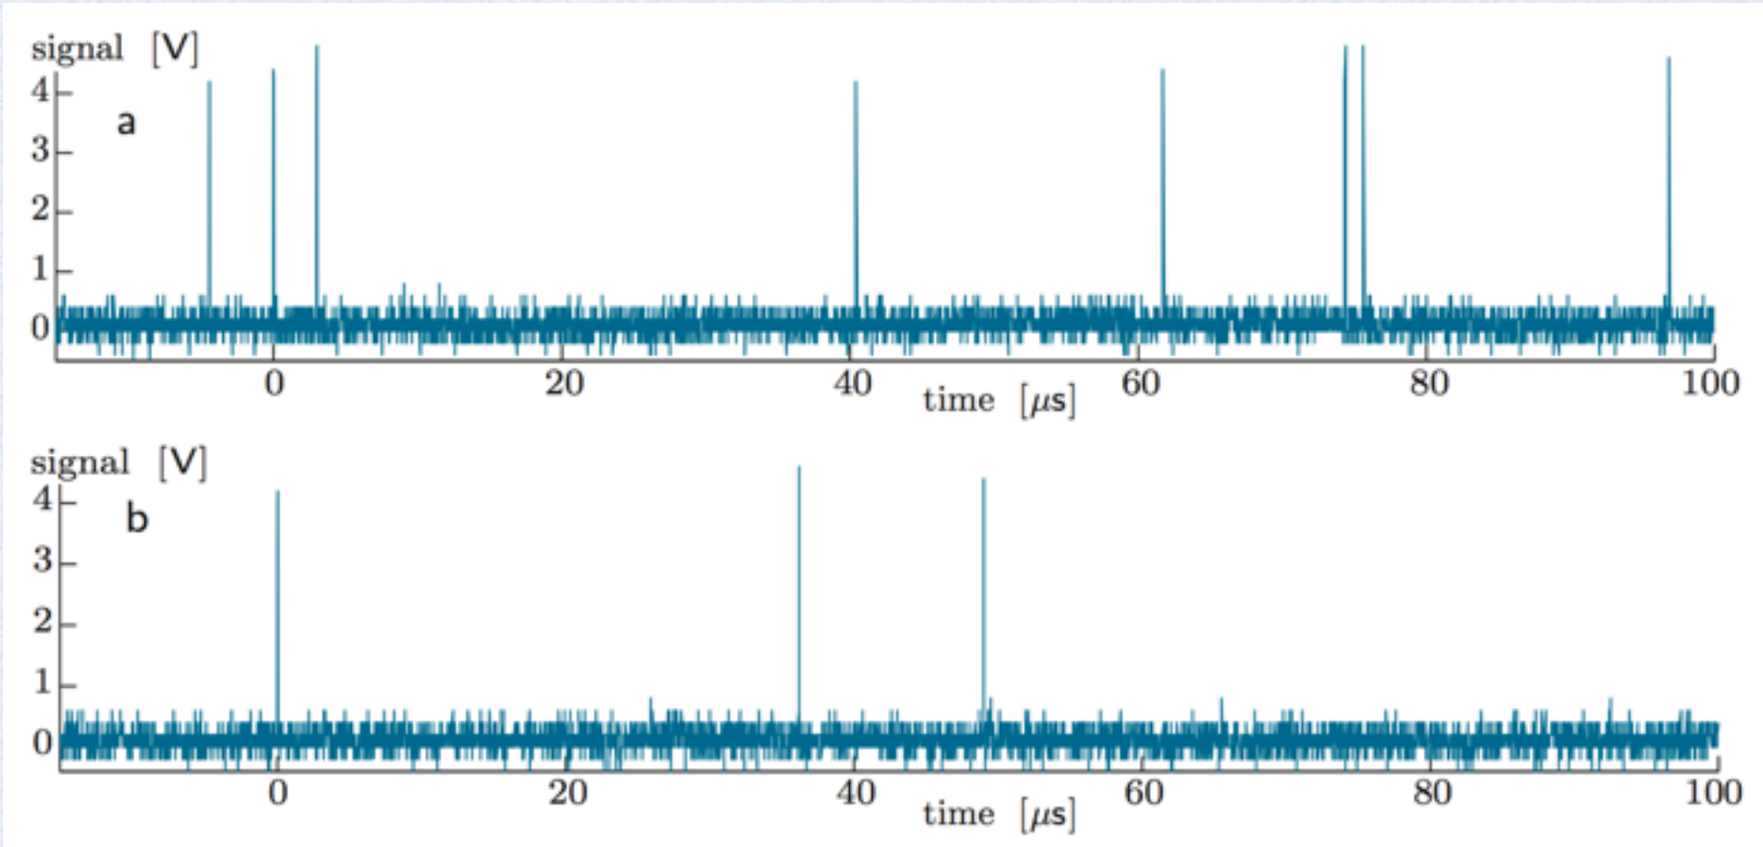

These are *not* uniformly spaced blips. Instead the clicks are *as random as possible* -- they are a "Poisson process." Something about light is intrinsically random.

click for uniform clicks audio *click for shot noise*

Moreover, when we shine dim light on *several* photodetectors, they *never respond in unison*: Each click comes from just *one* detector, even if the beam of light is spread out to cover them all.

From P. Nelson, ***From Photon to Neuron: Light, Imaging, Vision*** (in preparation).

Both digital and film cameras also expose one pixel at a time, at random:

Even classic diffraction effects turned out to be particulate in character.

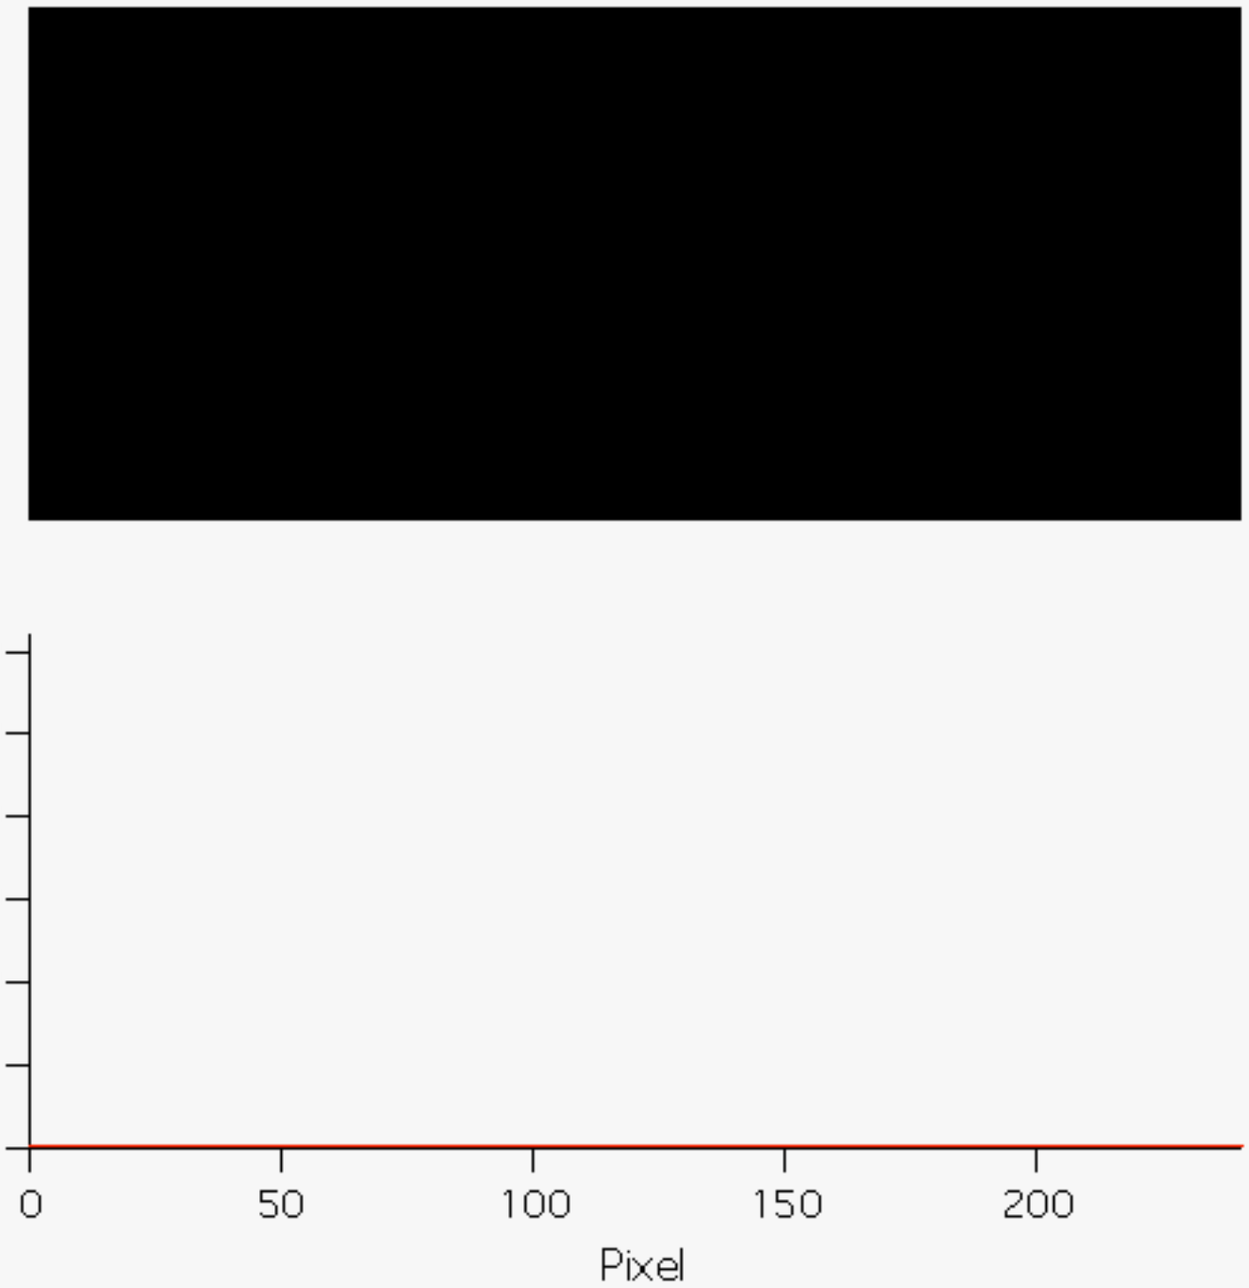

Jean-François Roch
François Treussart
Philippe Grangier
http://www.physique.ens-cachan.fr/old/franges_photon/interference.htm

Detection of light

So the *emission* of light is lumpy and random. Moreover, when light cruises by a single atom or molecule, its *capture* (absorption) is also random – a "Bernoulli trial." And typically the probability of capture is very low.

That's not too surprising, because light cannot be focused down to a spot smaller than a few hundred nanometers wide, whereas individual atoms or molecules are much smaller than that. And how can an incoming photon "know" whether it "should" be absorbed by a particular molecule? It just makes a random choice.

A high-efficiency detector may capture and signal nearly 100% of the photons that land on it, but only because it offers a very large number of molecules to the incoming light and – somehow – is able to signal whenever *any one* of the catches a photon.

So if we find that our own eyes are highly efficient, we'll need to understand how they manage that trick.

Plan

1. Indoctrination
2. Light
3. Vision: psychophysics
4. Vision: single-cell
5. Synthesis

Vision in dim light

Image Edith Widder

"What's all that theoretical stuff got to do with *vision,* a real biological process? Surely the inconceivably tiny energy in a single photon is irrelevant to a macroscopic organism like me?"

But on the contrary: If your competitor / threatening predator / prey has better night vision than you, that could be a problem. Survival is the mother of invention, big time. If there's a fundamental limit to vision, you need to be there if at all possible. Big eyeballs can help, but we can't all carry around basketball-sized eyes like *Architeuthis*. We need to detect *every photon,* or come as close to that ideal as we can. Lorentz, 1920: Are we anywhere near to that ideal?

Photons arrive randomly in time.

As mentioned earlier, single photon *absorptions* are also unavoidably random in character.

And any one molecule, even if engineered to be good at catching visible photons, will only catch a small fraction of those that arrive.

For all these reasons, the results from a series of identical trials will be different every time. How can we do science under such intolerable conditions, with such intrinsic irreproducibility?

Flash of light

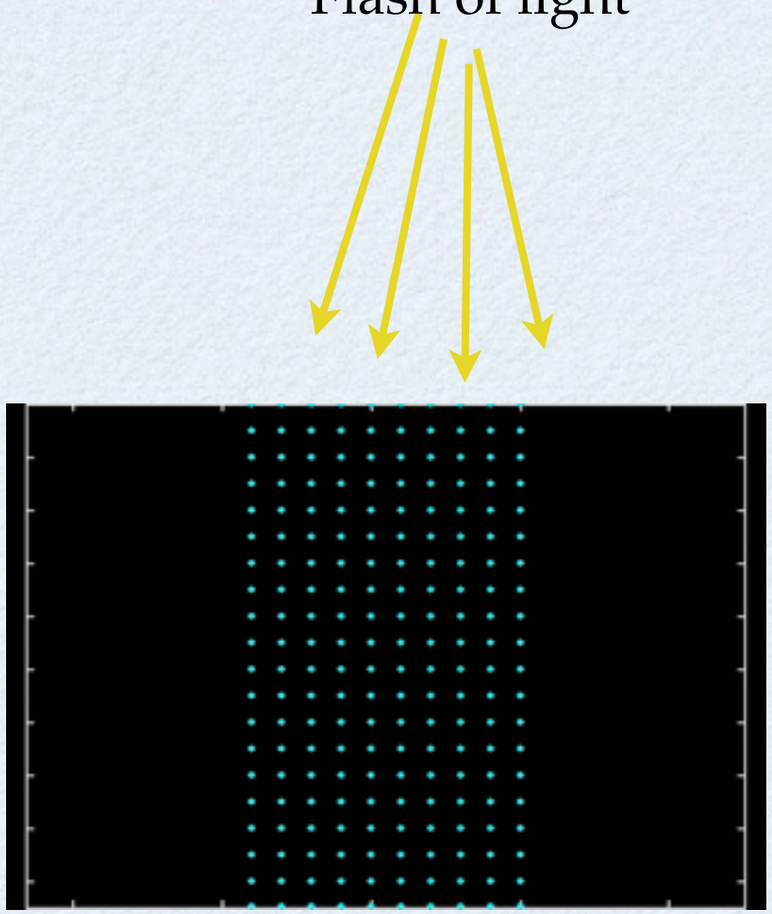

Photoreceptor cell packed with sensitive molecules

Identically prepared flash of light

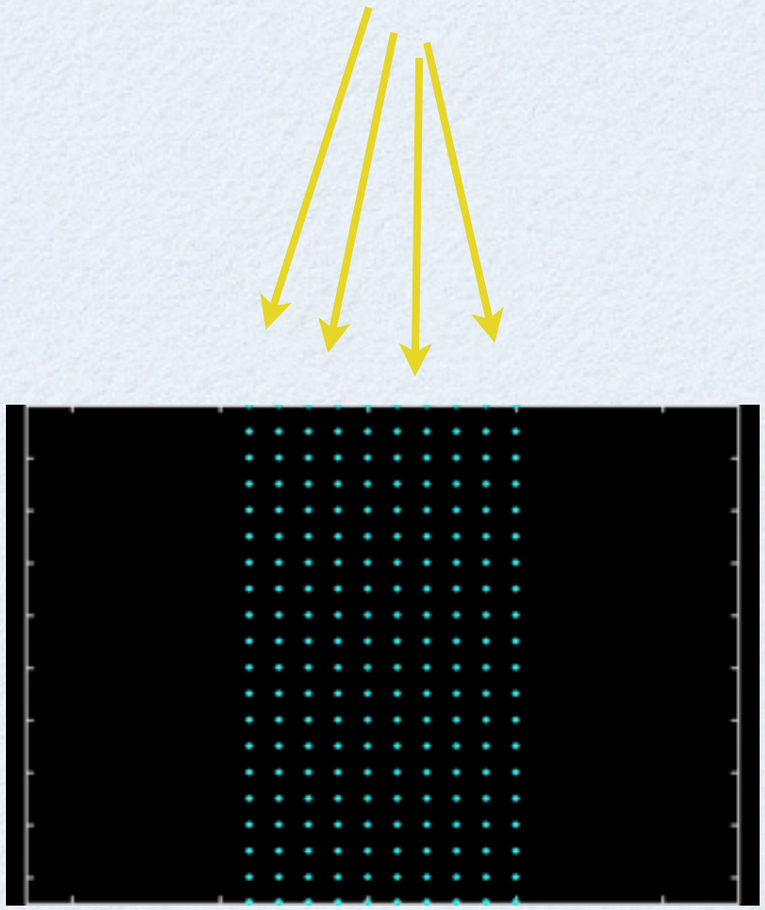

Identical photoreceptor cell

And again

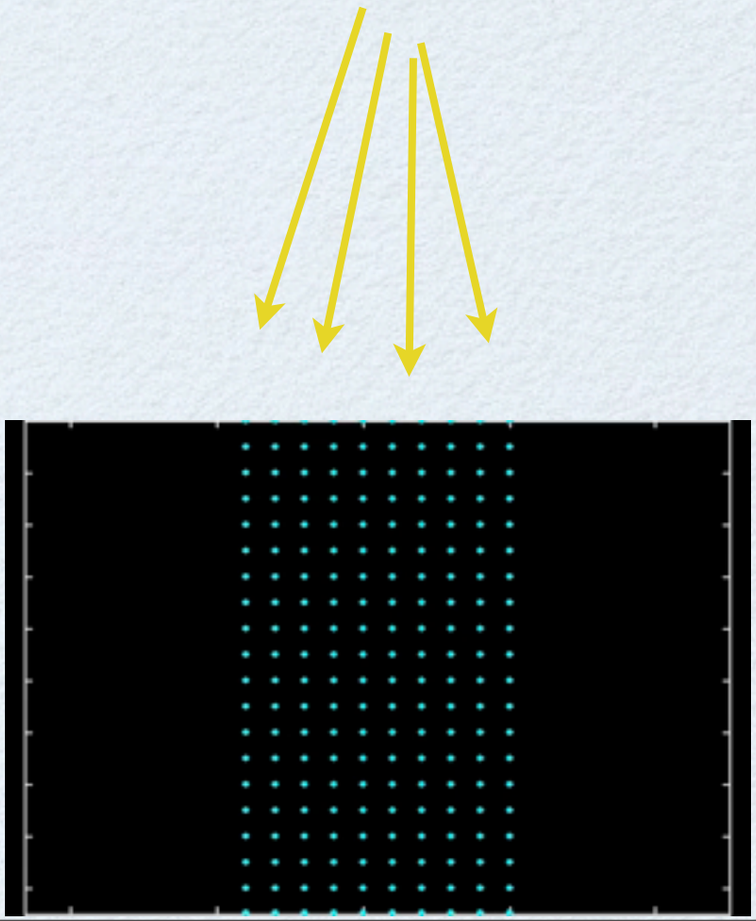

[Aside]

We need an experiment. Actually we need several.

An old experiment

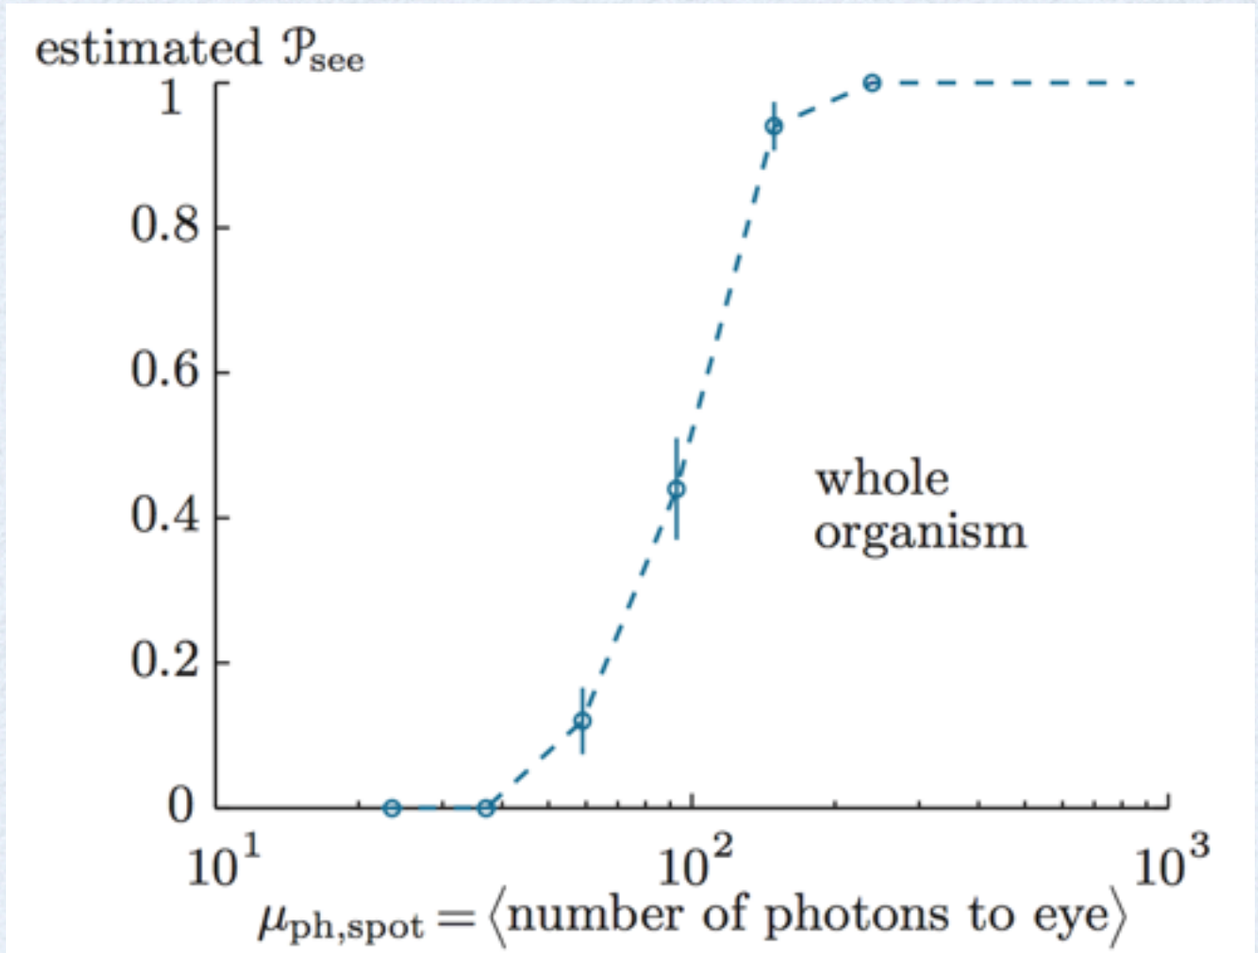

Hecht and coauthors sidestepped the "unrepeatability" problem by measuring the *probability* for a human subject to see a flash, vs intensity. They found that, at the lowest detectable flash intensities, seeing follows a sigmoidal curve. What can we conclude?

From P. Nelson, ***From Photon to Neuron: Light, Imaging, Vision*** (in preparation).

Hecht's amazing deduction

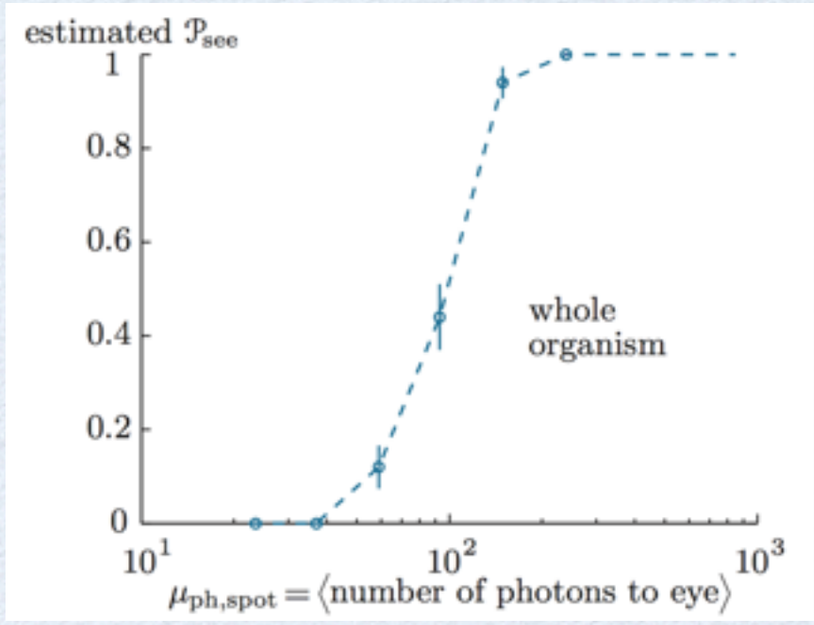

Hecht et al said, "If 150 photons to the eye is enough to see a flash fairly reliably...
and half of those are lost before they get to the retina (as is measured)...
and 4/5 of the reminder are absorbed by something other than rhodopsin...
then at around 15 absorbed photons is enough to get a reliable conscious impression.

But those 15 are spread over a spot served by thousands of rod photoreceptor cells...
so it's very improbable that any one cell caught more than one photon...
and yet this is a stimulus that reliably evokes a response...
so it must be the case that
an individual rod cell can generate a signal upon absorbing a single photon.

That is, ***we are all single-molecule biophysicists.*** right down to the lowly dung beetle.

Hecht et al. continued: In response to a faint flash, each individual rod photoreceptor signals or does not, independently with a small probability proportional to the flash strength. To avoid distraction, some neural system alerts the conscious brain whenever the number of such signals exceeds some "quorum requirement," *t*. That is, ***the probability of seeing is a cumulative Poisson distribution.***

That hypothesis has the advantage of having only *two fit parameters:* The probability per incident photon, *Q*, that a photoreceptor will signal ("quantum catch"), and the quorum requirement *t*.

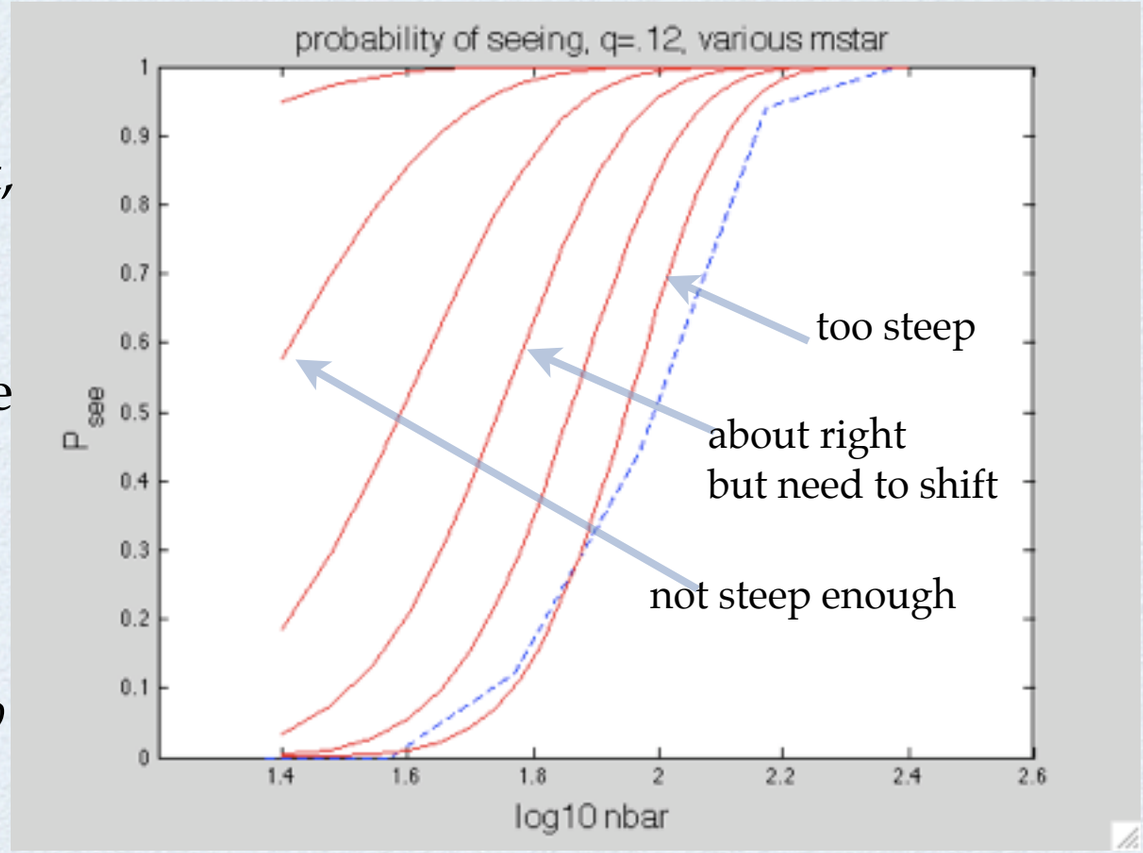

Above, *t*=7 has about the right steepness. Below, then we shift the graph by adjusting *q* to get it right.

```
% data from hecht et al page 835 table V 4th column
p=[0,0,0.12,0.44,0.94,1]; nphoton=[23.5,37,59,93,149,239];lnphoton=log10(nphoton);
plot(lnphoton,p,'o');
xlabel('log10 nbar','FontSize',16);ylabel('P_{see}','FontSize',16);title('probability of seeing','FontSize',16)
figure
plot(lnphoton,p,'--');hold on
for mstar=1:2:12,
   q=0.12;
for j=1:46,
   photons=20+5*j;
   mbar=q*photons;
   total=0;
   for i=mstar:50
      total=total+exp(-mbar)*(mbar^i)/factorial(i);end
   ptheory(j)=total;lphotontheory(j)=log10(photons);end
plot(lphotontheory,ptheory,'r');
end
```

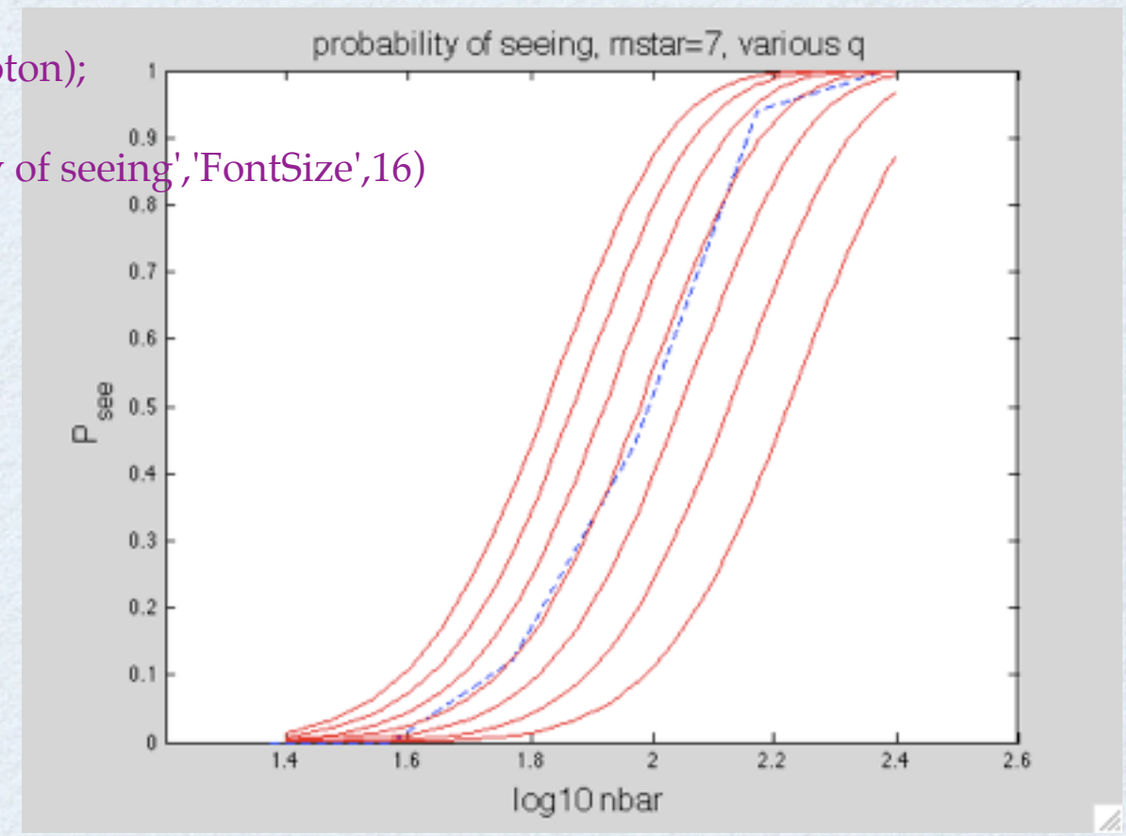

Sometimes simple is too simple

One often hears the takehome from the Hecht experiment expressed as "Our remarkable eyes can detect a single photon." Unfortunately that won't do:

- Quantum theory tells us that no laser, no filament, not shutter placed in front of the Sun, can deliver exactly 150, or any other exact number of photons–you always get a Poisson-distributed random variable. OK, the preceding analysis took that into account.
- But the Hecht data on previous slide shows that, for reliable seeing, we need to present about ***150*** photons to the eye – not ***1***. Why that number? We partially took that into account when we included losses from absorption...
- Still, we were led by the model to assume a quorum requirement of 7 rod-cell signals, not ***1***, before anybody upstairs in consciousness gets alerted. Why not 1?

So yes, our eyes are remarkable, but what's the precise statement of *how* remarkable? Only when we've got that can we start to discard some hypotheses about what's going on and retain others.

A better, though still old, experiment

Also there's a methodological problem in the Hecht experiment: What exactly does it mean to instruct subjects, "Don't make false positive reports"? Van der Velden did a similar experiment independently, and got different results.

Barlow** realized that the discrepancy could be due to the different level of false-positive reports tolerated in each experiment. Barbara Sakitt then performed a better psychophysics experiment than Hecht's: She measured a function of ***two*** variables, the probability of a subject assigning each subjective rating to the strength of a flash as a function of that flash's nominal strength. That gives a richer dataset:

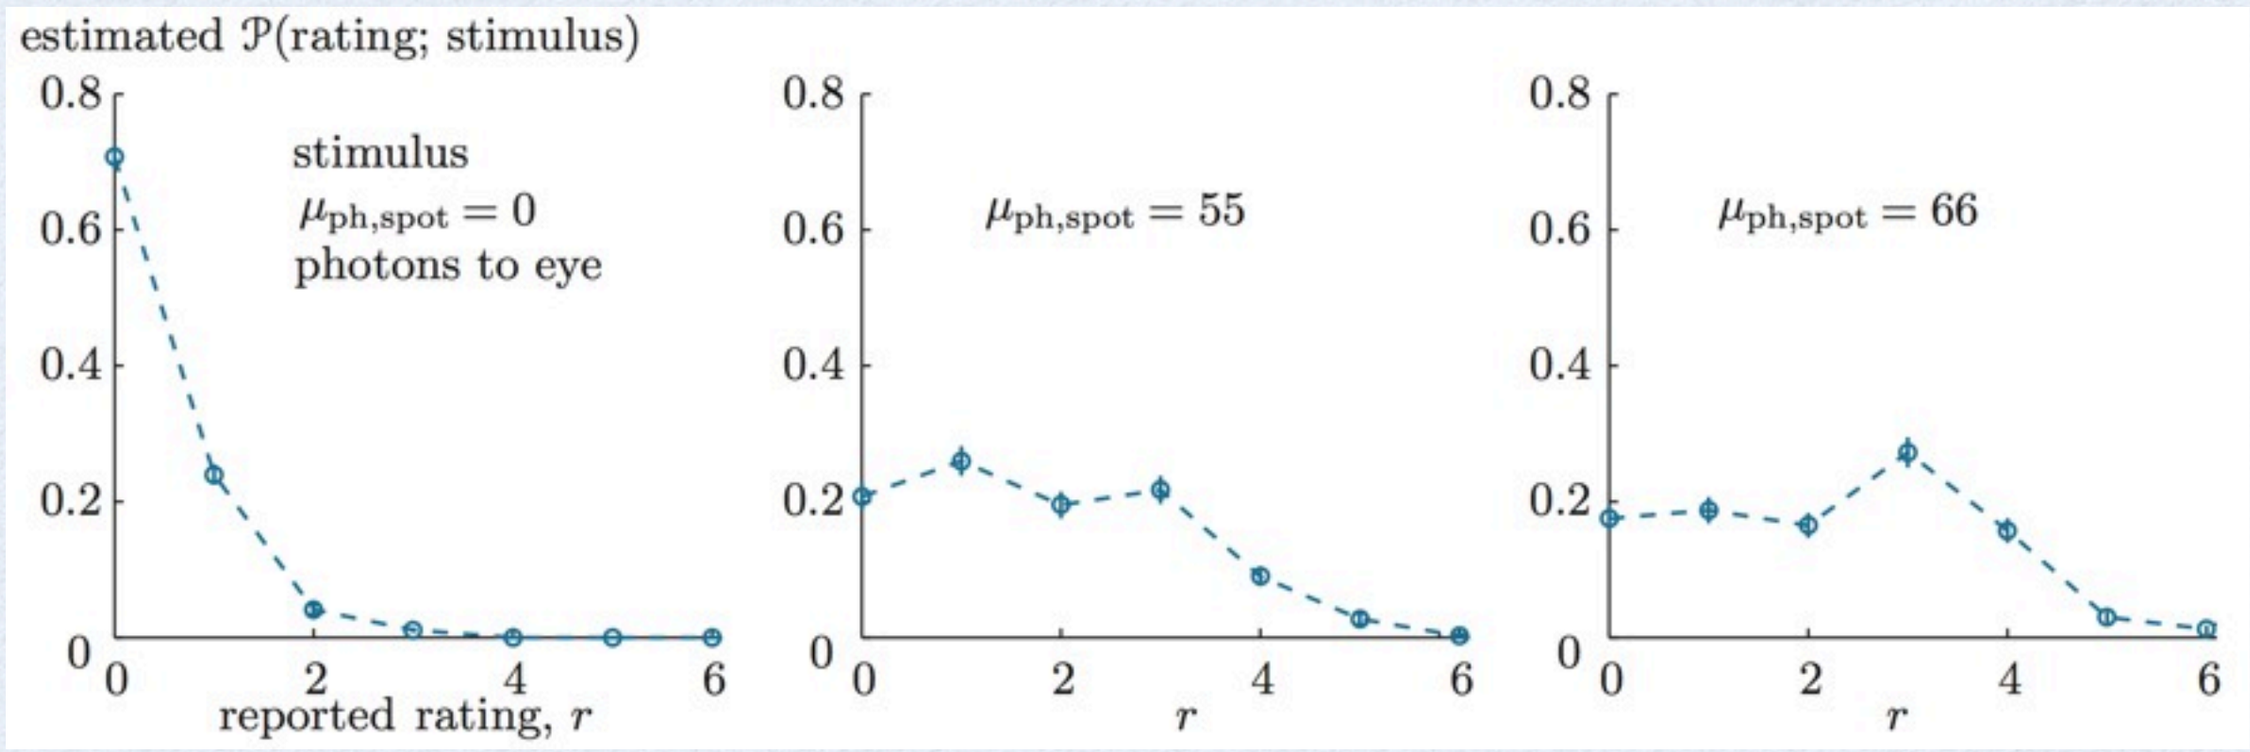

From P. Nelson, ***From Photon to Neuron: Light, Imaging, Vision*** (in preparation). Data Barbara Sakitt (1972).

Sometimes simple is too simple, part 2

Next turn to the puzzles in the model, such as the high quorum requirement. Won't I get better vision–and its fitness payoff–if I am aware of *every* photon?

Barlow pointed out that Hecht's model omitted something potentially huge–the role of spontaneous isomerizations. Barlow proposed that a trickle of fake photon signals, indistinguishable from the real thing, are folded into each rod cell's output at some fixed rate.

- The combined signal from both kinds of isomerization is again a Poisson process, but with mean rate given by a sum:
  - ✳One contribution is proportional to flash strength as before;
  - ✳The other contribution is independent of flash strength.
- As previously, some neural process downstream of the photoreceptors applies a quorum requirement to the total number of rod cell signals evoked by a flash. (If there are several allowed responses, as in Sakitt's experiment, then each one has its own quorum requirement.)
- The low-level hardware is characterized by *fixed* parameters (summarized by the quantum catch). But the high-level firmware is *flexible;* it can set quorum requirement(s) suitable for the assigned task.

Sounds great–we just have *more fit parameters* than before, no problem for our fitting software. I got a nice fit to all of Sakitt's data with this model. But...

So, do we pop the champagne now? No, we say the model is *promising*. Unfortunately, it's not very falsifiable. "Just one more fit parameter" is suddenly a lot: In fact, many different fits to the data are possible, with widely varying values of the parameters (Fred Rieke). So there's more to be done:

Is the fit value of quantum catch *reasonable*?
* E.g., how many chromophores would a single rod need in order to achieve that quantum catch? How many does a rod actually have, anyway? What mechanism in the rod cell could monitor all those chromophores and initiate signaling?
* Is there some independent way to obtain the value of the quantum catch, without fitting to psychophysical data? That would certainly make it harder for our model to give good-seeming results while being wrong. That is, it is good science to *overdetermine* the model as much as possible.

Is the fit value of spontaneous false-positives *reasonable*?
* E.g. how stable would each chromophore need to be in order to achieve it, and yet still be able to photoisomerize upon receiving a visible photon? Are chromophores actually that stable?
* Is there some independent way to obtain the value of the spontaneous false-positive rate?
* Are there other noise sources besides spontaneous isomerization that could contribute?

Are the fit values of the quorum requirement *reasonable?*
* E.g. they have got to be integers, and they will turn out to be *small* integers – not much freedom – quite possibly *no* reasonable values could be found that give a fit after we constrain the other parameters.

Summary: psychophysics

- Hecht et al.'s historic experiment did imply that, remarkably, a single productive photon absorption can elicit a response from a rod cell, without fancy electronics, decades earlier than they had any "right" to get this result.
- But experiments of this type by themselves admit a range of different explanations. Other kinds of experiments were needed to break this degeneracy.

Plan

1. Indoctrination
2. Light
3. Vision: psychophysics
4. Vision: single-cell
5. Synthesis

Onward to single cells

An individual rod or cone cell's response can be measured by gently aspirating its outer segment into a pipette electrode and stimulating it with 500 nm light (*green streak*).

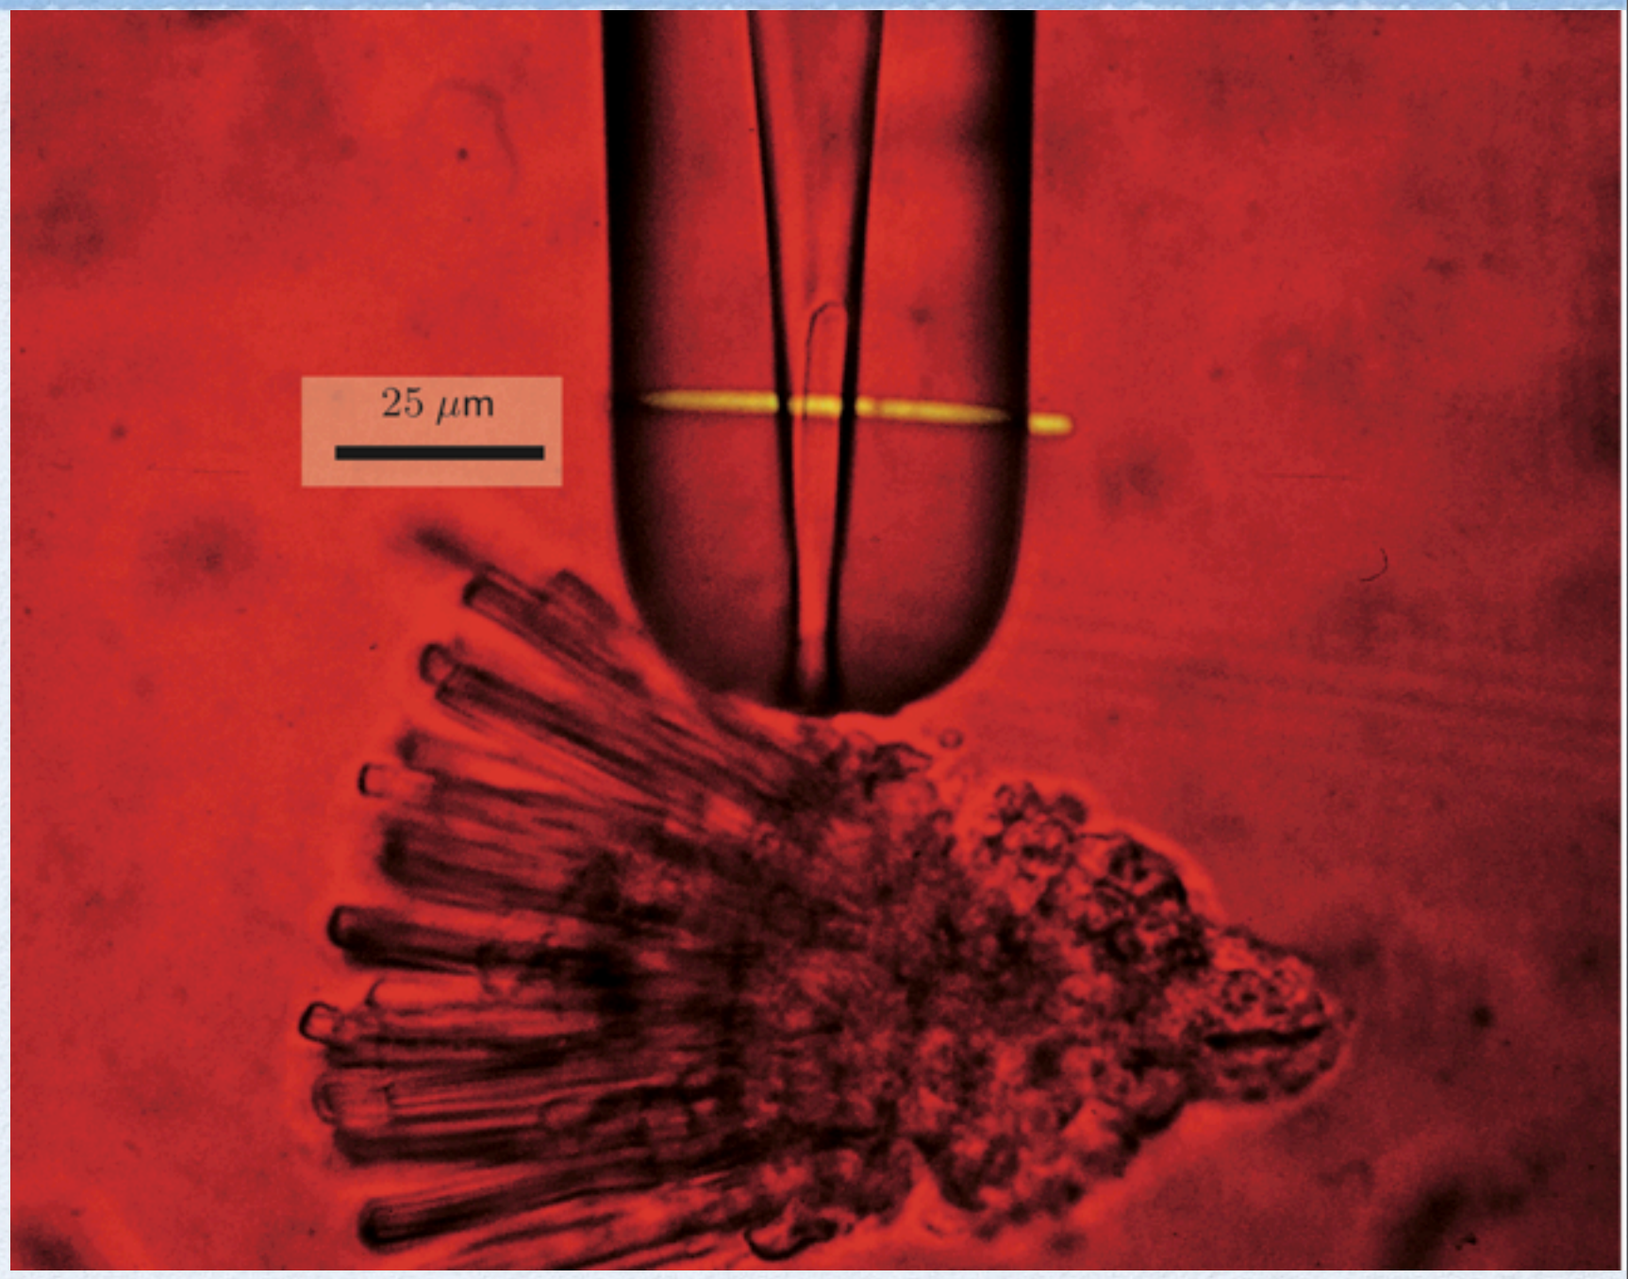

Image courtesy K-W Yau (1979).

Single-cell data

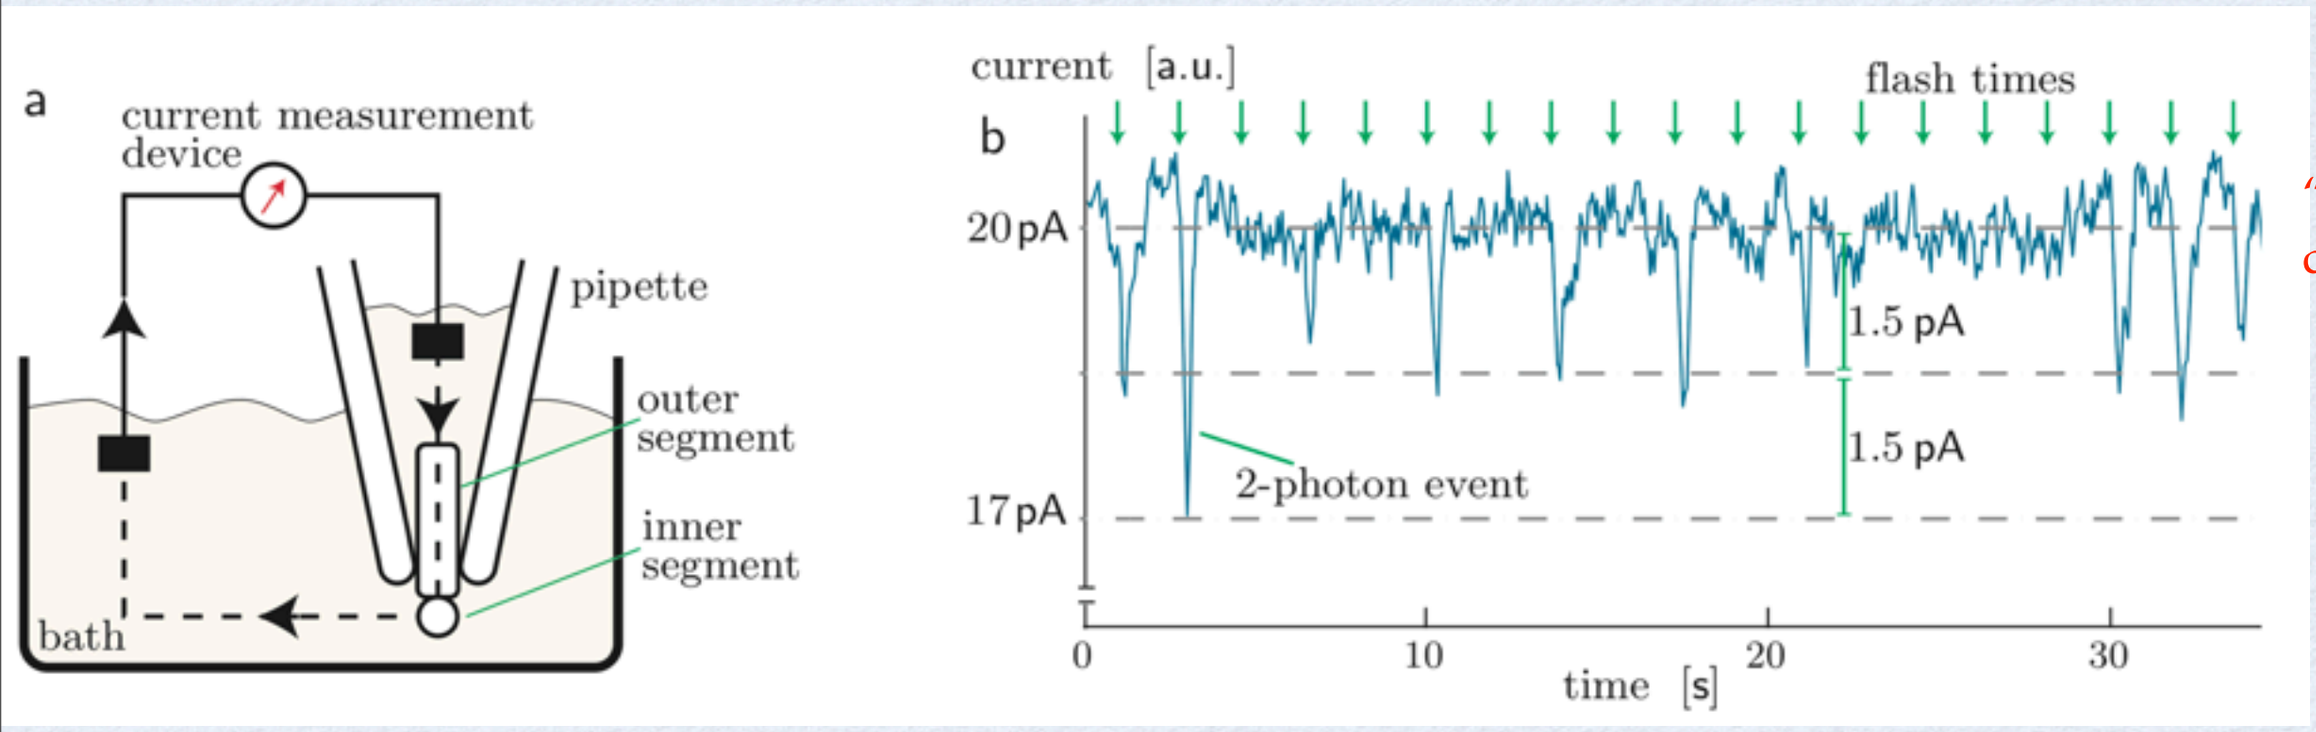

Baylor's experiments showed that although some flashes were missed altogether, others gave stereotyped responses.
The experiments also *two* kinds of noise: signals in the dark that look exactly like real single-photon responses ("spontaneous isomerizations") and a lower-amplitude rumble called "continuous dark noise."

From P. Nelson, ***From Photon to Neuron: Light, Imaging, Vision*** (in preparation). Data courtesy Greg Field (2008).

Another probabilistic argument

Baylor et al. found the probability of "seeing" curve for single rods. They knew the density of rhodopsin, its absorption cross-section, the rod cell dimensions. So from this data they obtained a meaningful (if partial) characterization of how good our visual apparatus is:

- *Most photon absorptions by rhodopsin are productive.*
- *A single productive absorption suffices to trigger a discrete signal.*
- *The false positive rate, though not zero, is small.*

It's a great example of indirect reasoning:
There's no way to fit these data to a model in which individual rods have a threshold of >1 photon for response, *even though there was no way to say how many photons were in any given flash.*

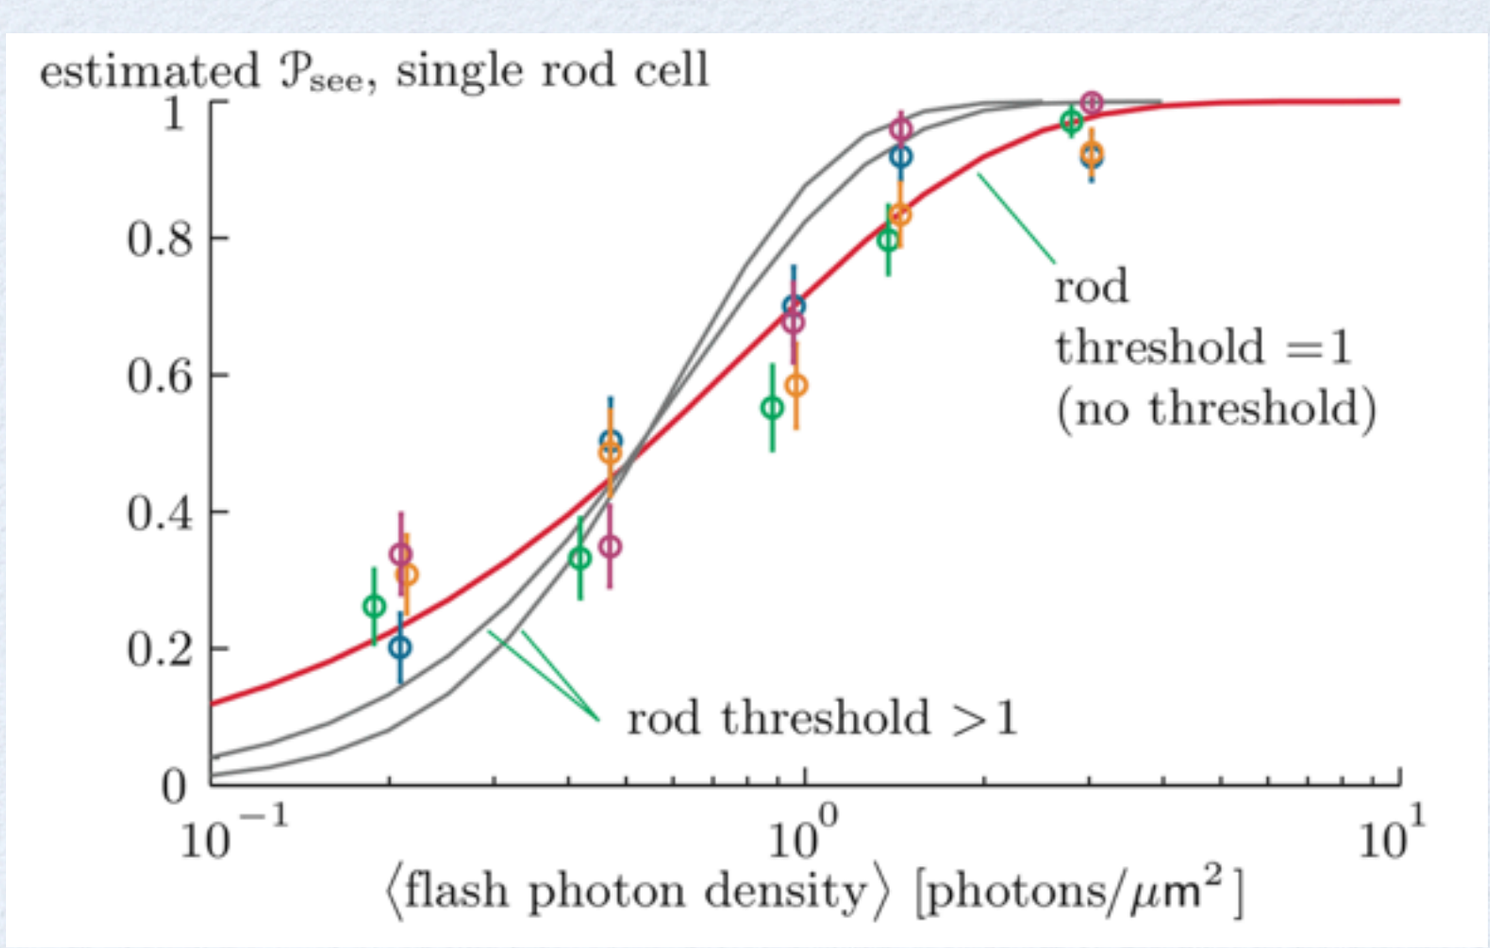

From P. Nelson, ***From Photon to Neuron: Light, Imaging, Vision*** (in preparation). Data from Baylor, Lamb, Yau (1979).

A more detailed prediction

The previous slide showed the probability of *any* rod response as a function of the flash strength. But the slide before that showed that rod responses can be broken down into a series of discrete strengths, strongly suggesting that each bin reflects 0, 1, 2, ... productively absorbed photons in the wake of a flash.

That is, the *full* Poisson distribution of counts, not just the cumulative distribution above 0, is measurable and should be compared to experiment.

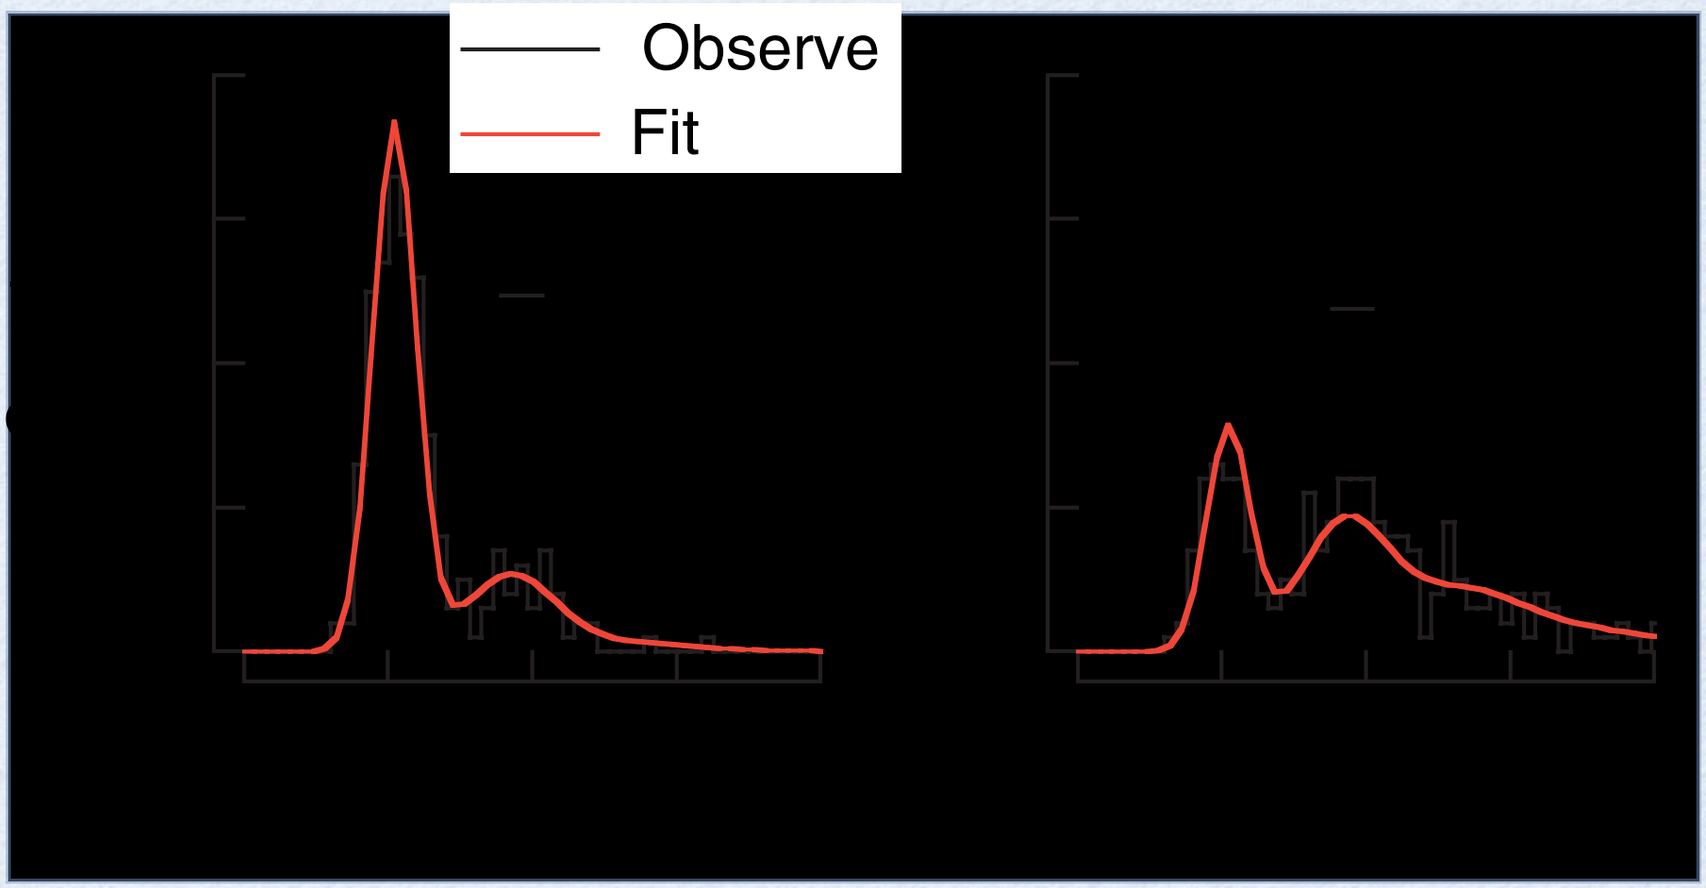

From F. Rieke, in *The Senses: A Comprehensive Reference*; experimental data from G. Fiel

Fast-forward to 21st century

Very recently, high-tech experiments confirmed directly that the rod cell threshold :

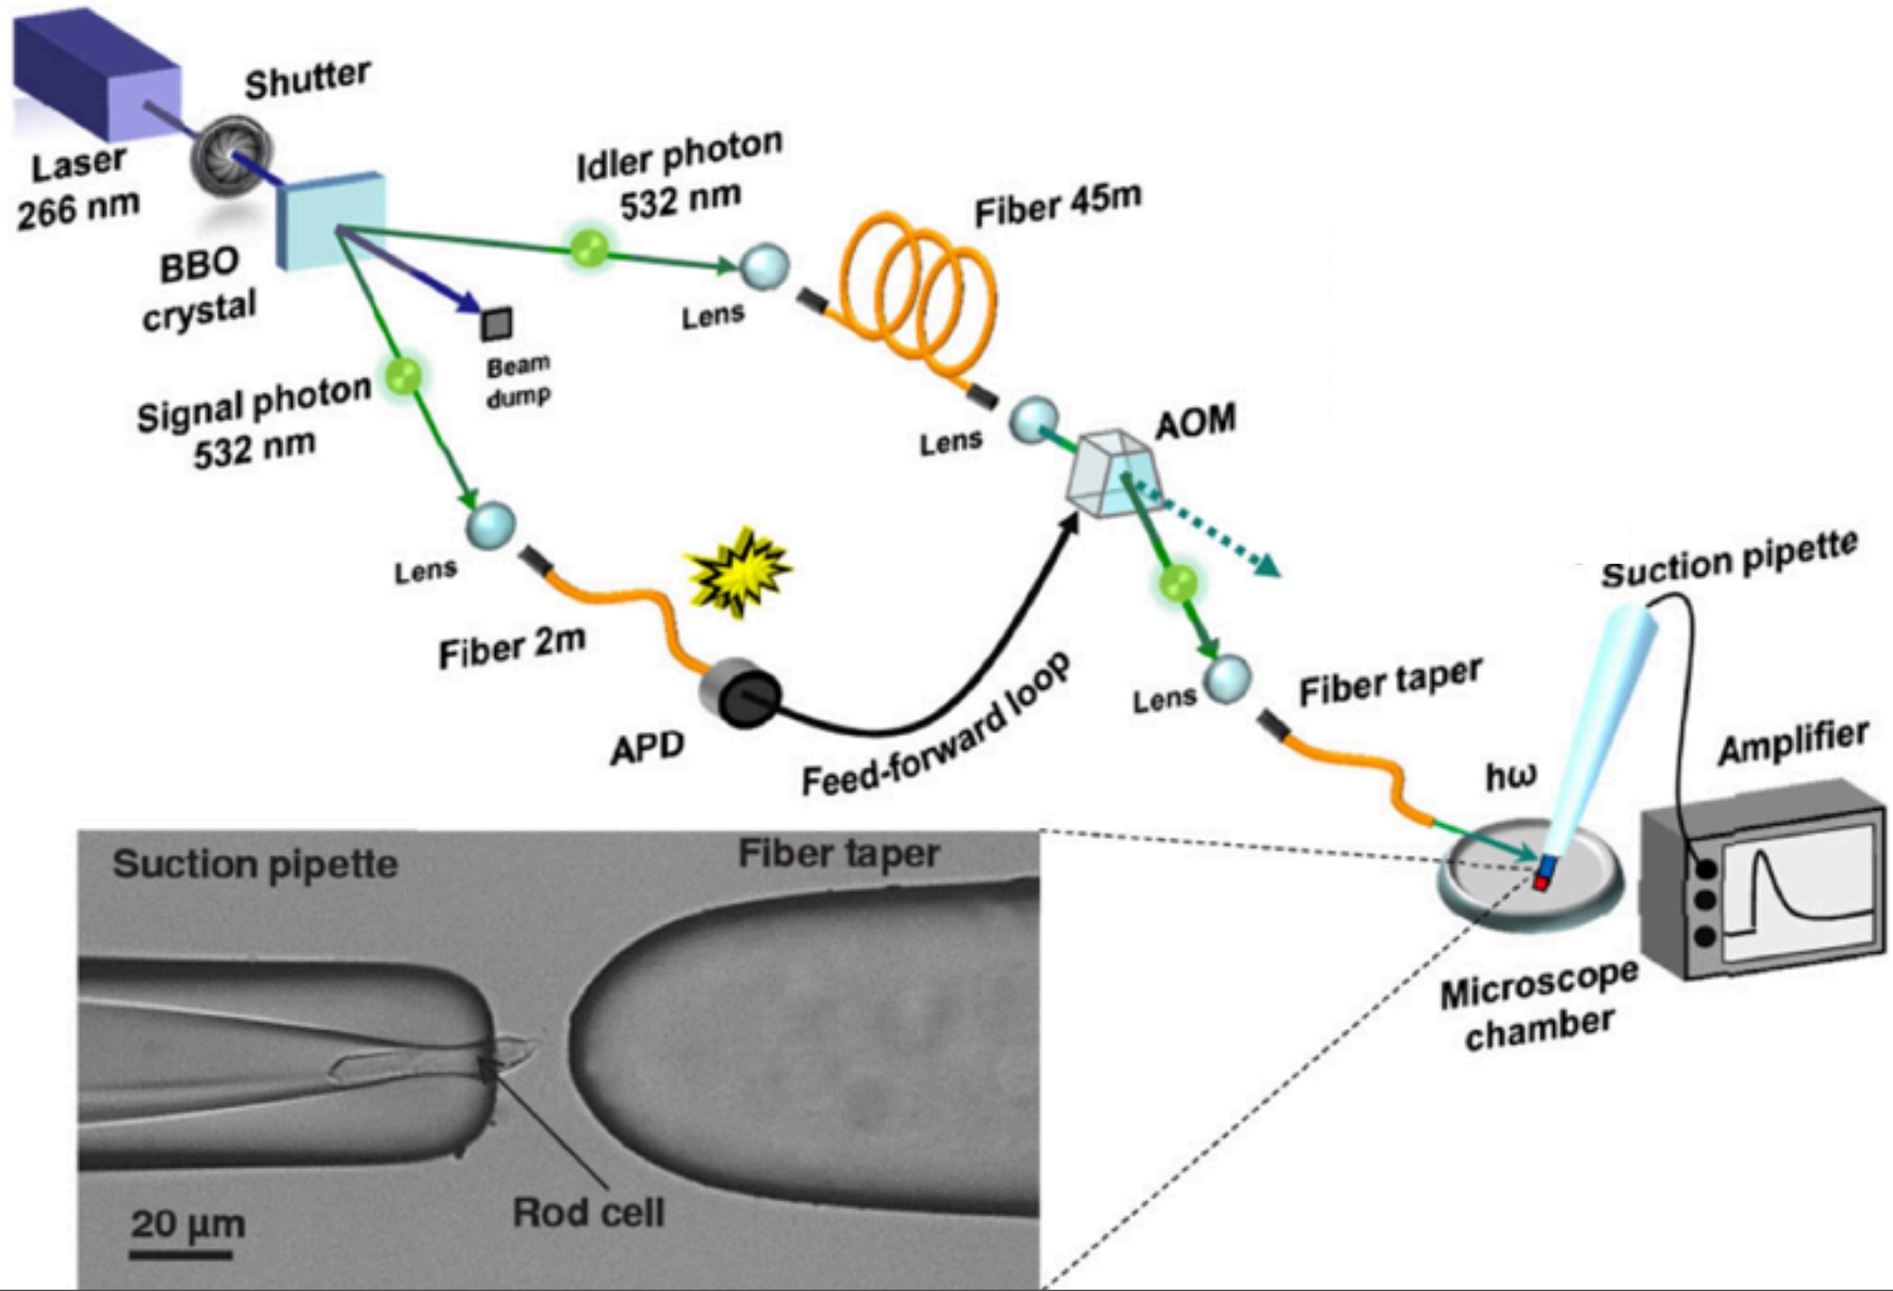

Phan, Krichevskiy, et al PRL 2014.

From P. Nelson, ***From Photon to Neuron: Light, Imaging, Vision*** (in preparation). Data Phan, Krichevskiy, et al PRL 2014.

Zero incoming photons:
pure continuous dark noise

One incoming photon:
some additional population of events

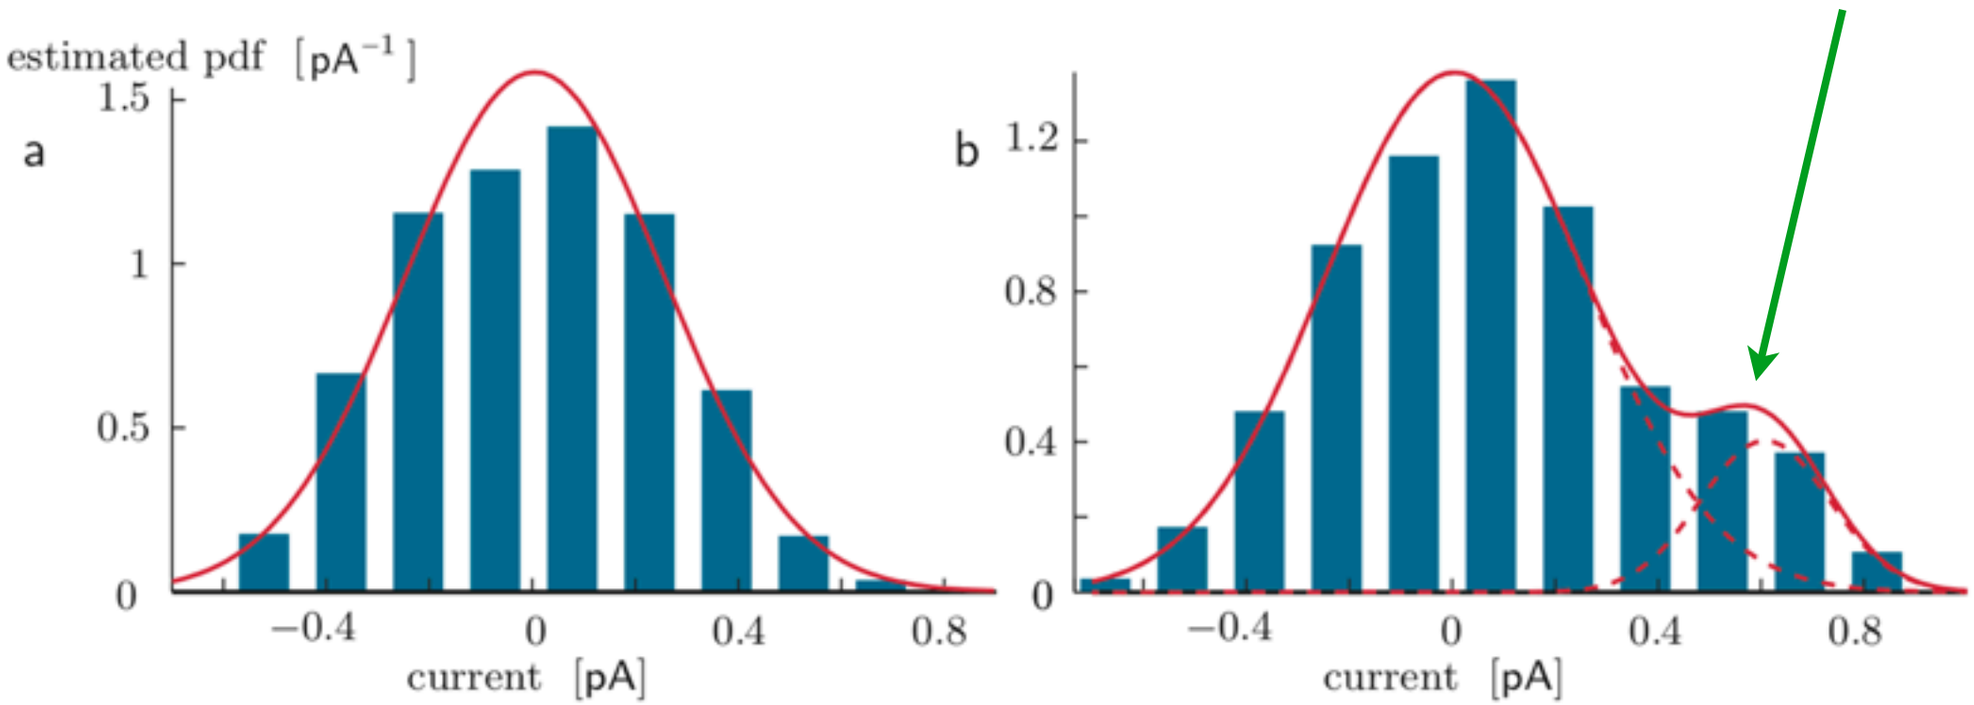

Now we know the real question

How can a cell monitor a hundred million light-sensing molecules, signal when *any one* catches a photon, and *not make* false-positive signals?
The answer's a long story, but my point is that nobody would have dreamed to *ask* the question – it's so crazy a claim – were it not for the pioneering experiments that said, crazy or not, it's true.

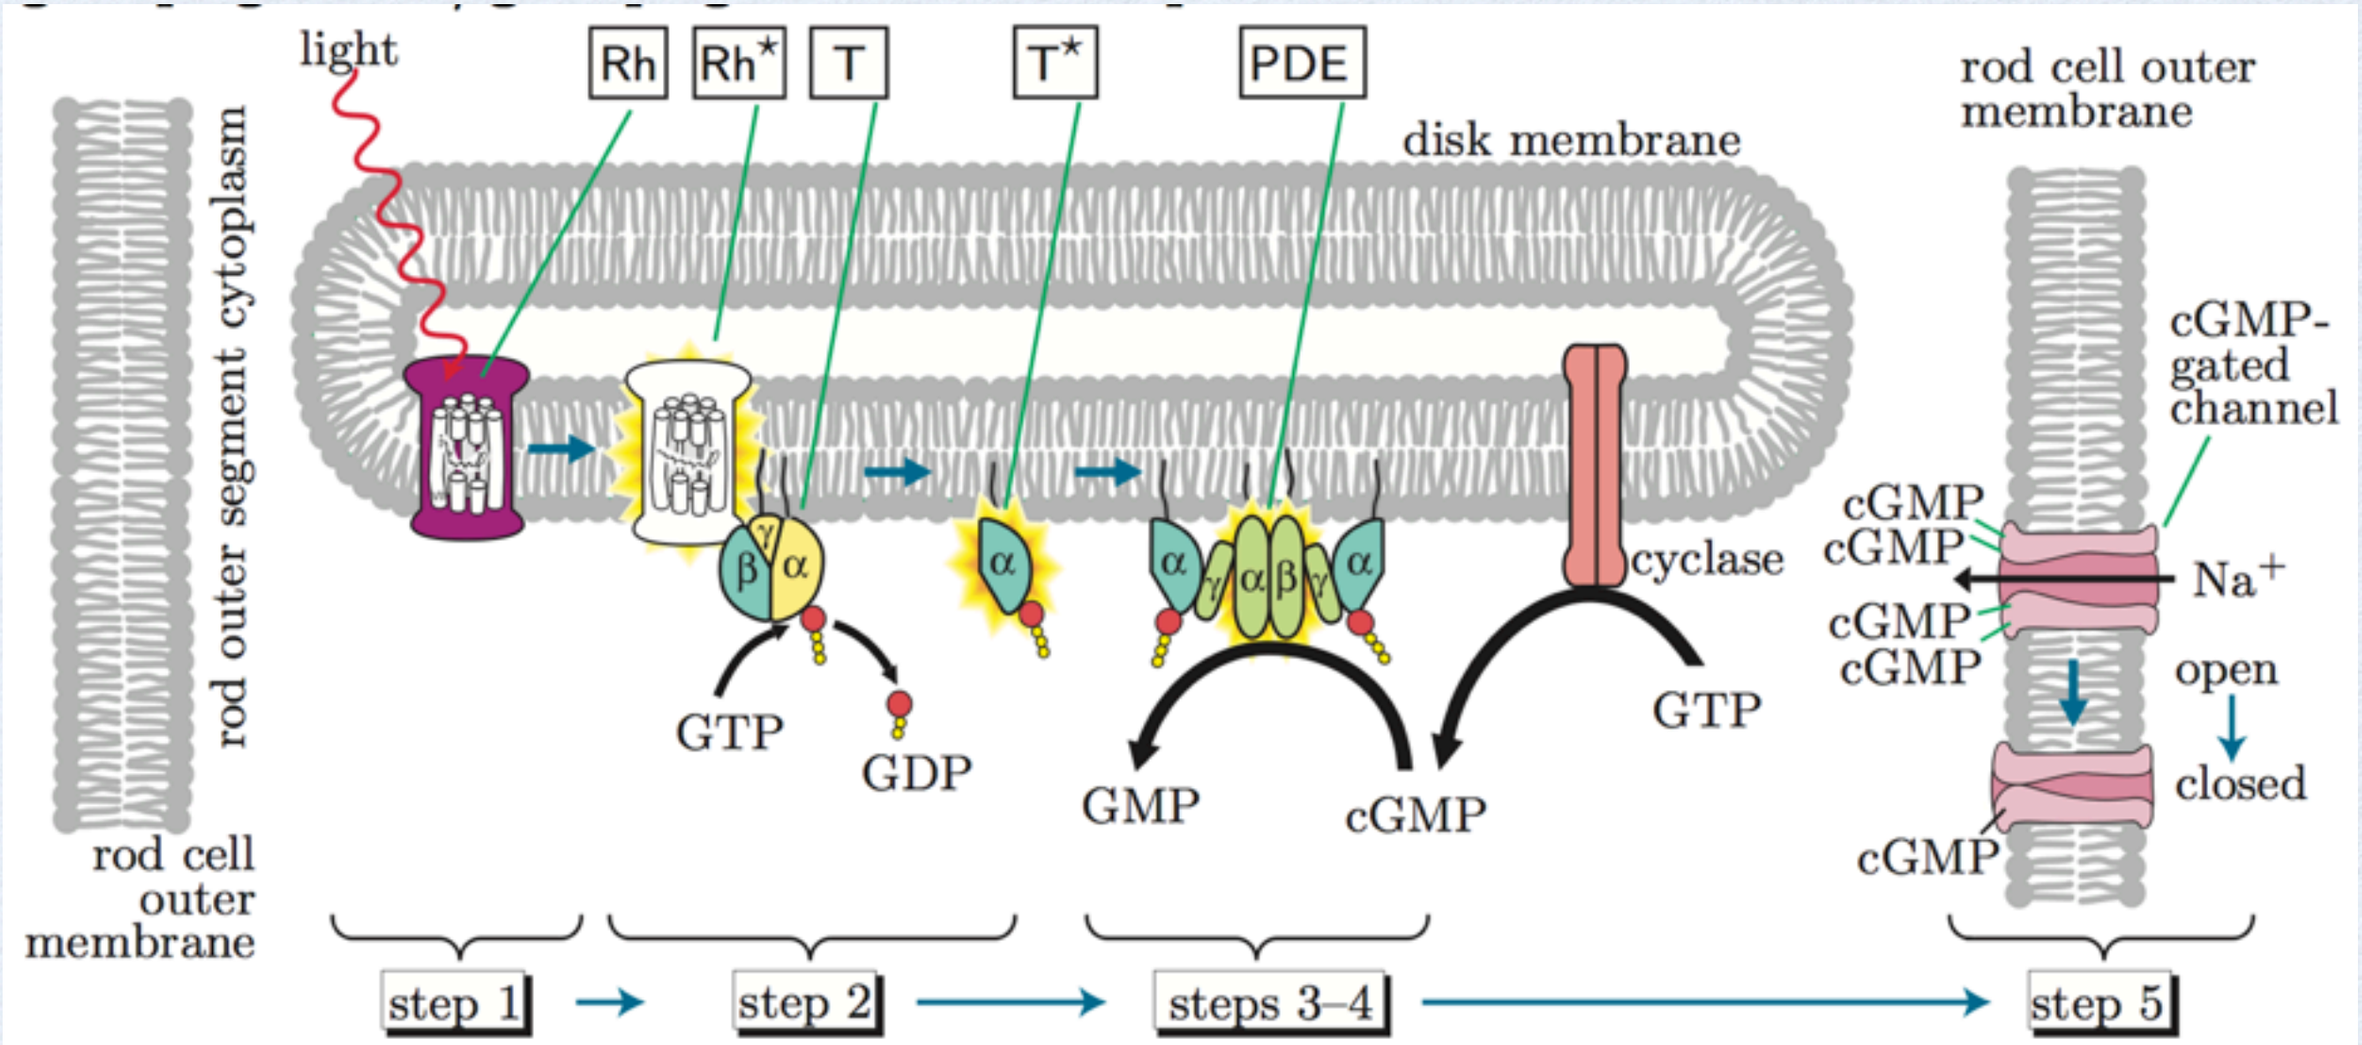

From P. Nelson, ***From Photon to Neuron: Light, Imaging, Vision*** (in preparation). Diagram courtesy Trevor Lamb.

The first synapse

Let's return to the higher level of signals and what we do with them.

Baylor et al also had a prescient insight: The very first synapse must discard some genuine photon signals. That's a great example of being *way ahead of your time* by using indirect reasoning: Indeed, recent work pegs that loss at 50%! That thresholding eliminates nearly all the continuous dark noise (which is now thought to be mostly spontaneous activation of phosphodiesterase).

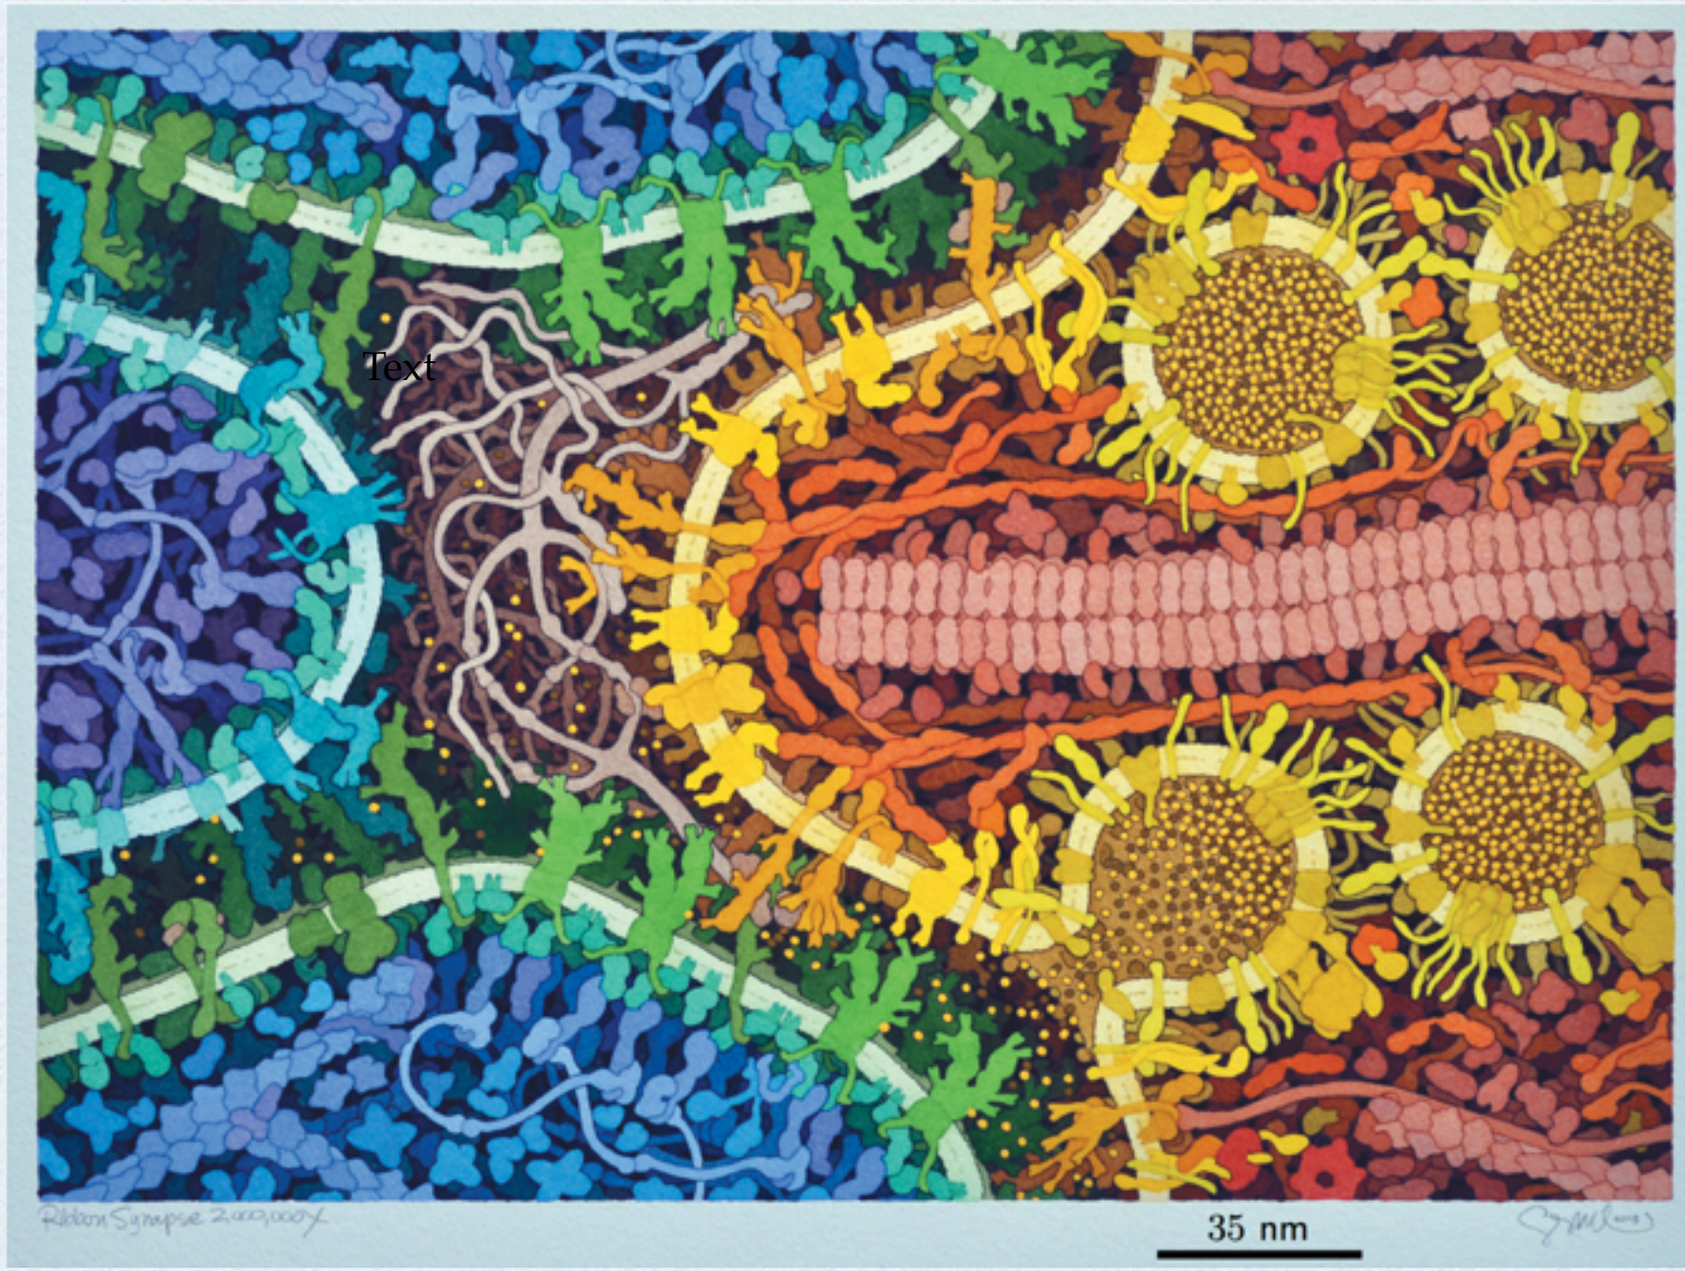

From P. Nelson, ***From Photon to Neuron: Light, Imaging, Vision*** (in preparation). Painting by David Goodsell.

The first synapse imposes a threshold

Because so many rod cells' signals are pooled, even a few extra false positive signals would be bad for our vision. So the first synapse sets a threshold to filter out continuous dark noise, cleaning up the signals fed to the next cells in the network (rod bipolar cells):

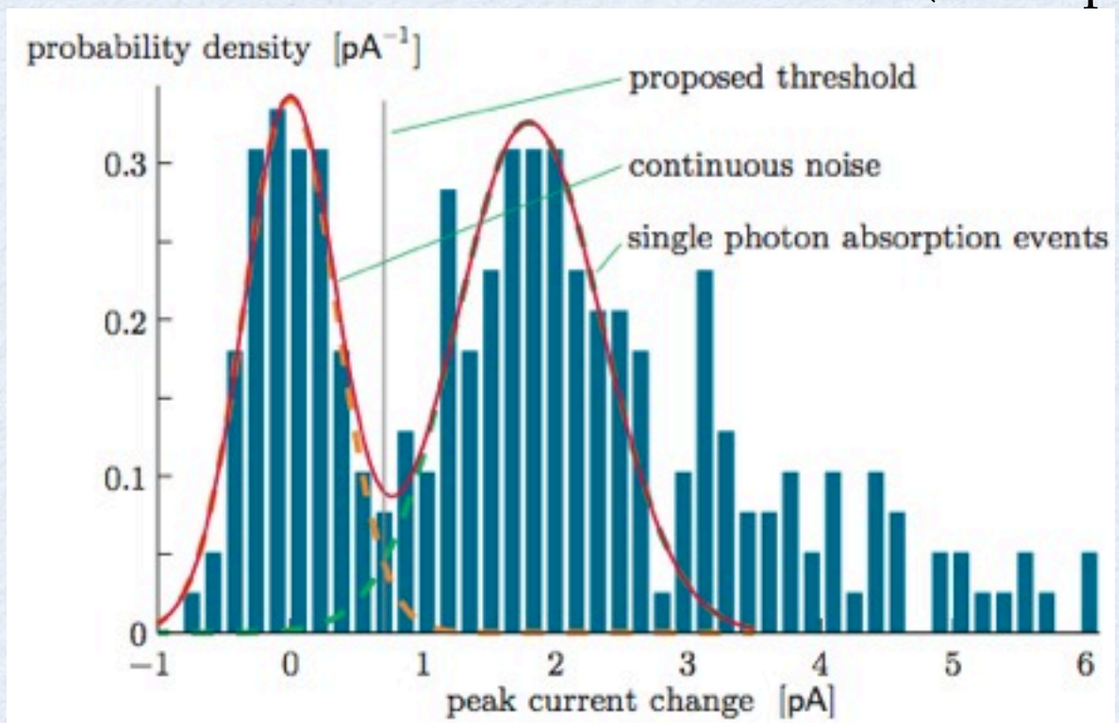

Recordings from mouse rod cells (before first synapse):

Recordings from mouse rod bipolar cells (after first synapse):

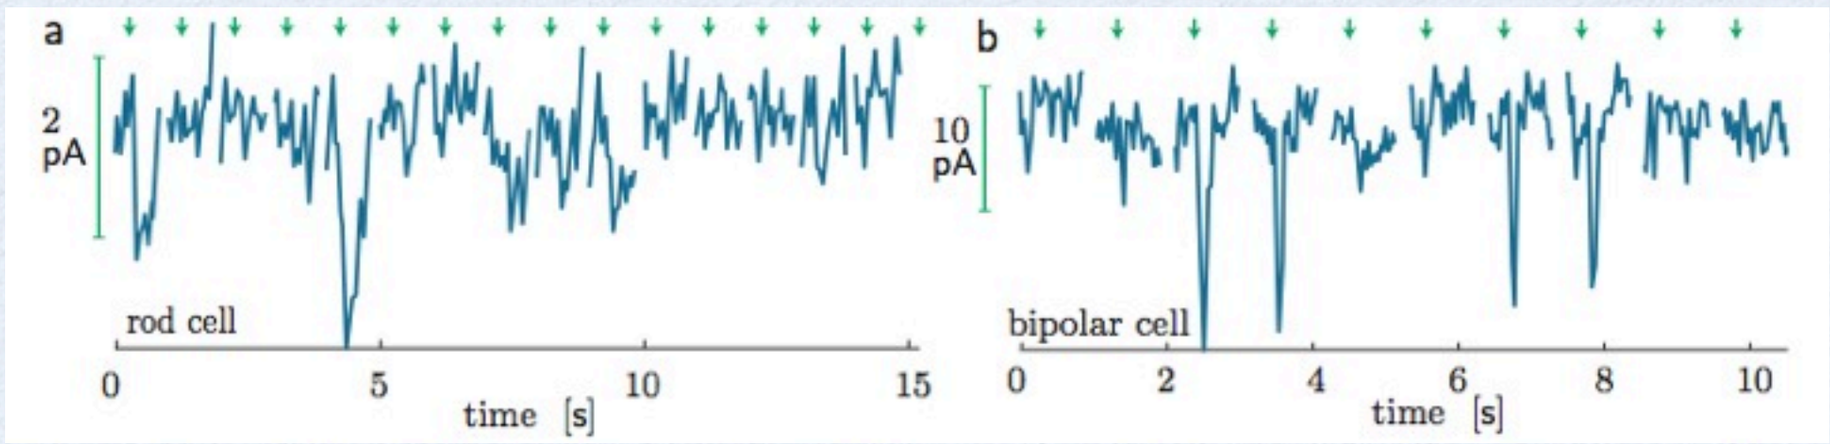

From P. Nelson, ***From Photon to Neuron: Light, Imaging, Vision*** (in preparation).

Summary: single cell

The pioneering single-cell experiments established that:

- One productive photon absorption is enough for a rod cell to signal.

The experiments also measured:

- The quantum catch for a rod cell *illuminated sideways*. (Not hard to convert this to the case of interest, axial illumination.)
- The mean rate of spontaneous isomerization.
- The rumble of continuous dark noise, which makes it necessary to threshold at the first synapse, even though it means discarding a lot of real signals.

It's time to ask whether Barlow's model (suitably amended to account for thresholding) can explain Sakitt's data, if constants that were initially fit parameters are frozen to their *measured* values.

Somehow this hadn't been done previously.

Plan

1. Indoctrination
2. Light
3. Vision: psychophysics
4. Vision: single-cell
5. Synthesis

The modeling challenge

Now we can do some serious science. It may be crazy-ambitious to attempt to model human psychophysics, corrupted as it is by all the complexity of the conscious and unconscious mind, with Barlow's ultra-simplified physical model–but let's try, after first extending in a simple way to account for continuous dark noise..

Barlow's model involves two modules:

Poisson processes thinned by various losses

Assume no random losses, just pool signals and apply quorum requirements

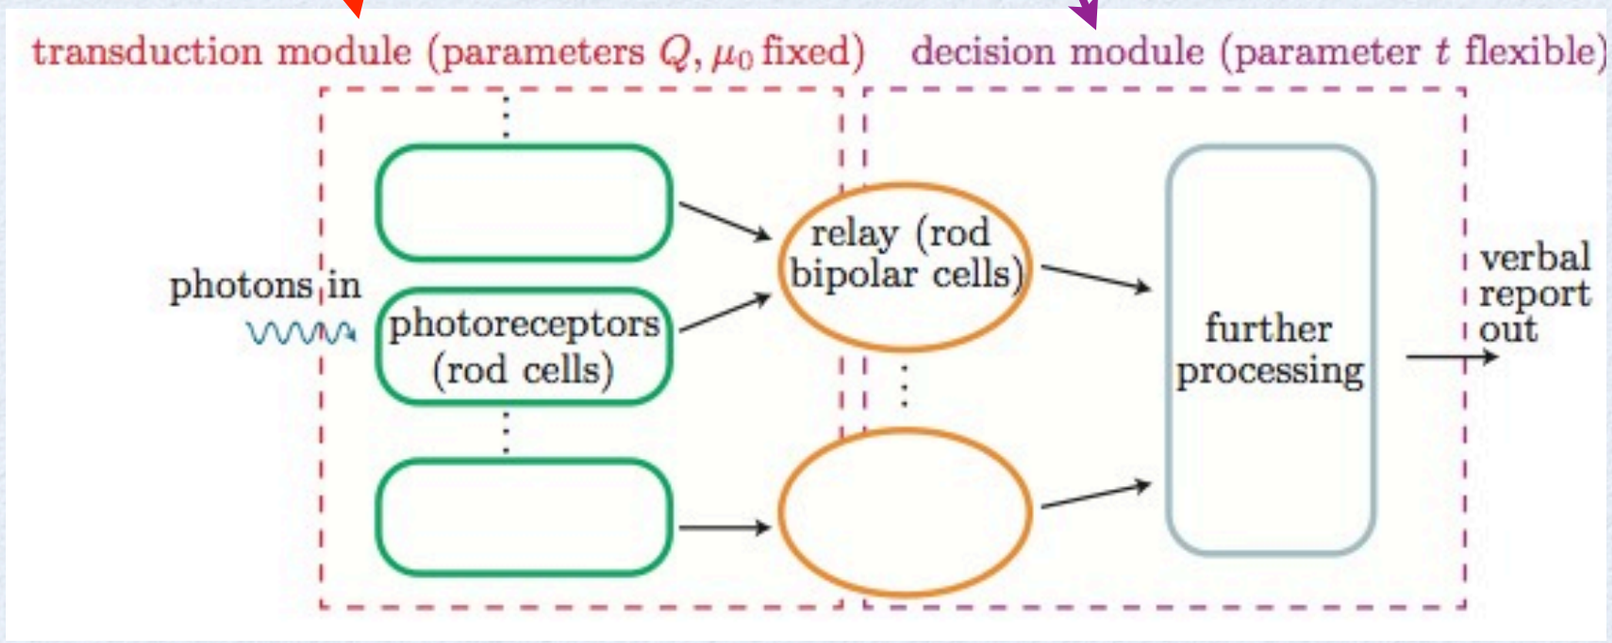

From P. Nelson, ***From Photon to Neuron: Light, Imaging, Vision*** (in preparation)..

Barlow model on the tightrope

- Single-cell physiology sets a *lower* bound on the visual system's loss and randomness.
- The incredibly good dim-light vision of an entire animal sets an *upper* bound.

Is there any room between these bounds, and if so, how much? Even if our model of the intermediate processing is naive, finding one model that works would establish a baseline for others.

From cartoon to model: outline

- In any time window, the signals crossing the first synapse are Poisson distributed, with mean related to the flash strength, reduced by absorption by ocular media, increased by spontaneous isomerization, and reduced further by rejection at the first synapse. Some loss and randomness is inevitable here. The parameters describing this "transduction module" are fixed.
- Some neural circuitry can count how many signals were generated anywhere in a "summation region" that includes the spot, in an "integration time" of around 200ms.
- Some neural circuitry applies a quorum requirement ("coincidence detector") to that total to decide if a "real" flash took place. "Don't bother the conscious mind unless at least t photon signals were generated in one integration time, in one summation region." At first we'll suppose this is a deterministic step. However, its parameter t is programmable according to the requirements of the task. There can even be multiple t's.

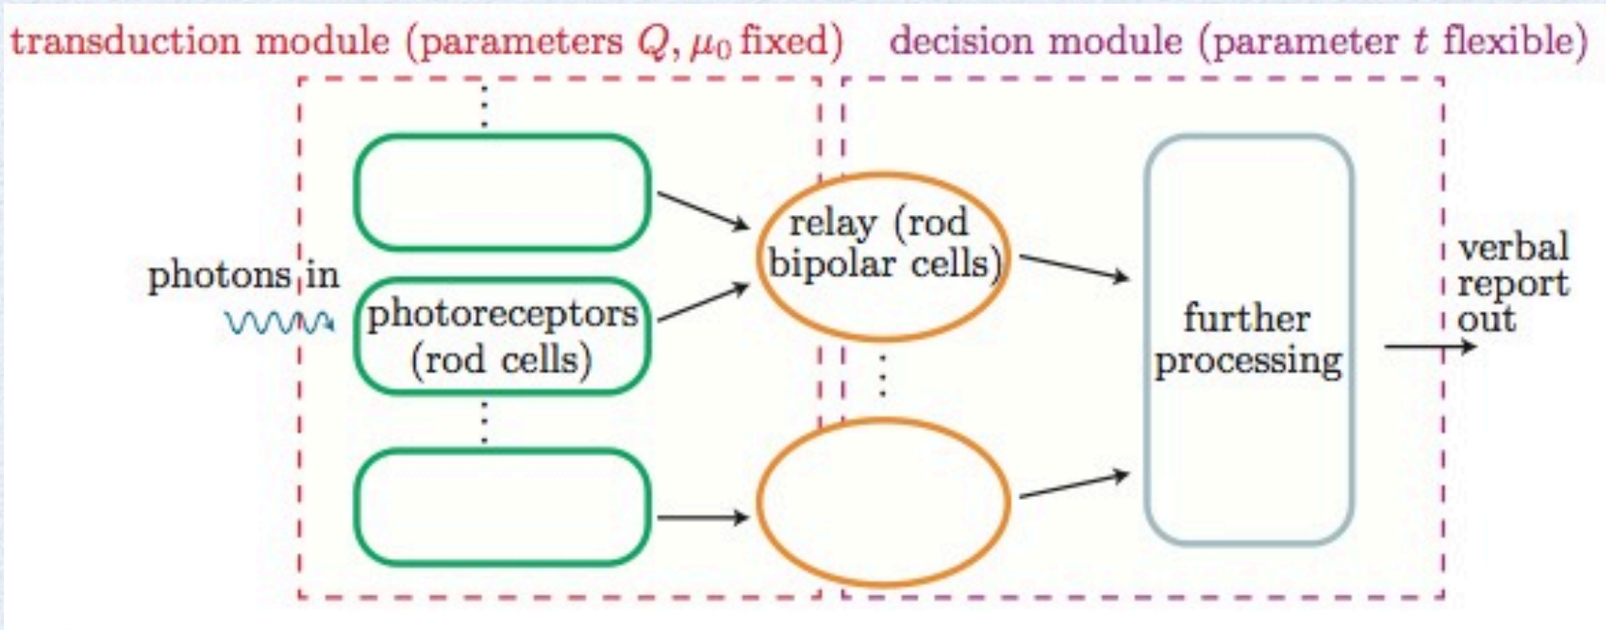

In detail

- Each flash sends in photons in a Poisson process with mean rate related to light intensity. ***We know that number.*** (It's purely physics.)
- Each rod cell in the illuminated spot has a fixed, small, probability of productively absorbing each incoming photon. ***We get that number ("quantum catch") from Baylor and some geometry.***
- Each rod cell initiates a "signal" in response to even one productive absorption. ***Now confirmed.***
- Some neural circuitry can "count" how many signals were generated anywhere in a "summation region" that includes the spot, in an "integration time" of around 200ms. ***We have quantitative estimates of the summation region size* and the integration time*.***
- But the light-sensitive molecules in rod cells can also spontaneously convert to their modified form, generating a dribble of false positive signals. ***We get that rate from Baylor.***
- The first synapse discards about half of all isomerization signals, both true and false ones. ***That fraction has been estimated*.***
- All that remains to be fit, then, are the quorum requirements t_i for each of the rankings in Sakitt's experiment. These must form an increasing sequence of small integers. ***And there are many more data points than t's, so the model is falsifiable.***

* = values are controversial

From P. Nelson, ***Physical Biology*** (in press). Data courtesy Heidi Hofer (2011).

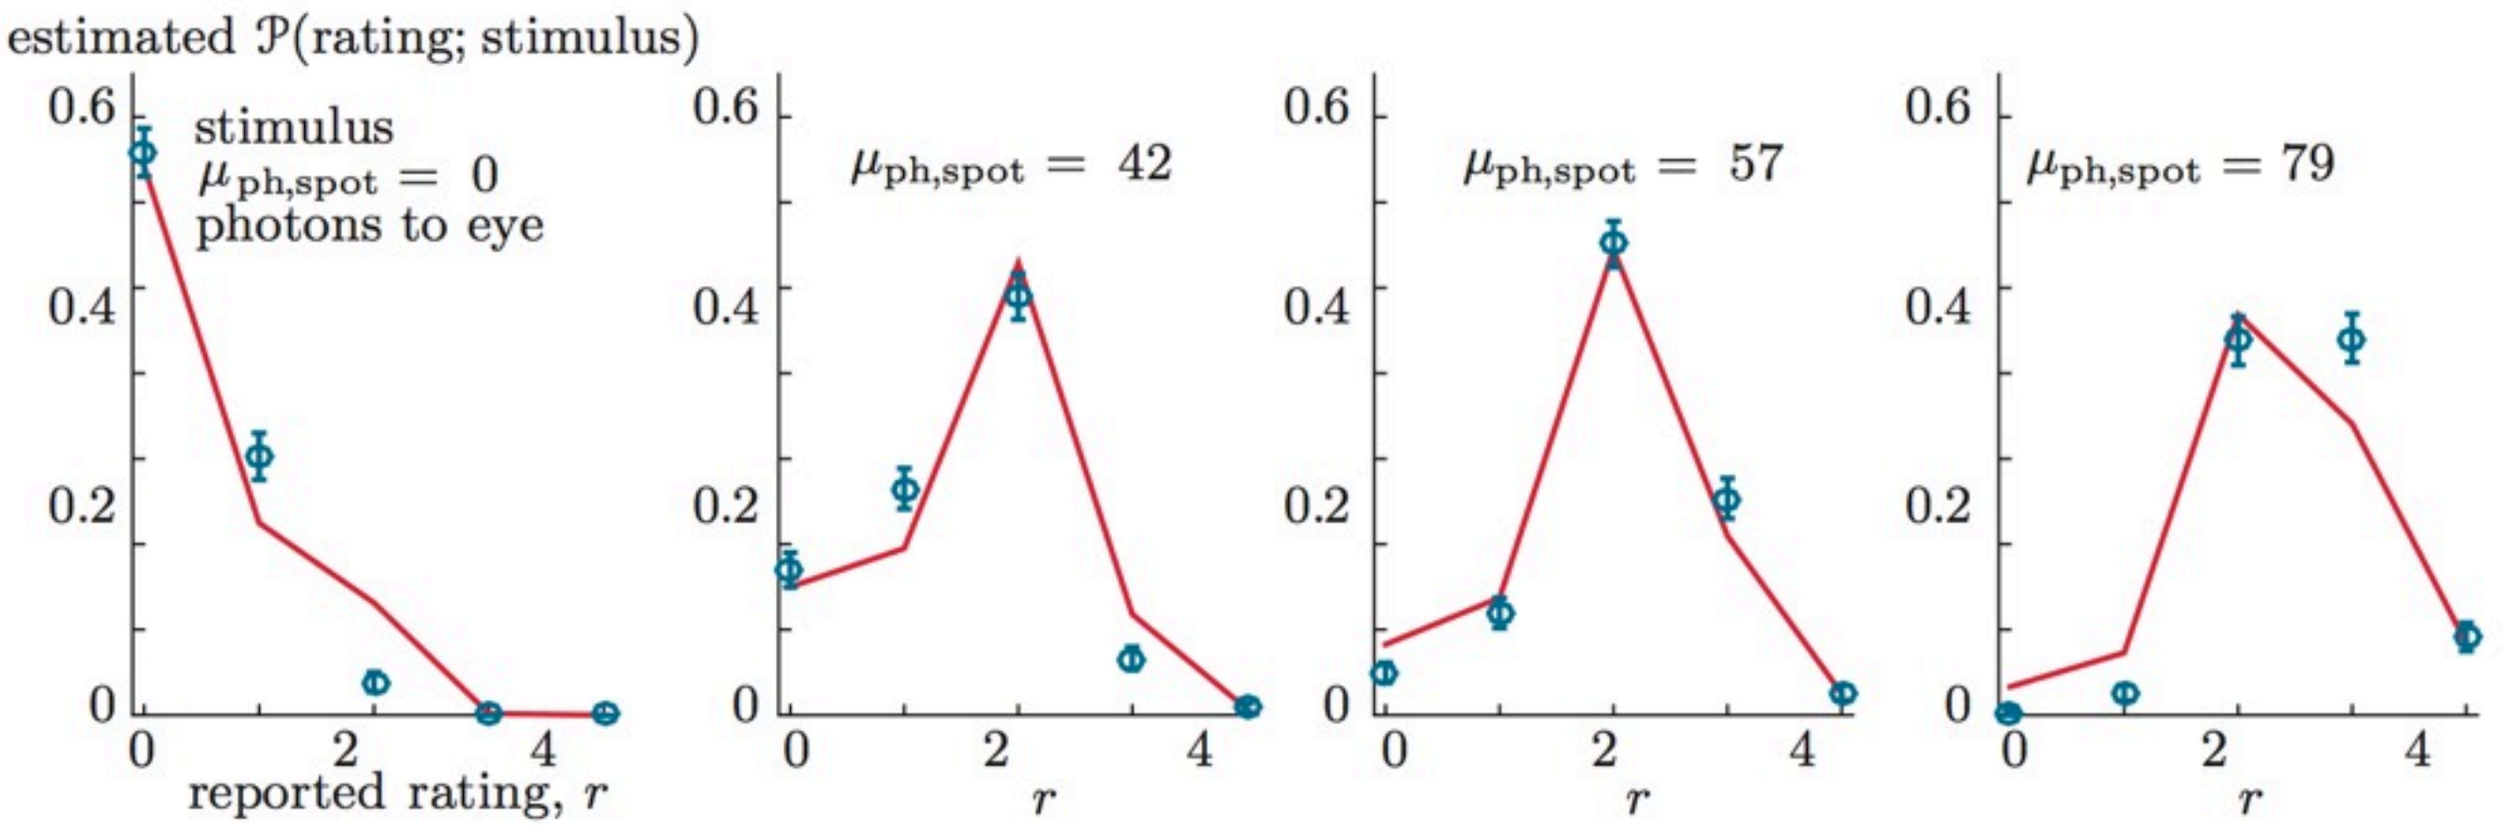

For this subject the best fit gives the first threshold as ***just two*** signals crossing the first synapse. That's enough to get a statistically measurable change in behavior. (Other subjects could also be fit, in some cases better with three, but certainly not, say, thirty.) (Sakitt's old data give similar fit.)

But for the highest rating ("certain"), the threshold was nine signals, close to our earlier rough estimate.

Moreover, ***no fit is possible*** if we change the model to add significant loss/noise after the first synapse. Everything beyond that point must be nearly perfect.

Understanding how that is possible is the ***next*** challenge/opportunity.

Summary: synthesis

It did prove possible to reconcile psychophysics with single-cell experiments in the context of a very simple model.

It was not obvious from the start that this would even be possible. For example, it is not obvious a priori that Barlow's separation into two "modules" is right.

Although today Barlow's picture of the "decision module" looks laughably oversimplified, we can regard this project as *characterizing* it in a way that any future, realistic model must explain.

That is, we found a way to take whole-organism (psychophysical) data and disentangle from it the first stage of processing, by itself fairly well understood now, so that we can focus on the next stages.

Technical details make this an interesting and challenging, but still doable, project. The best fitting model is the one that maximizes a likelihood function defined by the rather limited experimental dataset.

Plan

1. Indoctrination
2. Light
3. Vision: psychophysics
4. Vision: single-cell
5. Synthesis
6. Wrap

What is modeling, anyway?

Don't want to get all philosophical on you. I say, ***It's a Tetrahedron.*** Today I applied this approach to dim-light vision, but it's useful to think of any modeling challenge in this way:

```
for j=1:46,
   photons=20+5*j;
   mbar=q*photons;
   total=0;
   for i=mstar:50
      total=total+exp(-mbar)*(mbar^i)/factorial(i);end
```

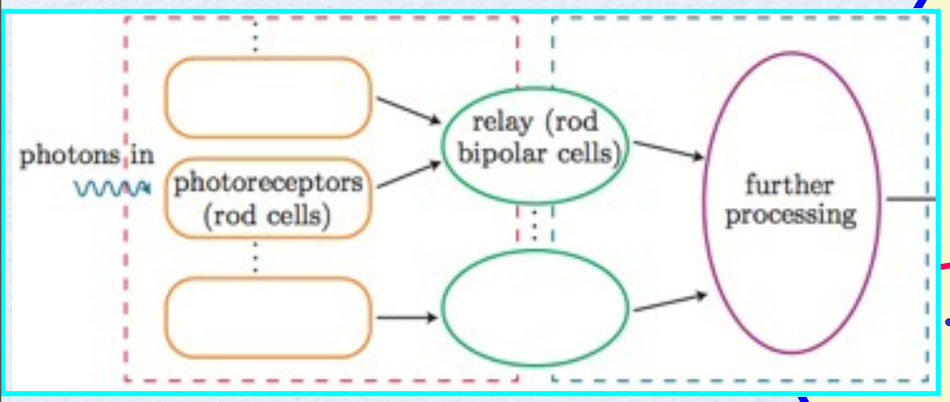

data

$$\mathcal{P}_{\text{see}}(\mu_{\text{ph,rod}}) = 1 - e^{-Q_{\text{rod,side}}\mu_{\text{ph,rod}}}$$

"Yadda, yadda... photons, yadda,... cumulative Poisson distribution...

Thanks

For these slides see:
www.physics.upenn.edu/~pcn
(or just google me)

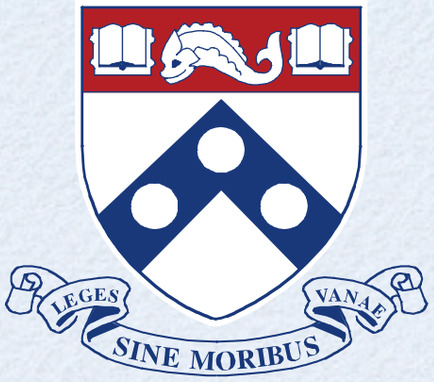

University of Pennsylvania

This material is in press at *Physical Biology*, and also the subject of a textbook in progress. See you at my poster.

Also see:

Physical models of living systems by PN (WH Freeman and Co., 2015) (www.physics.upenn.edu/biophys/PMLS).

A student's guide to MATLAB for physical modeling by Tom Dodson and PN (free at www.physics.upenn.edu/biophys/PMLS).

A student's guide to Python for physical modeling by Jesse Kinder and PN (Princeton University Press, 2015).

Let's take a moment to remember Kamal Shukla, NSF BIO

NSF NSEC