Consciousness and Cognition 117 (2024) 103624

Contents lists available at ScienceDirect

Consciousness and Cognition

journal homepage: www.elsevier.com/locate/concog

How to get rich from inflation

Simon Alexander Burns Brown [*]

Johns Hopkins University, United States
London School of Economics and Political Science, United Kingdom

A R T I C L E   I N F O

Keywords:
Consciousness
Richness of consciousness
Richness of experience
Subjective inflation
Temporal experience
Visual field
Periphery
Signal Detection Theory
Stability of experience
Perceptual reality monitoring

A B S T R A C T

We seem to have rich experience across our visual field. Yet we are surprisingly poor at tasks involving the periphery and low spatial attention. Recently, Lau and collaborators have argued that a phenomenon known as "subjective inflation" allows us to reconcile these phenomena. I show inflation is consistent with multiple interpretations, with starkly different consequences for richness and for theories of consciousness more broadly. What's more, we have only weak reasons favouring any of these interpretations over the others. I provisionally argue for an interpretation on which subjective experience is genuinely rich, but (in peripheral/unattended areas) unreliable as a guide to the external world. The main challenge for this view is that it appears to imply that experience in the periphery is not just unreliable but *unstable*. However, I argue that this consequence, while initially appearing unintuitive, is in fact plausible.

1. Introduction

What are you visually experiencing, at this very moment? You are focusing on these words. But that is not all you see. Perhaps in the background you also see an office, library, café, or idyllic beach scene, rich with details: books, plants, textured wallpaper, people, seagulls flapping across an azure sky, or whatever else populates the space in front of you. Or do you? A moment of visual experience gives us reliable access to far less of the world than many of us assume. Our performance on many tasks involving stimuli appearing in our periphery and in regions we are not attending to is surprisingly poor: indeed, we can fixate on an image for several hundred milliseconds without noticing radical alterations made to the stimulus in the periphery (Cohen et al., 2021). We can only simultaneously attend to and represent a small number of items in a way which is *access conscious* — available to a wide range of psychological systems, including introspection, reasoning, and report (Block, 2007). Given these limitations to access consciousness, how rich is visual phenomenal consciousness? That is, how many details are visually represented in a way which makes a difference to what it is like to be looking at this paper right now? Some, such as Bronfman, Jacobson, & Usher, 2018; Block, 2007, have argued for what I'll call "Traditional Richness": there is rich phenomenal experience across the visual field, outstripping what we can access and simultaneously attend to. Others, such as Naccache (2018), argue for what I will call "Traditional Sparseness", on which phenomenal consciousness is indeed sharply limited: while we may have access to some information about stimuli in our periphery, such as overall statistics about their distribution (Cohen et al., 2016), our visual experience does not include anything close to the level of detail across the visual field that we often assume. It is worth emphasising how radical a view this is: as Rosenholtz (2016, p. 438) points out, peripheral vision covers as much as 99.9 % of the visual field.

* Address: Centre for Philosophy of Natural and Social Science (CPNSS), Lakatos Building, The London School of Economics, Houghton St., London WC2A 2AE, United Kingdom.

E-mail addresses: simonabbrown@gmail.com, S.A.B.Brown@lse.ac.uk.

https://doi.org/10.1016/j.concog.2023.103624
Received 10 May 2023; Received in revised form 5 December 2023; Accepted 8 December 2023
Available online 26 December 2023
1053-8100/© 2023 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Not only are these issues of interest in their own right, but many have thought that they may help to decide between theories of consciousness. Theories which claim consciousness requires global broadcast (Dehaene & Naccache, 2001) or higher order representation (Lau & Rosenthal, 2011) are often thought to imply that consciousness is sparse. This is because they seem to imply that only content represented in frontal regions is conscious, and it is standardly thought that this content is sparse. By contrast, theories which associate phenomenal consciousness with certain kinds of local activity in more posterior regions (Lamme, 2010; Tononi, Boly, Massimini, & Koch, 2016), or even sub-cortical activity (Merker, 2007), seem to predict much more richness (Block, 2011).

Recently, Lau and collaborators have argued that we can occupy a middle ground in this debate, thanks to an empirical phenomenon dubbed "subjective inflation": signals for the periphery/unattended regions are lower quality, but subjects are more liberal in interpreting those signals. That is, a signal related to the periphery can be strong enough to elicit our reporting that we have seen something, even though a comparably strong signal pertaining to the fovea would not have been enough. Lau and his collaborators think that this phenomenon supports an account which captures the attractive features of both Traditional Richness and Traditional Sparseness: On such a view, we can accept evidence suggesting poor performance which seems to support Traditional Sparseness, whilst simultaneously doing justice to introspection-based intuitions favouring Traditional Richness (Odegaard et al., 2018; Knotts et al., 2019; Lau, 2022). Lau (2022) connects this idea with a theory of consciousness which he presents as a middle ground between the likes of Local Recurrence theory and Global Workspace theory — *Perceptual Reality Monitoring* theory.

I will argue there are multiple consistent interpretations of findings of subjective inflation, with strikingly different implications for the richness debate. Indeed, while Lau and collaborators sometimes characterise inflation as "denying overflow but explaining why it seems to occur" (Knotts et al., 2019, p. 53), I will argue, partly on the basis of Lau's own theory of consciousness, for a quite different interpretation: we do have experiences representing many details in our periphery, overflowing access consciousness. Contra Traditional Richness, such details will often be inaccurate, and the fact that they are conscious may well depend on sophisticated mechanisms located frontally in the brain. But Contra Traditional Sparseness, they are part of conscious visual experience. The resulting "Inflationary Richness" view has an interesting implication: phenomenology in the periphery is in constant flux. While this flux may seem unintuitive, I argue it is in fact plausibly an aspect of our experience we do not notice.

After explaining the relevant empirical findings in more detail (§2), §3 enumerates three different interpretations of inflation: Degrees of Consciousness, Phenomenal Threshold, and Inflationary Richness. §4 further distinguishes different varieties of Degrees of Consciousness, but argues that, although some suggest important avenues for further research, each seems to make implausible predictions about phenomenology. §5 argues that Phenomenal Threshold makes unwarranted and possibly untestable assumptions, while also struggling to actually explain why we have the impression of rich perceptual phenomenology throughout the visual field. §6 shows that Inflationary Richness avoids these problems and gives a satisfying account of our experience. It does make a prima facie implausible claim about the time-course of experience — that we should expect "flickering" phenomenology in the periphery. However, §7 shows that there may be ways of complicating the view to avoid making such predictions, and that even if future research does not vindicate such options empirically, flickering of the specific variety predicted by the view should be seen as a surprising discovery rather than as conflicting with our everyday experience. §8 deals with one problem for Degrees of Consciousness in more detail.

2. Inflation

For the purposes of this paper, "inflation" will be defined as subjects having more liberal decision criteria for reporting stimuli in areas of the visual field that are more eccentrically located and/or accorded lower attention.[1] For example, subjects' criteria are more liberal when asked whether they saw a grating or just mere noise, when the stimulus in question was presented in lower-attention (uncued) rather than high-attention (cued) (Rahnev et al., 2011), and peripheral rather than central, regions (Solovey et al., 2015).

Liberal criteria are defined in terms of Signal Detection Theory (SDT). The important ideas of SDT for our purposes can be summarised briefly. According to SDT, perceptual decisions (e.g. whether a stimulus of a certain kind is present) turn on whether a *level of activity* exceeds a *criterion* (see Fig. 1).

Activity is a variable which is available to the perceptual-decision-making system, and which correlates with the presence of the stimulus. Importantly, while the relevant activity could consist in a simple neural magnitude such as a single neuron's firing rate or average firing rate across a neural assembly, it could be (indeed is likely to be) implemented by details of the *patterns* of activity. All the framework requires is that whatever it is, activity is driven by two factors: whether there really is a stimulus (with a *signal* being generated as a result), and random noise. Noise makes it possible for there to be high levels of activity when there is no stimulus, and low levels of activity when there is a stimulus, but higher levels of activity are generally an indicator that the stimulus is present.

Criteria can be set at different levels: they are *liberal* if they result in "yes" decisions even for relatively low levels of activity, and *conservative* if "yes" requires a high level of activity. Liberal criteria cause more *hits* and fewer *misses* — more cases where there really is a signal present will provoke a "yes" rather than a "no" — but also more *false alarms* and fewer *correct rejections* — more cases where there is really only noise present will provoke a "yes" rather than a "no".

By analysing rates of hits, false alarms, etc., we can determine subjects' criteria and underlying signal to noise ratio (Hautus et al., 2021). Inflation (as I use the term here) is the phenomenon, discovered through such analyses, that for low attention stimuli (Li et al.,

[1] "Inflation" is sometimes used to name a phenomenon which is assumed to both underlie these liberal criteria, and independently explain other phenomena, such as lower metacognitive efficiency for peripheral stimuli (e.g. Odegaard et al., 2018). The relationships between these latter findings and liberal criteria will be discussed below.

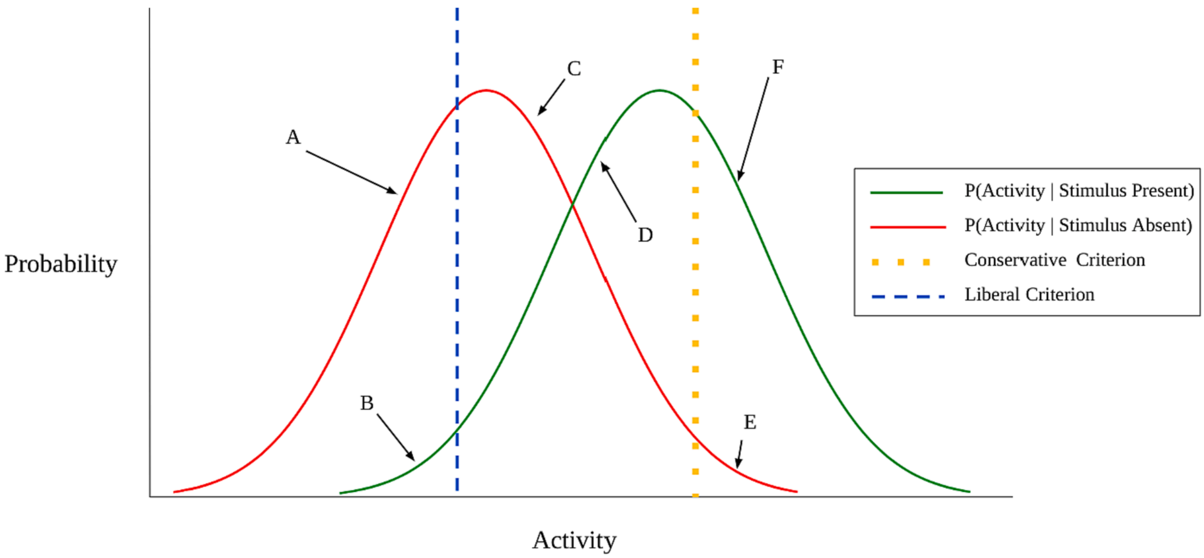

Fig. 1. Signal Detection Theory. Activity depends on both whether the stimulus is present and random noise. When the stimulus is present, higher activity is more likely, but not guaranteed. The subject reports seeing the stimulus when activity crosses a criterion. More conservative criteria require higher activity. A-F represent possible outcomes, with the distribution the point is drawn from and its position relative to each criterion determining whether it is a hit, miss etc. on each criterion. A: correct reject for both criteria; B: miss for both criteria; C: conservative correct reject, liberal false alarm; D: liberal hit, conservative miss; E: false alarm for both criteria; F: hit for both criteria.

2018; Rahnev et al., 2011), and peripheral vision (Odegaard et al., 2018; Solovey et al., 2015), subjects have more liberal criteria coupled with worse signal to noise ratio — resulting in a similar proportion of hits but more false alarms. Inflation has been found for detecting gratings as opposed to mere patches of noise, for detecting whether a patch of short lines included a group all oriented in the same direction (Odegaard et al., 2018), and whether a passing pedestrian's shirt was yellow in a driving game (Li et al., 2018). While, as I will discuss below, much more research is needed to establish the full extent of this phenomenon, here I will assume that it has been established in at least some cases, in order to investigate its implications for the richness debate.[2]

Low attention and peripheral stimuli are precisely those at stake in the richness debates: Traditional Richness claims we do experience (many) such stimuli; Traditional Sparseness denies this. For brevity, below I will only talk about the periphery; parallel considerations apply to low attention stimuli, unless otherwise noted.

Underlying the richness debate is poor performance for detecting and categorising such stimuli. Several mechanisms contribute to this poor performance. Some are very low-level. For example, with increasing eccentricity, densities of cones and retinal ganglion cells decline, receptive fields are broader, and cortical magnification factors decrease (i.e., less neural real estate is devoted to any given location) (Rosenholtz, 2016) — and this decline is even more marked when moving vertically from the centre as opposed to horizontally (Himmelberg, Winawer, & Carrasco, 2023). Spatial attention can counteract these effects for specific locations, for example shrinking receptive fields and boosting the gain of the relevant neurons, and recruiting neurons with neighbouring receptive fields — but at a cost to perception of neighbouring, unattended locations (Anton-Erxleben & Carrasco, 2013). In addition to these low-level phenomena, there is a "bottleneck" associated with working memory and some forms of attention, such that typically the details of only 3–4 items make it to working memory (Sperling, 1960), although the exact nature of this bottleneck is complex (Gross & Flombaum, 2017).

Advocates of Traditional Richness can point to evidence that the limitations arising from low-level mechanisms are not absolute: for example, we can still perform colour judgements in the periphery, provided stimuli are large enough (Hansen et al., 2009). And they will claim that the attentional bottleneck does not affect phenomenology itself, but only *access* consciousness (see also Amir et al., 2023 for a recent approach purporting to demonstrate phenomenal consciousness without access consciousness).

Furthermore, there is evidence for strong performance remaining intact for certain kinds of stimuli even when they are presented peripherally and with attention directed to a demanding task. For example, we seem to be able to determine whether a scene contains an animal or a vehicle (Li et al., 2002), and certain properties of faces, such as their gender (Reddy et al., 2004, Matthews et al., 2018) and even which individual faces we have seen before (Reddy et al., 2006), despite attention being engaged elsewhere. And Koch &

[2] Throughout this paper, I will talk about inflation as resulting from criteria *shifting* to become more liberal. The findings were originally explained with a model that posited a fixed criterion while activity distributions get broader — which, in many situations, results in the same profile of behaviour (e.g. Rahnev et al., 2011, Solovey et al., 2015). This model was also applied to other important phenomena (e.g. Ko & Lau, 2012). However, Denison et al. (2018) and Lee et al. (2023) convincingly argue, through a combination of intricate mathematical modelling and an innovative experimental paradigm, that a model where criteria do in fact shift provides a better explanation.

Tsuchiya (2007) place these findings in a broader theoretical context suggesting attention is not necessary for consciousness. However, these same papers imply that only some stimuli are immune to low attention in this way: their arguments rely on contrasting attention's lack of an effect on tasks like detecting faces, with its much more marked effects on tasks like determining whether a disc has green on the left and red on the right, or vice versa, and discriminating letters from one another. It remains to be established exactly which kinds of features are, and which are not, easily detectable in the absence of attention: one possibility is that faces and certain other stimuli are very special cases and attention is required for detecting most stimuli well; but another is that attention is required only for those stimuli whose detection requires certain kinds of binding (e.g. Green-Left). Nonetheless, for at least some stimuli performance is worse in the periphery, many more details are missed, and there seems to be less easily accessible phenomenal experience, than many people's subjective impressions might seem to suggest, and this paper is concerned with the impact of inflation on experience in cases of *this* sort.

Lau and his collaborators take liberal criteria for low attention, peripheral stimuli to explain why we are tempted to think they are experienced. For example, Odegaard et al., (2018, p. 2) gloss the implications of inflation as follows:

> "Inflation can be defined as the subjective overestimation of the reliability or quality of the sensory representations themselves … in inflation, the representations themselves are not necessarily filled in with details but are subjectively misestimated to be rich in content. Across the entire visual periphery, it is unlikely filling-in processes provide all the fine details in early sensory regions in a precise, pixelated representation instantly as soon as we view a scene"

Lau (2022, p. 92) puts it this way:

> "the details are never represented as such — neither in explicit nor compressed summary forms. Instead, we just interpret the sensory representations as if they are rich in detail, even when they are not"

Meanwhile, Cohen et al., (2021, p. 1705) characterise the view suggested by inflation as "observers systematically overestimate the richness of their perceptual experience". I take these passages to be endorsing an interpretation of inflation on which it supports thinking that phenomenal content is sparse but misinterpreted as rich. However, it is far from obvious that this is the only — or the best — interpretation of these results.

3. Alternative interpretations of inflation

There are multiple potential interpretations of inflation. Fundamentally, this is because it is not entirely clear how Signal Detection Theory maps onto consciousness. Different mappings imply different positions on inflation.

Denison et al. (2022) helpfully distinguish three possible mappings from SDT to consciousness, although they do not discuss inflation directly.[3] On "Option 1", the criterion only affects whether we *report* a stimulus: phenomenal experience is mapped to the activity level (cf. Abid, 2019; Phillips, 2021 fn. 20.). Activity level comes in degrees, so this option implies that consciousness of the stimulus does too: we are not simply aware or unaware of a stimulus, but aware to degree 0.8 one moment, 0.5 the next, and so on. "Option 2" maps experience to the activity level, but only when that activity crosses some threshold (which need not correspond to the SDT criterion); below that threshold, the representation of the stimulus is simply unconscious, rather than conscious to a low degree. "Option 3" eschews the idea of consciousness coming in degrees, and claims that there is experience when activity crosses the threshold, and not otherwise. While this taxonomy of mappings from SDT to consciousness is important, we will see that it conceals some importantly different options when it comes to interpreting inflation.

"Option 1" suggests a view we can call *Degrees of Consciousness*. According to this interpretation of inflation, we have low-degree awareness of many stimuli across the visual field (in addition to high-degree awareness of some stimuli); inflation's liberal criteria simply make us readier to *report* some of that low-degree awareness. Degrees of Consciousness vindicates aspects of both Traditional Richness and Traditional Sparsity. With Traditional Richness, our experience does include conscious representation of many different features throughout the visual field. However, with Traditional Sparsity, many of these representations have a much lower degree of consciousness than representations of centrally presented stimuli.

Option 3 can have very different implications for the interpretation of inflation, depending on how exactly the phenomenal threshold (the level of activity required for experience) relates to the measured report criteria found to be more liberal in inflation. If these liberal criteria reflect a more liberal phenomenal threshold, then inflation impacts conscious experience itself. If inflation results in a lower phenomenal threshold in the periphery, then even weak activity in the periphery will generate a great deal of experience. We will consciously represent many more details in the periphery than reliability of performance or signal-to-noise ratio alone indicate. We will have *rich*, albeit unreliable, conscious experience. Call such a view *Inflationary Richness*.

However, another possibility is that the *report criterion*, which determines the level of activity required to prompt a report (i.e. the criterion directly measured by analysing report behaviour using SDT), is independent of the *phenomenal threshold*, which determines whether activity is sufficient to generate experience (Fig. 2). On this view, it is possible that in inflation, the report criterion shifts, but

[3] Denison *et al.*'s "mapping" terminology allows these debates to remain largely neutral on certain metaphysical issues. SDT's variables are all measured by analysing behaviour, in turn presumably determined by neural states. As such, the view that consciousness "maps" to one of these elements is compatible with a range of views about the precise relationship between mental and neural states, from identity theory, to views which deny consciousness can be reduced to or identified with any physical process but accept that it metaphysically supervenes on such processes, to dualist theories on which consciousness merely reliably correlates with such processes.

the phenomenal threshold remains put. If so, then lower activity in the periphery implies we will have fewer stimuli represented in experience (insofar as the phenomenal threshold is conservative) — experience really will be sparse — despite our being willing to report the presence of many stimuli. Call this view *Phenomenal Threshold*. Both Phenomenal Threshold and Inflationary Richness are compatible with Denison *et al*.'s "Option 3": as Denison *et al*. do not discuss inflation, their options do not determine whether it involves the phenomenal threshold shifting.

Denison *et al*.'s Option 2 is similarly ambiguous about inflation. It claims that consciousness only occurs when activity is above some threshold, but comes in degrees above that threshold. If that threshold shifts during inflation, then the result is a hybrid between Inflationary Richness and Degrees of Consciousness: inflation directly results in richer experience in the periphery than there otherwise would be, but that experience comes in a relatively low degree. If the threshold does not shift during inflation (and if it is conservative), the result will be similar to Phenomenal Threshold: inflation may affect what we are willing to report, but does not result in even low-degree experience in the periphery.

Thus, multiple interpretations of inflation are available. It is not always clear which Lau and collaborators endorse. The quotes at the end of §2 above suggest adherence to either Phenomenal Threshold or Degrees of Consciousness, and the following seems to deny that experience is rich (although even under Phenomenal Threshold, the claim that details are "not actually represented anywhere in the brain" is puzzling):

> "richness may not be 'real'. Some of the details may not actually be represented anywhere in the brain…[Inflation suggests] a mechanistic account of the appearance of stable richness, despite its lack of actually rich content" (Lau, 2022, p. 94)

However, Lau sometimes seems to assume something closer to Inflationary Richness, for example when concluding from the fact that the criteria in question do not respond to changes in incentives "that they were not cognitive. Instead, they likely reflect what people actually consciously saw" (Lau, 2022, p. 95).

This prevarication arguably reflects the fact that these alternative interpretations of inflation — and their starkly different implications for the richness debate — have not been clearly distinguished in the literature. Having distinguished these interpretations, we can now ask about the strengths and weaknesses of each in turn. Denison *et al*. pursue several lines of evidence for deciding between their mappings, relating to metacognitive sensitivity, visibility ratings, and neural correlates of the different components. However, as they summarise their conclusions, this evidence (p. 256) "to date does not definitively rule out any of the possible mappings". As such, I will largely leave these lines of evidence aside, and focus on the strengths and weaknesses of the different interpretations of inflation specifically, beginning with Degrees of Consciousness.

4. Degrees of Consciousness

According to Degrees of Consciousness, we can identify the degree of consciousness of a stimulus with the level of activity associated with that stimulus — a continuous rather than dichotomous variable. It follows that we do have conscious representation of many different features throughout the visual field, but representations related to the periphery typically have much lower degrees of consciousness.

An immediate issue for this view is making sense of perceptual experiences coming in degrees. There are many options for doing so, but each faces challenges, and the most promising options have yet to be spelled out in detail.

There are some unproblematic senses in which experience comes in degrees, which will not fit the bill. For example, we can talk about "levels of consciousness" of an individual as a whole. This could mean their state — their being awake, sleepy, close to comatose etc. (Bayne et al., 2016). Or it could mean the amount of perceptual, emotional, and other experiences simultaneously undergone. But *overall* states of a subject cannot elucidate differences between simultaneous representations of individual stimuli in different parts of the visual field.

An alternative account of degrees of consciousness casts them as felt degrees of confidence that the feature is instantiated. On this view, the report criterion determines how confident you have to feel about the presence of the feature to deem it present. This proposal is a natural way to interpret SDT: it is key to the framework that higher activity correlates with a higher probability of the stimulus being present. Unsurprisingly, similar suggestions have been made for interpreting the phenomenology of related models (e.g. Pereira et al., 2022, p. 458).

This view seems to require shouldering the following commitments. First, every conscious perceptual state involves a feeling of confidence in the relevant sense. Next, this feeling can coexist with what we might ordinarily think of as feelings of *low* confidence, as when we lose trust in our senses thanks to reading Descartes (or to awareness of being in misleading lighting conditions). Such a view might claim that the relevant feeling of confidence is computed by a sub-personal perceptual monitoring system (Lau, 2022) operating independently from more intellectual feelings of doubt. However, if the sub-personally driven felt "confidence" dissociates from familiar feelings of doubt, more needs to be said about what exactly it does feel like.

We might instead appeal to representation of a degreed variable, such as brightness. However, while there is a *relationship* between brightness and level of activity, it is possible to have low degrees of activity representing a stimulus as bright, and high degrees of activity representing a stimulus as dim (this is why we can be very confident that there is a dim stimulus as opposed to either a bright stimulus or no stimulus at all). Similar issues face any proposal which reduces activity level to representation of any single feature. The best candidate for such a feature might derive from experiences' having a property of *phenomenal intensity*, which correlates with features like brightness, confidence etc., but is not exhausted by them. However, even if we can get an independent grip on phenomenal intensity, it would not help Degrees of Consciousness. One recent defender of the idea of pure phenomenal intensity explicitly claims that in inflation, we have *high* phenomenal intensity despite *low* activity (Morales, 2021, pp. 21–22).

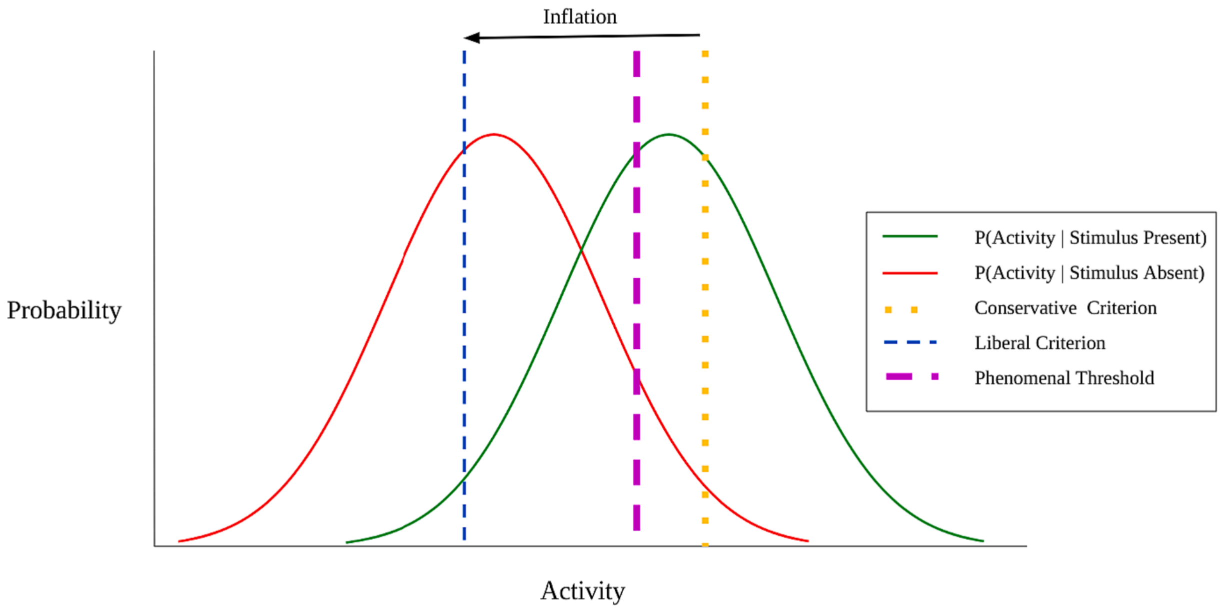

Fig. 2. Inflation in the Phenomenal Threshold view: the report criterion is more liberal than it otherwise would be, but the phenomenal threshold, hence consciousness, is unaffected.

A more promising approach posits that activity level shows up phenomenologically as the representation of *collections of details* of the item being detected. Intuitively, when we see many details of an object, we are more likely to say we saw it. For example, in Odegaard et al.'s (2018) Experiment 2, subjects are briefly shown a diamond composed of small lines, and then asked if there was a group of lines all tilted in the same direction: plausibly, willingness to report such a group is tied to how many of the lines' orientations are experienced. So perhaps the degree of consciousness of an item corresponds to how many of its details we see.

Statements by Lau and others suggest something like this idea. Lau (2022, p. 164) describes one case of an experience with a low activity level as follows:

> "Suppose we have the first-order sensory state of a cat that is rather impoverished. As a pictorial representation it lacks details. The content barely reflects a real cat. There is no way to tell if it is a Persian cat or a Bengal, or neither. So, an unbiased optimal system may well consider this to be noise."

Kouider et al. (2010), meanwhile, suggest that we have *partial awareness* of our periphery, understood as having access to some levels of the processing hierarchy (associated with representations of certain features), but not to others.

A related idea is that the degree of consciousness associated with some representation can be understood in terms of its *precision*: the narrowness of the range of values compatible with the representation being veridical. A very precise colour representation might narrow the stimulus down to being within a tiny sliver of the colour wheel (e.g. the scarlets); a less precise representation would be compatible with a broader range (e.g. all the reds).[4] We are *typically* readier to say that we saw something red if we are able to commit to a more specific shade of red. Perhaps *in general*, activity level for detecting a feature *F* corresponds to degree of consciousness of *F*, which corresponds to level of precision with respect to *F*.

This account is promising, but has yet to be carefully developed in a way which provides a satisfying, complete account of degrees of consciousness in the relevant sense. This would require committing to a specific account of the relevant kind of precision for each feature. Stazicker (2011) does develop a closely related account of experience in Sperling tasks in terms of how *determinate* the represented properties are, but this would need to be generalised. This may be complicated by slightly different accounts looking more promising for different features. For example, precision may mean something rather different in the cases of *red* and *cat*, given that colour is a continuous variable with a well-defined space of possibilities, while categorising objects as cats likely relies on a number of lower-level features along multiple dimensions, as well as broader context. A related issue is how closely such a view should tie precision and degrees of consciousness. It is very plausible that (i) low levels of activity related to a certain feature, (ii) low confidence that the feature is present, and (iii) only having a very imprecise representation of the feature, often go together. But to provide a mapping from SDT's activity level to experience, one would need to commit not just to a tight correlation, but to something more like *identity*, between (i)-(iii). It would need to rule out cases where we might have high activity but low precision or vice versa, or supplement the view with accounts of how these can come apart.

[4] This characterization of precision makes it distinct from vagueness (the fuzziness of the boundaries of the represented range), and (as Block, 2015 emphasises) from spatial resolution or acuity (being able to discriminate different colour patches which are very small).

Relevant ideas about precision have also been developed very differently, without relying on representational properties at all. Block (2015) argues that phenomenal precision sometimes varies even while representational precision is constant. His argument for this claim centres on the idea that the contrasts of an unattended Gabor patch and an attended Gabor patch might have different phenomenal precision, in the sense that the attended Gabor's contrast is experienced more "crisply", even though they are represented with the same amount of representational precision (i.e. there is one range of possible contrasts, such that all the representational content of experience specifies about each Gabor is that it is somewhere within that range). His argument for this is rather intricate, and rests on a number of controversial empirical and theoretical assumptions, so discussing it in detail is beyond the scope of this paper (see Fink, 2015 for a commentary). But even if the argument is successful, it is not clear that it can help us cash out what degrees of consciousness *are*. Block (2015, p. 8) denies that phenomenal precision can be defined, except through introspection on certain cases (in addition to the metaphor of "crispness" (11, 12, 35)). And the only cases of phenomenal and representational precision coming apart that he provides relate to the experience of contrast in the periphery and unattended stimuli — i.e. precisely the kinds of experience that we are trying to illuminate in the case of inflation. So again, this proposal would require further development to provide an account of degrees of consciousness.

The upshot of all this is that there are many quite different ways of articulating the idea of Degrees of Consciousness, but some are mutually incompatible with one another, and all await detailed development and defence. All will also face further issues.

Firstly, insofar as noise is constantly fluctuating, the degrees of activity associated with a stimulus will be constantly fluctuating. Degrees of Consciousness therefore seems to predict that all our phenomenology is constantly fluctuating (at least when this noise relates to the representation of lower-level properties like colour), in a way which seems introspectively implausible. We will return to this objection in detail below, as it also affects Inflationary Richness.

Another problem more unique to Degrees of Consciousness is that, partly due to noise, we have (low) levels of activity for a *vast* wealth of possible contents of perception. If we are looking at a group of lines all tilted in one direction, our visual system's activity will have a representation with this content, with a fairly high degree of activity, plus a slightly lower degree of activity associated with a group of lines all tiled in a slightly different direction, and so on. But it will also be effectively assigning tiny but non-zero probability to there being a group of circles, a skull shape, a solid shiny blue texture and all manner of other possibilities. Yet, while we should accept that such contents are weakly, briefly represented much of the time, Degrees of Consciousness predicts that these contents are *conscious*, albeit to a low degree. And this prediction seems phenomenologically implausible: is there really only a matter of degree between our consciousness of these possible stimuli and of the group of lines? The defender of Degrees of Consciousness will need to make some serious commitments here, either in their theory of how contents become conscious, or in revisionary claims about ordinary phenomenology. The options for Degrees of Consciousness here are discussed in more detail in Section 8 below.

5. Phenomenal Threshold

According to Phenomenal Threshold, we consciously represent a feature when activity relating to that feature crosses some threshold, which is not more liberal for the periphery, even though the report criterion is. If the phenomenal threshold is conservative, this implies sparse phenomenology. What advantages and disadvantages does such a view have?

One might worry that Phenomenal Threshold posits an additional component without any motivation — the phenomenal threshold itself. Would it not be more parsimonious to only have one threshold, the report criterion, and identify this with the phenomenal threshold, as a simple version of Inflationary Richness does?

The critic of Phenomenal Threshold should not lean too heavily on this worry. Firstly, Michel & Lau (2021, p 589) commit to such a threshold for independent reasons.[5] Secondly, everyone should agree that not *every* shift in response criterion results in a change to how easily representations can become conscious. If hits are rewarded with $1,000,000 and false alarms go unpenalised, you should be readier to say you see a stimulus than if hits earn $1 and false alarms cost $20. But we would not expect you to be correspondingly more likely to consciously experience the stimulus (Sánchez-Fuenzalida et al., 2023). It is natural to explain such cases with a phenomenal threshold remaining static while your response criterion shifts. So the *existence* of an independent phenomenal threshold should not be disputed (Denison et al., 2022, pp. 251–252; cf. Peters et al., 2016).

However, the view still must make substantive assumptions to generate any claims about the richness of consciousness. I characterised it as implying that phenomenology is sparse. But in fact, it only implies this given two questionable conditions.

Firstly, we need to assume that the threshold is relatively conservative. If the phenomenal threshold is liberal, we end up with a view that amounts to Traditional *Richness*: many stimuli are conscious, throughout the visual field, as activity frequently crosses the phenomenal threshold. Phenomenology is *richer* in the fovea than the periphery, but is rich everywhere. Indeed, the richness of phenomenology may be under-appreciated by experiments that rely on reports, given that the phenomenal threshold might be more liberal than the report criterion, especially in the fovea.

[5] They characterise the relevant criterion as a *perceptual* criterion, and so focus on whether the relevant phenomena are perceptual as opposed to cognitive, citing results like Rolfs et al.'s (2013) finding that some criterion effects lead to retinotopically specific adaptation, and Jin & Glickfield's (2019) finding that criteria could be shifted in mice just by optogenetically manipulating V1 neurons. In principle, it might be that such shifts relate to perception, whether conscious or unconscious, and that there is yet another criterion, which really is fixed, determining *conscious* perception. But absent some reason to assume this, such a possibility seems unlikely: the main reason to insist on a fixed phenomenal threshold is the idea that criteria are entirely determined by cognitively weighing up the costs and benefits of hits, false alarms etc., and the evidence they cite does seem to undermine this.

Secondly, we would need to assume that the phenomenal threshold, unlike the report criterion, is identical between the fovea and the periphery. If the phenomenal threshold did shift, then despite differing from the report criterion, we would end up with a view amounting to Inflationary Richness: conscious representation is achieved with less activity in the periphery, so despite lower signal and performance, stimuli are frequently consciously represented.

Unfortunately, it is extremely difficult to determine what the phenomenal threshold is and when it moves, insofar as it is independent of the response criterion. We *can* ask if a given shift in the observed response criterion can be explained by — or at least is responsive to — factors which everyone agrees do not affect the phenomenal threshold, such as changes in incentives. Inflation does *not* seem to respond to changes in incentives — or at least, to training and instructions not to use liberal criteria (Rahnev et al., 2011; Solovey et al., 2015). This rules out one alternative explanation for the shift in report criterion, raising the probability that it is due to a shift in phenomenal threshold. But other alternatives remain live options, such as those shifts in consciousness-irrelevant perceptual factors, or in cognitive factors outside voluntary control. One approach here might be to use other methods for measuring consciousness, such as the Perceptual Awareness Scale (Overgaard & Sandberg, 2021), in conjunction with the kinds of evidence Michel & Lau (2021) use to argue for perceptual criteria (see fn. 8 above), or a suitably modified version of the methods used by Sánchez-Fuenzalida et al., 2023, to try to tease out exactly what causes the shifts in criteria found in different cases of inflation. Peters et al. (2016) likewise point to cases in which conscious perception seems to be affected by criterion effects, and point to further strategies for uncovering which criteria are due to irrelevant features of subjects (see especially p. 6). However, disentangling the exact contributions to the criteria shifts in different cases of inflation remains a task for future research.

Another approach might be neural. Insofar as we can identify phenomenal consciousness with certain kinds of ignition, whether local (Fisch et al., 2009; Noy et al., 2015) or global (Dehaene & Naccache, 2001), we might look for evidence concerning whether the threshold for the relevant kind of ignition shifts during inflation. This would be useful to know, and evidence that the relevant threshold did shift would constitute converging evidence for the theory of consciousness in question and for Inflationary Richness as an interpretation. But a negative result would not be compelling without a clear account of what *is* driving the shift. And even a positive result would be far from decisive. For one thing, the data with respect to ignition are more complex and difficult to interpret than sometimes presented: Noy et al. (2015) find that consciousness as they measure it is associated with both local ignition and a kind of activity which spreads quickly through frontal areas but is less content-selective and powerful than global ignition, which they associate merely with reportability. More fundamentally, any move of this kind would rely on antecedently accepting that the relevant kind of activity is associated with phenomenal consciousness, rather than allowing results about inflation and richness to help choose between theories of consciousness.

There is a further problem with Phenomenal Threshold. If Inflation involves a shift in response criterion but not in phenomenal threshold, then experience is not rich, but we are more likely to say that we saw stimuli in the periphery. If so, how does inflation explain the impression of *perceptual* phenomenology? Why do we have the impression that we perceptually experience many details across the visual field, rather than just feeling more inclined to report that there are stimuli present without a corresponding perceptual experience?

One might answer that the liberal criterion, despite not making more representations phenomenally conscious, affects more than our inclination to report: it affects *access* consciousness, allowing us to *cognitively* represent many items across our visual field. Perhaps this cognitive richness is mistaken for rich perceptual consciousness. This view seems unpromising. It is rather introspectively implausible that we have a genuinely rich *non-perceptual* phenomenology, and sparse *perceptual* phenomenology. Even worse, this view struggles to account for the findings which motivated Traditional Sparsity in the first place: there is a bottleneck on conceptualization and access consciousness, such that cognition *cannot* represent such a rich array of features on the basis of sensation.

Alternatively, inflation might be connected to our sense of richness not through our representing many features across the visual field in any way, but through our having a metacognitive representation of first-order perceptual experience as a whole as rich. A general impression of our experience as rich might be comparable to other metacognitive feelings which can go wrong, such as a feeling of fluency of processing, or a hunch that a memory is old. Perhaps such a false metacognitive impression of our experience as rich is produced through our feeling inclined to report items in the periphery on many occasions — even though this inclination is not in fact due to our actually having difficult-to-access phenomenal experiences, but due to a liberal criterion applied to a weak signal. However, it is not clear *why* we would mistake mere inclination-to-report with actual experience so pervasively. And any explanation of this mistake, such as appealing to a general prejudice that experience is rich, is liable to simply appeal to resources which Traditional Sparsity theory already has at its disposal, suggesting that inflation does not do much work to advance the debate after all.

It is true that Odegaard et al. (2018) find poorer metacognitive efficiency in the periphery in addition to liberal criteria, which they take to indicate “a failure to introspect correctly” (p. 2). This is an interesting finding. But it is unclear that it tells in favour of Phenomenal Threshold. This is partly because, as Lau puts it, “empirical evidence at the level of metacognition is somewhat weak” compared to the evidence for liberal criteria, with his and other laboratories also finding null effects or even effects in the other direction (Lau, 2022, p. 94, Matthews et al., 2018). But it is also because it is far from clear what findings about the kind of “introspection” or “metacognition” involved in tracking how good performance is in specific tasks (i.e. what Odegaard et al. measured) really tell us about the kind of “introspection” involved in judging how many details are phenomenally consciously perceptually represented — including by inaccurate representations — across the visual field as a whole.

6. Inflationary Richness

Given the problems facing the other families of interpretations of inflation, we should take Inflationary Richness seriously. According to Inflationary Richness, the difference in criterion between fovea and periphery indicates a shift in how much activity is required for a stimulus to be phenomenally consciously represented. This implies we have *rich*, just *unreliable*, conscious experience. Increased susceptibility to noise means peripheral experience will frequently represent things which are not actually out there. But inaccurate experience is still experience, and there will be enough of it to give genuine richness: conscious perceptual representations of many features throughout the visual field.

It is worth emphasising a potential ambiguity in terminology here. In ordinary language, we often use "*S* sees *f*" in a way that implies that there really is an *f* which is reflecting light into *S*'s eyes. If many peripheral experiences are false alarms, it may therefore be *misleading* to say "we see many features throughout the visual field", even if we do have experiences as of features throughout the visual field. Relatedly, talk of "richness" sometimes conflates the amount of content represented and the reliability of representation. To be clear, the relevant sense of richness for our purposes — which Inflationary Richness secures — is conscious experience as of many features, throughout the visual field, even if this experience is unreliable.

This view avoids the need to make sense of degrees of consciousness, does not require positing a phenomenal threshold that is conservative irrespective of eccentricity, validates our introspective impression of genuinely rich *perceptual* phenomenology, and can neatly avoid claiming that extremely low-but-non-zero-activity bizarre representations enter into phenomenology at all. Indeed, insofar as it posits a flexible phenomenal threshold, it can allow for lower thresholds for some sorts of content than others, allowing a particularly neat solution to this problem of ubiquitous bizarre contents.

Inflationary Richness also coheres neatly with Lau's Perceptual Reality Monitoring account of consciousness (Lau, 2022). Roughly, this claims that a perceptual state *P* is conscious just in case a Reality Monitoring mechanism tags *P*, labelling it as being caused by external stimulation as opposed to mere noise, imagination, memory etc., on the basis of sub-personal processing that is inaccessible to the subject rather than an explicit higher order judgement. On this theory, if the threshold used by the Reality Monitor for determining whether to tag *P* as perception shifts, it will render *P* phenomenally conscious for lower activity levels. And it would make sense for the Reality Monitor to use more liberal criteria for the periphery: with a more conservative criterion, it would *almost never* tag peripheral representations as veridical perception, effectively leading the sensitivity to the outside world we do have to be ignored and thereby wasted. Note that here, unlike for Phenomenal Threshold, we need not simply assume a theory of consciousness and use this to choose a mapping to SDT in the particular case of inflation: rather, the idea is that the Reality Monitoring theory *predicts* inflation.

Inflationary Richness agrees with Traditional Richness on several important issues. Both posit rich phenomenal experience across the visual field, outstripping access consciousness. Both lend themselves to views on which the contents of sensory experience are associated neurally with sensory rather than frontal areas: Traditional Richness lends itself to views like Local Recurrence Theory, while Inflationary Richness lends itself to Perceptual Reality Monitoring Theory, on which frontal areas do have a role in rendering sensory contents conscious, but not in determining those contents themselves.

However, there are also important differences between Traditional and Inflationary Richness.

Firstly, the Perceptual Reality Monitoring account twinned with Inflationary Richness does give frontal regions an essential role in richness: they are presumably where criteria are shifted. Giving frontal regions such a role has important consequences. For example, it may lead us to attribute consciousness to a much narrower range of species than if consciousness were supported solely by sensory cortices or even non-cortical regions (Barron & Klein, 2016; Lamme, 2022; Lau, 2022, pp. 166–169). Even variants of Perceptual Reality Monitoring Theory which allowed for consciousness in species without Prefrontal Cortices would need to do so via less sophisticated forms of Reality Monitoring housed elsewhere, which might in turn not be sophisticated enough to produce inflation, leading (according to Inflationary Richness) to such animals having sparser experience than humans.

Secondly, Inflationary Richness has a more compelling account of poor performance in the periphery than Traditional Richness. Inflationary Richness posits that there is a weak signal but liberal bias for the periphery, which *predicts* poor performance even as it validates our intuition of rich experience. By contrast, Traditional Richness seems to predict that we should have good performance in the periphery, at least insofar as we can base our behaviour on the relevant (admittedly difficult to *access*) rich representations. Traditional Richness will posit an attentional bottleneck diminishing performance in many circumstances, and can allow *some* decline in richness and hence performance with eccentricity while still insisting it is somewhat rich. But to the degree that it posits richness that can drive behaviour, and does not (like Inflationary Richness) have an account implying that this experience is unreliable, it is forced to predict that the relevant experience should result in accurate behaviour.

Things are more complicated when it comes to another major criticism of Traditional Richness: it appears to predict *unstable* phenomenal experience, as it associates experience with activity. Activity will fluctuate, for several reasons. Every neural representation is subject to fluctuations in noise, and if more contents are represented neurally at any given time, there will be more noise-driven fluctuations in content during any given time period. In addition, there will be fluctuations in activity thanks to changes in attention, including fluctuations in spatial attention occurring periodically at around 4–8 Hz (Van Rullen, 2013, 2018). Fluctuations in attention might be thought to affect access consciousness rather than activity, but there is reason to think they will affect both: Even theorists who endorse the access-phenomenal distinction and use it to argue that the attentional bottleneck is compatible with richness, have been convinced by work by Carrasco and colleagues that attention also affects low-level activity (for details, see Block 2015, pp. 17-18).

If noise and attention imply that activity is unstable, tying experience to low-level sensory activity will result in experience dramatically varying from moment to moment. However, our experience "seems relatively stable. A subtle change in the perceptual task does not seem to change this overall impression so much." (Lau, 2022, p. 89). Yet this is not just a problem for Traditional Richness and Degrees of Consciousness. Insofar as Inflationary Richness also ties experience to low-level sensory activity, it seems to face a parallel problem.

However, the time-course of experience under Inflationary Richness is more complex.

7. The time course of experience

7.1. Predictions about the time course of experience

How stable is our experience over time, and are the different views under consideration compatible with different answers to this question?

We should immediately distinguish two sources of instability of experience. Firstly, as our eyes move, a stationary item's position relative to the visual field correspondingly changes, and with it sensitivity and activity relating to that item. If the criterion/threshold determining whether the object is experienced remains static through these changes, then some items that are seen when foveated will no longer be seen as they move into the periphery (and vice versa for items initially unseen in the periphery). Under Inflationary Richness, inflation could *reduce* this kind of instability.[6] Criteria shifting in a liberal direction *compensate* for activity declining as gaze moves away from an object, ensuring the subject continues to experience it. Likewise, insofar as spatial attention shifts around the visual field at 4–8 Hz in a way that causes working memory and performance to similarly oscillate (Van Rullen, 2013, 2018), it is possible that inflation could somewhat smooth the effects of these cycles on felt experience itself. On views where these criteria do not affect experience, by contrast, at best only reports will exhibit this sort of smoothing.

A second source of fluctuations in experience may be exacerbated by inflation under Inflationary Richness. Fluctuating internal noise could induce changes in experience even of a static object under a fixed gaze. Liberalising criteria in the periphery, where the signal-to-noise ratio is worse, will lead to more false alarms, which, for Inflationary Richness, will mean more conscious experiences driven solely by noise. If noise varies more than signal over time, and especially if it "flickers", inflation will cause corresponding "flickering" in experience in the periphery.

To visualise this point, we can plot the components of SDT over time (Fig. 3). SDT itself is not dynamic: we need to add further assumptions to generate predictions about the time course of perceptual decision-making and experience. The simplest approach (the impact of alternative assumptions is explored below) assumes that the criterion remains fixed over time (for a given kind of stimulus and location)[7]; and that at each time point, the level of activity is drawn from fixed distributions depending only on whether the stimulus is present or absent at that moment, and the stimulus is deemed present by the system if activity crosses the criterion at that moment.

Given these assumptions, we can then apply the three views about how SDT maps onto experience. Under Degrees of Consciousness, the degree to which the stimulus is experienced fluctuates at each moment with activity (and hence with noise). Under Inflationary Richness, the stimulus is conscious whenever activity crosses the criterion. Under Phenomenal Threshold, there is conscious experience as of the stimulus whenever activity crosses the phenomenal threshold, which is conservative despite a liberal *decision* criterion.

What do these views imply about the time course of experience?

Under Degrees of Consciousness, conscious experience tracks activity (the pink line in Fig. 3), and hence is pushed around by noise: it fluctuates every moment, and in that sense is highly unstable. However, this fluctuation will generally not be between experience and complete absence of experience: instead, there will be fluctuation between different intermediate degrees of consciousness.

Under Inflationary Richness, conscious experience occurs every time activity crosses the liberal criterion, and ceases when it dips back below it. Due to the periphery's inflated liberal criterion and poor signal-to-noise ratio, experience at each particular moment strongly depends on fluctuating noise levels. Consequently, there will be frequent 'flickering' — where activity is briefly pushed above criterion (e.g. see t_3 in Fig. 3). Overall, experience will be rich, in the sense that it represents many features throughout the visual field; but in the periphery, experiences will be unstable.

This instability differs to that predicted by Degrees of Consciousness. It consists in flickering between "on" and "off", rather than between intermediate levels of experience. Further, it will be slightly less frequent: some changes in activity will not result in activity rising above or falling below the criterion. But more strikingly, on Inflationary Richness, instability will be concentrated in the periphery: experience in the fovea will be less susceptible to fluctuations in noise and hence more stable (this also sets the view apart from Traditional Richness).

Phenomenal Threshold predicts a much more stable experience than either of the other two views, in the sense that it predicts experience is consistently sparse: there will rarely be experience in the periphery, and *ipso facto* there will rarely be changes between experience and absence of experience. Noise will still affect exactly when activity crosses the relevant threshold, so there may still be situations where we get some "flickers", as at t_4 in Fig. 3. But such flickers will be far rarer than under Inflationary Richness. On the other hand, there will be more instability across shifts in gaze and attention than on the other views, as such shifts will be more likely to push a signal below a static phenomenal threshold than below Inflationary Richness's liberalised threshold.

What considerations tell between these three views of the time course of experience?

We might initially think that Phenomenal Threshold clearly does better than its rivals on this score. Introspectively, we report *stable*

[6] A complication: if inflation also causes additional eye movements, e.g. by making us more likely to see items in our periphery which then attract our gaze, the overall result may still be more instability of this kind. Nonetheless, any given eye movement will be less likely to result in changes to what is experienced.

[7] This abstracts from fluctuations in gaze and attention, to show that even with stable gaze and attention, there will be fluctuations in experience. If we were to incorporate attention cycles, for example, we would see systematic oscillations in signal strength, which given inflation would be matched by corresponding oscillations in criteria.

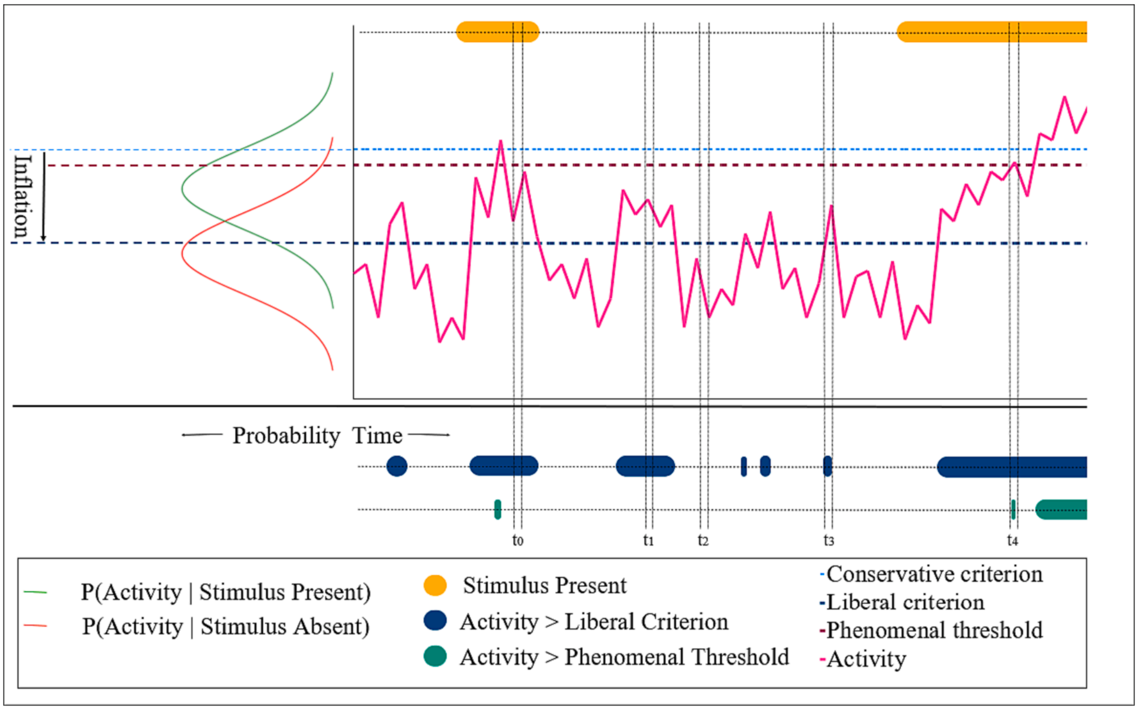

Fig. 3. The time course of experience: Left: the probability distributions from which activity is drawn, as in Figs. 1 & 2. Right: the time course of the stimulus, activity, and experience (fictional "data", purely for illustrative purposes). The stimulus is present during the periods covered by the yellow blobs, otherwise absent. Activity (pink line) fluctuates moment to moment, being drawn anew each moment according to (a) whether the stimulus is present, and (b) fluctuating noise. On Degrees of Consciousness, the degree of consciousness of the stimulus at a moment tracks activity. Dark blue blobs mark when activity is above the liberal criterion, i.e. when experience occurs according to Inflationary Richness. Green blobs mark where activity is above the phenomenal threshold, i.e. when experience occurs according to Phenomenal Threshold. t_0: veridical experience for Inflationary Richness, no experience (miss) for Phenomenal Threshold, relatively weak experience for Degrees of Consciousness. t_1: non-veridical experience for Inflationary Richness, no experience (correct reject) for Phenomenal Threshold, relatively weak experience for Degrees of Consciousness. t_2: no experience (correct reject) for both Inflationary Richness and Phenomenal Threshold, weak experience for Degrees of Consciousness. t_3: "flicker" for Inflationary Richness, small fluctuation for Degrees of Consciousness, and stable absence of experience for Phenomenal Threshold. t_4: sustained veridical experience for Inflationary Richness, fluctuations for Degrees of Consciousness, and "flicker" for Phenomenal Threshold. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

experience. This phenomenon should be particularly worrying for Inflationary Richness: not only does it seem to predict even less stable experience than Degrees of Consciousness (in the sense that the latter only predicts slight fluctuations in degree, where Inflationary Richness predicts flickering between "on" and "off"), but one of the main attractions of Inflationary Richness is supposed to be validating our introspective impression of genuinely rich perceptual phenomenology. So conflict with introspection seems not just to be a point against Inflationary Richness in itself, but to undercut its supposed strengths. This is a serious problem indeed.

However, an advocate of Inflationary Richness has two potential kinds of response available: offer a more complicated model which is in the spirit of Inflationary Richness, but which avoids instability, or show that its prediction of instability in the periphery is not so introspectively implausible after all.

7.2. Complicating the model to avoid instability

One might hope that turning to models on which there is less short-term fluctuation in noise might help. The experimental evidence supporting inflation implies that activity in the periphery is heavily affected by noise, and that noise fluctuates considerably between trials, but it is compatible with the relevant kind of noise evolving smoothly over shorter timescales: indeed, it is in principle compatible with noise being completely stable for durations up to the length of a trial. One way in which activity might be smoothed over short periods is through recurrent connections allowing activity from a moment ago to influence activity now. Some prominent theories associate local recurrent activity with consciousness (Lamme 2006, Tononi et al., 2016), and other theories agree that there is *typically* recurrent activity in posterior regions during conscious perceptual experience, generally only denying that such activity is sufficient for consciousness, and sometimes that it is necessary in *every* case (Lau, 2022; Seth & Bayne, 2022).

The activity of individual content-selective neurons does vary somewhat on short timescales despite recurrent connections. So this view would need to commit to the idea that the "activity" appealed to in the relevant SDT models does not simply consist in any individual neuron's activity. There are independent reasons to think this: for example, for clearly visible, foveated, attended stimuli, which everyone agrees are consciously perceived for several seconds, individual neurons' activity coding for that stimulus typically falls away well before the end of conscious experiences (usually only 200ms after stimulus onset) — suggesting more complex distributed *patterns* of activity in higher visual areas, which are more stable over time, are better candidates for the relevant kind of activity (Broday-Dvir, Norman, Harel, Mehta, & Malach, 2023; Vishne, Gerber, Knight, & Deouell, 2022). Unfortunately, even if we are willing to infer from these results concerning clearly visible, foveated, attended stimuli, to the kinds of peripheral experiences at issue for Inflationary Richness, we would not find fully stable activity: the distributed patterns studied by Broday-Dvir, Norman, Harel, Mehta, & Malach, 2023 and Vishne, Gerber, Knight, & Deouell, 2022 do still fluctuate *somewhat* moment to moment.

At present, it would be an unacceptable empirical bet for advocates of Inflationary Richness to simply assume that activity, whatever it turns out to be, is smooth over the relevant time period. A better approach would appeal instead to models in which experience is determined more dynamically, i.e. taking multiple periods' activity levels into account. The simplest way of doing this would be to have consciousness determined by whether the *average* activity over several time periods crosses the criterion (Fig. 4).

Averaged activity will be smoother, and hence using it to determine consciousness will result in fewer flickers, as individual spikes in noise will generally not result in conscious experience. Furthermore, from an engineering perspective it makes sense to make signal detection decisions based on averaging over multiple samples, thereby avoiding any individual decision being hostage to noise. Indeed, this sort of model effectively constitutes a simple version of an Evidence Accumulation (specifically Leaky Integration) model, whose application to conscious experience has been explored by several recent discussions (Dehaene, 2011; Denison, Block, & Samaha, 2022; Kang, Petzschner, Wolpert, & Shadlen, 2017; Pereira, Perrin, & Faivre, 2022; Shadlen & Kiani, 2011).[8]

However, this model also predicts that conscious experience will be slower to change in response to genuine changes in the stimulus: because previous activity will continue to influence the average for several time periods, averaged activity will only rise above criterion a few moments after stimulus onset is reflected in unaveraged activity, and fall to below criterion a few moments after stimulus offset is reflected in unaveraged activity.[9] Furthermore, *genuine* flickers in the *stimulus* will be unlikely to be reflected in experience at all, assuming they result in only brief changes in activity that barely shift the average: essentially, averaging across time means losing out on temporal sensitivity.

A further worry is that averaged activity may result in fewer flickers partly because it flickers less, but partly because it is less frequently above the criterion (as reflected in Fig. 4). So averaging may also result in less richness (fewer experienced items across the visual field) at any given time. This might be combated by lowering the criterion even further, but this in turn will only guarantee richness via resulting in more flickers, albeit slightly longer-lasting ones than those that result from unaveraged activity.

Simply averaging activity across time is not the only way to allow activity to determine experience more dynamically. For example, we could stipulate that activity must cross the criterion for a certain amount of time to produce consciousness. Or we could make any number of other complications to our assumptions. But all options that smooth experience over time by effectively smoothing the signal, rather than by improving signal-to-noise ratios, face the same tradeoff: reducing the amount of flickering that results from noise tends to reduce the temporal sensitivity and/or richness of experience.

This tradeoff can be illustrated by the "leakage" parameter in some Evidence Accumulation models. Despite my introducing such models as ways of *reducing* flickering, both Denison et al., (2022,p. 261) and Pereira et al. (2022) specifically complain that such models *predict* flickering, for some levels of leakage, as at least *sometimes* the decision variable (the weighted average/cumulative level of recent activity levels) hovers around threshold, being pushed above and below it through fluctuations in noise. Leakage is the rate at which previously accumulated evidence is forgotten — approximately, the inverse of the number of periods being averaged over. High leakage will make the decision more dependent on very recent activity levels, i.e. vulnerable to fluctuations in noise. Lowering leakage will diminish this problem, but it will do so by making experience more temporally coarse-grained, missing out on very brief changes in the stimulus, and slow to respond to sudden changes. There is no value of leakage which produces high temporal resolution whilst also ruling out flickering: there is always a tradeoff between the two. There is evidence that tradeoffs of this kind are ameliorated somewhat by the brain flexibly adapting parameters like leakage based on the typical durations of the stimuli important to a task (e.g. Ossmy et al., 2013), but even if the brain makes these tradeoffs *optimally*, they need to be made, implying phenomenal flicker could only be reduced so much.

Some kinds of evidence accumulation model can avoid these issues in a different way. Suppose that experience arises when accumulated evidence crosses a decision threshold, but that once the threshold is crossed, the experience subsequently plays out in a predetermined way, without being affected by further fluctuations in activity. Flickers would be ruled out by default: even if activity dips above and below threshold, once started, the experience would play out. A related phenomenon has been found empirically. Kang et al. (2017) had subjects indicate the overall direction of random dot motion and then estimate on a clock the time at which they felt they made their decision. Evidence accumulation models positing a decision at the indicated times outperformed rival models in

[8] Such models are typically presented as based around whether *total* or *accumulated* activity crosses a threshold, rather than average activity. But these approaches are equivalent, in all respects that are important to the issue at hand, (especially where the evidence accumulation models in question are *leaky*, i.e. only recent activity affects the decision variable). Both use a weighted sum of activity over recent periods, which smooths out the effects of momentary variations in noise, and both have similar strengths and weaknesses. But thinking in terms of average activity allows for direct comparison to a model where only activity at one moment counts, as in Fig. 4.

[9] This lag could be ameliorated somewhat if predictions about the stimulus in a few moments' time are incorporated into activity, rather than activity only resulting from feedforward processing of actual stimulation — but only when such predictions are accurate.

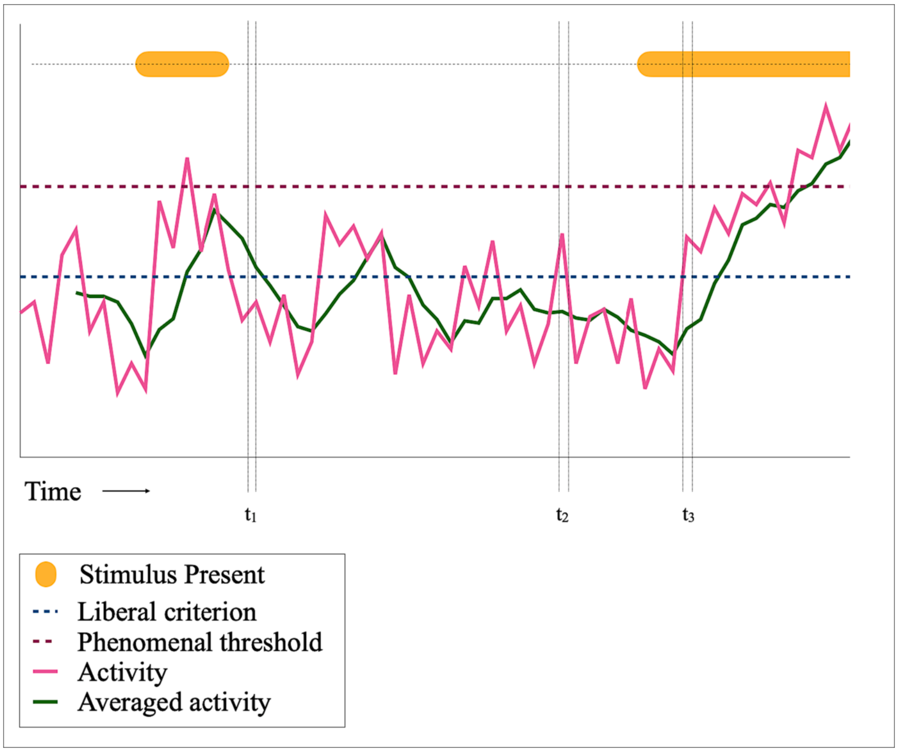

Fig. 4. The Time Course of Averaged Activity. The green line is the mean of activity levels (pink line) over the current and previous 4 periods. Averaged activity fluctuates less moment to moment in response to noise, and hence will avoid flickers (e.g. t_2), but also is slower to change in response to genuine stimulus onsets (e.g. t_3) and offsets (e.g. t_1). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

predicting reaction times. Crucially for our purposes, those models also indicated that subjects became less responsive to new evidence after making their decision. If consciousness tracks the decision (in the spirit of Inflationary Richness),[10] flickering of experience would be correspondingly reduced.

This view still faces a version of the tradeoff we have been discussing. It seems to have trouble accounting for cases where our consciousness does rapidly change, especially for cases where items *do* flicker in and out of consciousness. The idea that experience simply plays out in a predetermined way as soon as some threshold is crossed only looks plausible in constrained cases where it is clear how experience *should* play out if there is a stimulus, i.e. in extremely constrained scenarios like a well-defined, simple task of estimating the direction of random-dot motion by pressing a button once. Often, we do not have a well-defined task at all, and consciousness unfolds anyway, revealing changes in the external world as they unfold, however they unfold, from a vast number of possible alternatives. Even when we passively look at a random dot motion screen, we are capable of seeing overall motion smoothly or abruptly changing speed or direction, and there may be no subjectively salient decision time of the kind Kang et al. (2017) asked subjects to report. Kang *et al*. are sensitive to the fact that experience does not simply consist in a single representation piercing consciousness while otherwise remaining static, and suggest that there may be many provisional "mini-decisions" to report about the target property (overall motion direction) or different features of the display as the experiment unfolds (Kang et al., 2017, p. 2293). However, making this move in our context raises a dilemma: if such mini-decisions show up in phenomenology, are frequent enough to explain our phenomenology changing over time, arise as a result of the accumulation of noisy evidence, and do not damp down future evidence from affecting further mini-decisions, then flickering reappears: if the incoming stream of evidence is noisy, we will get fluctuation between mini-decisions which contradict one another. If, on the other hand, one or more of these assumptions is violated, it is unclear if we can explain why consciousness does change as rapidly as it does in the face of genuine change in the world.

It may turn out to be possible to avoid such tradeoffs through an even more complicated model. Options of this kind deserve to be fully worked out and tested. But they will face at least the following worry: as they demand more and more complexity of processing, the theory starts looking like it might run into the very bottlenecks that motivate Traditional Sparseness, instead of genuinely enabling spatial richness.

But the tradeoff, and indeed the whole apparent problem posed by flickers for Inflationary Richness, might be avoided via a very

[10] This would also be in the spirit of Kang et al. (2017), who explicitly suggest that the decision responsible for the measured report is also responsible for the item "piercing conscious awareness".

different strategy, which the next section explores. Rather than trying to reconcile experience of genuine temporally fine-grained changes with stable phenomenology, it may be possible to embrace the view that phenomenology *is* unstable, at least in the periphery.

7.3. Embracing instability

To begin, we should recognise that it is at least possible that instability in phenomenology might not be noticed — or even noticeable — as such. Change in experience is not necessarily experienced as change. James (1890, p. 629) famously claimed that "a succession of feelings, in and of itself, is not a feeling of succession". Likewise, the flickering of experience — a succession of absence of experience, followed by experience, followed by absence — need not result in an experience of anything as flickering.

What's more, while Inflationary Richness predicts flickering, it also predicts that trying to detect that flickering through carefully *attending* to the periphery will systematically underestimate the extent of that flickering. Recall that, unlike some forms of Traditional Richness and Degrees of Consciousness, which predict forms of flickering across the visual field, Inflationary Richness only predicts flickering for regions where there is inflation. And inflation occurs especially where spatial attention is lower. The very act of attending to the periphery boosts the signal-to-noise ratio, leads criteria to become more conservative — reduces inflation — and thereby reduces flickering.

This effect of careful introspection destroying the very details of experience it is supposed to reveal is closely related to a phenomenon many richness theorists are already committed to. One of the main arguments for the existence of rich phenomenal consciousness overflowing access consciousness (Block, 2007) revolves around findings by the likes of Sligte, Scholte, & Lamme, 2008 and Sperling (1960) which suggest (though see Gross & Flombaum, 2017 for an alternative interpretation) that we have a kind of short-term memory carrying a wealth of information about briefly-presented stimuli. For example, we seem to have access to the identities of many individual letters on a briefly-presented grid. Theorists like Block take this wealth of information to reflect a wealth of phenomenal experience. But they also think that, while each part of that information can be accessed if cued in the right way, it is highly fragile, in the sense that cueing one part of the information leaves us unable to access the others, and that this fragility makes these contents phenomenally but not access conscious. Block's (2007) interpretation of these experiments is controversial (Kouider et al., 2010; Phillips, 2011; Stazicker, 2011), but it does show that the possibility of flickers likewise going unnoticed despite entering into consciousness is not an impossible pill for a pro-richness theorist to swallow.

One might think that even so, we should notice flickers: do we not typically notice when there is a flickering light in our periphery?

Things are not so straightforward. Firstly, we do *not* always notice changes in experience. Change blindness, where subjects miss sometimes large changes in stimuli, is a well-documented phenomenon, and is more likely in the periphery and in unattended areas of the visual field (Henderson & Hollingworth, 1999; O'Regan et al., 1999, 2000; Pringle et al., 2001; Rensink et al., 1997). These results are sometimes presented as indicating that perception is not as rich as we typically suppose (e.g. Rensink et al., 1997), and *sometimes* such "blindness" may be because experience does not represent the items which change. But richness theorists will insist that in at least some of these cases, it does: it is just that as our experience changes, this is not experienced as change, and goes unnoticed, for example because we do not remember what our experience was like moments ago robustly enough to compare it to what it is like now. Such interpretations are enough to make unnoticed flickers a similarly theoretically viable option.

Secondly, we should ask what happens when we do consciously notice flickering lights presented to our periphery: these flickers draw exogenous attention to the relevant region of the visual field, plausibly only resulting in access conscious experience of flickering when higher attention processing confirms that there really is flickering out there. But as we already saw, in our case attention will do the opposite: it will reveal that any flickers in signal above a liberal criterion were simply noise.

Thirdly, while I have been talking about 'flickering' in phenomenology, it is worth being clear about what exactly this amounts to. It is tempting to suppose that flickeringly seeing a dim light would be like seeing a flickering dim light. But it would not be. The latter would consist in seeing a succession of darkness, light, and darkness again. But flickering experience would not involve seeing darkness, absence, or black, at any point in the sequence: except when seeing the stimulus, there would be no experience at all. The overall effect would merely be fluctuation between experience of a stimulus as present, and an absence of experience. This would not be nearly as striking as experience of absence, and would in no way compromise one's sense that one's environment is stable.

In fact, this kind of instability can be assimilated to a much broader instability in our phenomenology. Presumably, small details of our perceptual experience change from moment to moment given subtle changes in lighting conditions, slight changes in the precise regions of objects' surfaces we are currently seeing and the angle from which we are seeing them, small changes to our eyes and visual systems, faint afterimages fading as we look away from bright surfaces, and on and on. Sometimes we can become aware of instability in our phenomenology, as when we attend to the warmth of crackles and pops in a vinyl record, or actively try to paint glistening, rippling water. But such instability typically goes unnoticed, and noticing it can require special training or explicit instruction. The stable elements of the environment, including the shape and colour of the boat that remain constant as it bobs, and even the overall dynamic texture of the water as rippling, are nearly always more salient than transient features like the *precise* shape of the water at a specific time. They are also nearly always far more important to the many uses experience has for us. This is why we have constancy mechanisms dedicated to extracting many of these stable elements of our surroundings.

It is not unreasonable to think that there may be many ways in which experience changes from moment to moment — including noise giving rise to extremely brief peripheral flickers — which, given our normal purposes, would be pure distraction nearly all of the time were we to notice them. Neither is it unreasonable to suspect that such changes are, for this reason, filtered out so effectively that we *cannot* come to notice them, even in the way that a trained painter, through difficult training, can come to appreciate subtle changes from light dancing off the waves — although the view would also be compatible with certain kinds of training in introspection, such as in meditation, coming to reveal some such flickers.

Ultimately, the important question for using introspection to decide between views of richness is how any introspective pull towards thinking that phenomenology is flicker-free compares to the introspective pull towards richness. The latter — the apparent feature of experience which Inflationary Richness claims to validate — seems to me stronger than the former. This finds support from a final point: Inflationary Richness is far from alone in making surprising predictions about the time-course of experience. For example, even a standard sparse Global Workspace theory appears to predict "temporal islets of consciousness separated by brief periods of unconsciousness" in the *attended fovea* (Naccache, 2018, p. 7).

Given all this, it should be plausible that the precise time-course of experience might not be readily introspectable to us, and that this might be an area where a surprising prediction by a theory of consciousness might be more likely to be a surprising discovery than an error. As such, the main worry for Inflationary Richness need not be a worry at all: it might well be that, in the periphery, there are brief noise-driven flashes of phenomenal experience which mostly go unnoticed and in many cases are difficult or even impossible to notice as such.

8. Ubiquitous low-degree bizarre phenomenology

While accounting for the time-course of experience is the main challenge for Inflationary Richness, a major independent problem for Degrees of Consciousness is that it seems to predict ubiquitous bizarre conscious experience. We always have low but non-zero levels of activity for many possible contents of perception, and if degree of consciousness is simply dictated by level of activity, this implies that we have some degree of consciousness for a large range of contradictory contents at any given time: a group of circles, a skull shape, a solid shiny blue texture etc. This prediction seems phenomenologically implausible. One of the major advantages of Inflationary Richness over Degrees of Consciousness is its avoiding this consequence (as most such activity will not cross its phenomenal threshold, even when it is more liberal due to inflation). It is therefore worth considering possible replies the defender of Degrees of Consciousness might make to this worry.

They can point out that the problem of the brain encoding many possibilities which do not seem to show up in experience is not unique to their view: a parallel problem seems to be generated by evidence for probabilistic representation in perception. Neurons do not simply respond to unique contents, but have tuning curves: they often respond maximally to one value of a variable, but respond somewhat to a wide range of values. Any model where the content of experience is determined straightforwardly by such activity would therefore seem to imply that, when looking at a vertically oriented line in good lighting conditions, we should experience not just the vertical orientation, but also a very wide range of other orientations, associated with much lower probabilities — including even horizontal orientations. A similar prediction can be generated at the purely computational level, after abstracting away from neural details of implementation like tuning curves: Bayesian models of perception seem to have very similar implications. Yet it is extremely implausible that our experience of vertical lines in good conditions is like this. Such experiences seem to have, as Block puts it (2018: 2), "no hint" of orientations far from the vertical: "conscious perception does not normally reflect the probabilistic hypotheses other than the dominant one". The Degrees of Consciousness theorist might reasonably hope that the solution to *this* problem, which does not depend on endorsing Degrees of Consciousness, can be transposed to provide a solution to their specific problem of ubiquitous bizarre contents.

Unfortunately, this hope may go unfulfilled. It is not clear that the most promising solutions to the problem of probabilistic perception can be successfully transposed to the problem of ubiquitous bizarre contents.

Consider the discussion of the problem of probabilistic perception in Block (2018). Block convincingly rejects multiple potential solutions, including: (i) feelings of perceptual confidence (which fail to explain why we do not have unconfident but conscious perceptual awareness of low-probability possibilities); (ii) claiming experience is indeterminate between different possibilities (which fails to explain why our experience seems to only include the high probability option); and (iii) claiming that our experience reflects sampling from probability distributions for perceptual decision-making (which, to simplify considerably, Block argues cannot explain why we experience determinate contents in the absence of any specific task). Block suggests a solution with two main components. Firstly, Bayesian models should be interpreted merely instrumentally, not as specifying computations literally performed by visual systems (Rahnev and Block, in Rahnev, Block, Denison, & Jehee, 2021, make related suggestions on which perceptual representations really specify just the most likely estimate and an associated uncertainty value). Secondly, sensory representations only result in conscious experience when they "win" over competing representations, with losing representations getting suppressed (one natural interpretation of what "winning" involves is *getting selected for global broadcast*, but this is not the only possibility, with at least one other alternative more amenable to Block's other views: local ignition (Fisch, et al., 2009; Noy, et al., 2015)).

Suppose this account successfully explains why we perceive vertical orientations rather than a broad range of possible but improbable orientations. To solve *their* problem in like manner, the advocate of Degrees of Consciousness would need to specify which representations are in competition with which other representations. For a particular line on the screen we can find neat Gaussian distributions over possible orientations, or over possible colours, and the competing hypotheses seem straightforward: the line could have any orientation between 0° and 180°, or any hue on the colour wheel. But what exactly is a very weak encoding of a skull shape in the bottom half of the visual field competing with? Given that more than one feature can be consciously represented at once, it cannot be that every representation is simply competing against every other representation. Again, a much better-specified theory would be called for to adopt a reply like this.

One potentially independent way of avoiding the problem of ubiquitous low-degree phenomenology is to abandon *pure* Degrees of Consciousness in favour of Denison et al.'s (2022) "Option 2": consciousness comes in degrees corresponding to SDT activity levels, but only occurs when activity is above some phenomenal threshold (a similar view is discussed in Pereira et al., 2022, p. 457). This would avoid the problem at hand; but such a view would inherit several disadvantages of both Phenomenal Threshold and of Degrees of

Consciousness. For example, it would need to commit to an account of what degrees of consciousness actually are in this context, like Degrees of Consciousness. Yet it would also, like Phenomenal Threshold, need to provide a motivation for any commitments concerning how liberal or conservative threshold is and how it behaves during inflation. It could not avoid making such commitments: they are crucial to generating any predictions about richness. Yet making such commitments will be tricky, not only because it is hard to empirically determine the phenomenal threshold independently of the report criterion, but also because no commitment here seems to allow such a position to maintain its distinctive advantages over all the other views. If the threshold is uniformly high throughout the visual field, then the view becomes extremely similar to a standard conservative Phenomenal Threshold view: it is no longer the case that there is rich phenomenology throughout the visual field. If instead the threshold is uniformly low, the view will be very close to Degrees of Consciousness without any threshold, and may not solve the problem of ubiquitous contents after all. Yet if the threshold is flexible and affected by inflation, then the view effectively just becomes a version of Inflationary Richness.

The problem of ubiquitous bizarre contents, then, remains a pressing one for Degrees of Consciousness, and there does not seem to be a solution to it that is as compelling as the solutions Inflationary Richness can offer to the problem of flickering.

9. Conclusion

There are multiple possible interpretations of subjective inflation, each consistent with the empirical data, and each with distinctive costs. Degrees of Consciousness predicts low-degree conscious experience of many features in our periphery, but implies that the degree of consciousness of different stimuli is constantly fluctuating *throughout the visual field*, and seems to implausibly imply our having *too many* low-degree experiences. Phenomenal Threshold predicts that experience is sparse in the periphery, claiming that pro-richness intuitions result from misinterpreting a tendency to *report* stimuli which we do not really *experience*, but needs to explain why we mistake these rather different kinds of experience. Inflationary Richness allows for rich (albeit inaccurate) phenomenology throughout the visual field, and naturally coheres with the Perceptual Reality Monitoring theory of consciousness, but implies flickering phenomenology in the periphery. This is a real cost: accepting it requires emphasising a certain kind of fallibility in our introspection. But given that, as I have shown, this implication is neither shared by all versions of Inflationary Richness, nor quite as implausible as it might appear, the balance of evidence currently favours Inflationary Richness. The next most plausible option would solve some of Degrees of Consciousness's issues by effectively combining it with Inflationary Richness: as in Dension *et al.*'s "Option 2", consciousness of specific features comes in degrees (perhaps specified in terms of a certain kind of felt confidence and/or precision), but only occurs when activity is above some threshold which (with Inflationary Richness) is more liberal in the periphery due to inflation — predicting rich, unstable, low-degree phenomenology throughout the visual field, but far fewer bizarre representations than standard Degrees of Consciousness, especially in the fovea.

We should be cautious before decisively endorsing either view, however, especially insofar as each predicts instability of a kind that involves accepting sharp limits to introspection's capacity for revealing the temporal structure of phenomenology. Progress on numerous relevant open empirical and theoretical questions could tip the balance differently in the future. Does inflation always obtain, for all unattended or peripheral stimuli, or is it limited to certain kinds of stimuli? Does it play out differently for more complex stimuli, or stimuli that come with certain kinds of well-defined dimensions? If performance in the periphery is less affected by attention for faces than for other stimuli like colours and letters (Reddy et al., 2004; 2006), this is one reason to suspect that inflation will similarly play a role for only some stimuli (indeed, the metacognitive aspects of inflation seem to not apply to some of these stimuli — Matthews et al., 2018). One avenue for further investigation is this: if inflation is due to a Perceptual Reality Monitoring system, a view which supports Inflationary Richness, we might expect to find inflation compensating for differences in sensitivity that result from other factors besides spatial attention and eccentricity. For example, given that performance differs not just with eccentricity, but also with polar angle (Himmelberg, Winawer, & Carrasco, 2023), one might expect differences in criteria for different regions of the periphery (e.g. vertical vs. horizontal) too. It would be particularly informative to test whether inflation not only occurs differentially for these different regions, but does so in a way which tracks the changes across development in these performance differences (Carrasco et al., 2022). Furthermore, given the flexibility of the SDT framework, it may be unrealistic to expect a uniform mapping to consciousness across all contexts: some shifts in criteria may reflect shifts in factors having only to do with reporting, while others reflect differences more directly relevant to consciousness (Peters & Lau, 2015; Peters, Ro, & Lau, 2016). Teasing out the precise reasons for a particular shift can help.

Perhaps the most important lines of future research lie in fleshing out more sophisticated models of the time course of experience throughout the visual field and finding ways of deciding between them. Both the sorts of mechanisms for reducing flicker discussed in 6.2, and the possibility of unnoticed flicker discussed in 6.3, might be at play (and could even interact with one another), and studying them in greater detail would shed a great deal of light on experience throughout the visual field. It would be particularly useful to find ways of isolating how much the "activity" invoked by SDT actually fluctuates moment to moment, and how these fluctuations get converted into changes in experience. Various papers have tried to pin criteria to different brain regions (e.g. Crapse et al., 2018, Luo & Maunsell, 2018), and huge numbers have measured fluctuating activity leading to one-off decisions through evidence accumulation (reviewed in Pereira et al., 2022). But given that we have emphasised that multiple factors can determine behaviourally expressed criteria in different situations (including incentives, reality monitoring etc.), determining the time-course of experience in open-ended situations without a defined task might require much more sophisticated models of the dynamics of experience, which separate out different reasons for changes in criteria (e.g. reality monitoring, changes in incentives etc.), show exactly how tradeoffs between flickering and temporal sensitivity are dealt with, and incorporate any differences in how these systems deal with different kinds of stimuli. Such a comprehensive picture of experience across the visual field, across time, will require integrating different measures of consciousness, including neural measures and non-SDT measures like the Perceptual Awareness scale, in combination with sophisticated modelling of the computations involved — it will require a true "mesh" between different disciplines' insights into visual experience (Block, 2007).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgements

I would especially like to thank two reviewers: Marius Usher, and another who remained anonymous . I would also like to thank Ian Phillips and Caroline Myers, for detailed comments on earlier versions of this work, and Steven Gross, Andrew Lee, Matthias Michel, Brian Odegaard, and audiences at Johns Hopkins University, at the Association for the Scientific Study of Consciousness 2022 meeting in Amsterdam, and at the Society for Philosophy and Psychology/European Society for Philosophy and Psychology 2022 joint meeting in Milan, for helpful discussions.

Funding

This research is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, Grant Number 851145.

References

Abid, G. (2019). Deflating inflation: The connection (or lack thereof) between decisional and metacognitive processes and visual phenomenology. *Neuroscience of Consciousness, 2019*(1), niz015. https://doi.org/10.1093/nc/niz015

Amir, Y. Z., Assaf, Y., Yovel, Y., & Mudrik, L. (2023). Experiencing without knowing? Empirical evidence for phenomenal consciousness without access. *Cognition, 238*, Article 105529. https://doi.org/10.1016/j.cognition.2023.105529

Anton-Erxleben, K., & Carrasco, M. (2013). Attentional enhancement of spatial resolution: Linking behavioural and neurophysiological evidence. *Nature reviews. Neuroscience, 14*(3), 188–200. https://doi.org/10.1038/nrn3443

Barron, A. B., & Klein, C. (2016). What insects can tell us about the origins of consciousness. *Proceedings of the National Academy of Sciences, 113*(18), 4900–4908. https://doi.org/10.1073/pnas.1520084113

Bayne, T., Hohwy, J., & Owen, A. M. (2016). Are There Levels of Consciousness? *Trends in Cognitive Sciences, 20*(6), 405–413. https://doi.org/10.1016/j.tics.2016.03.009

Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *The Behavioral and Brain Sciences*, *30*(5–6), 481–499; discussion 499-548. doi: 10.1017/S0140525X07002786.

Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences, 15*(12), 567–575. https://doi.org/10.1016/j.tics.2011.11.001

Block, N. (2015). The puzzle of perceptual precision. In J. Windt, & T. Metzinger (Eds.), *Open MIND* (pp. 1–52). Frankfurt am Main: MIND Group.

Block, N. (2018). If perception is probabilistic, why does it not seem probabilistic? *Philosophical transactions of the Royal Society of London. Series B, Biological sciences, 373*(1755), 20170341. https://doi.org/10.1098/rstb.2017.0341

Broday-Dvir, R., Norman, Y., Harel, M., Mehta, A. D., & Malach, R. (2023). Perceptual stability reflected in neuronal pattern similarities in human visual cortex. *Cell reports*, *42*(6), 112614. Advance online publication. doi: 10.1016/j.celrep.2023.112614.

Bronfman, Z., Jacobson, H., & Usher, M. (2018). Impoverished or rich consciousness outside attentional focus: Recent data tip the balance for Overflow. *Mind & Language, 34*, 423–444. https://doi.org/10.1111/mila.12217

Carrasco, M., Roberts, M., Myers, C., & Shukla, L. (2022). Visual field asymmetries vary between children and adults. *Current Biology, 32*(11), R509–R510. https://doi.org/10.1016/j.cub.2022.04.052

Cohen, M. A., Dennett, D. C., & Kanwisher, N. (2016). What is the Bandwidth of Perceptual Experience? *Trends in Cognitive Sciences, 20*(5), 324–335. https://doi.org/10.1016/j.tics.2016.03.006

Cohen, M. A., Ostrand, C., Frontero, N., & Pham, P.-N. (2021). Characterizing a snapshot of perceptual experience. *Journal of Experimental Psychology: General, 150*(9), 1695–1709. https://doi.org/10.1037/xge0000864

Crapse, T. B., Lau, H., & Basso, M. A. (2018). A Role for the Superior Colliculus in Decision Criteria. *Neuron, 97*(1), 181–194.e6. https://doi.org/10.1016/j.neuron.2017.12.006

Dehaene, S. (2011). In *Conscious and Nonconscious Processes: Distinct Forms of Evidence Accumulation?* (pp. 141–168). Springer. https://doi.org/10.1007/978-3-0346-0428-4_7.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition, 79*(1–2), 1–37. https://doi.org/10.1016/s0010-0277(00)00123-2

Denison, R. N., Adler, W. T., Carrasco, M., & Ma, W. J. (2018). Humans incorporate attention-dependent uncertainty into perceptual decisions and confidence. *Proceedings of the National Academy of Sciences of the United States of America, 115*(43), 11090–11095. https://doi.org/10.1073/pnas.1717720115

Denison, R. N., Block, N., & Samaha, J. (2022). What Do Models of Visual Perception Tell Us about Visual Phenomenology? In F. De Brigard, & W. Sinnott-Armstrong (Eds.), *Neuroscience and Philosophy* (pp. 241–284). MIT Press.

Fink, S. B. (2015). Phenomenal precision and some pitfalls: A commentary on Ned Block. In J. Windt, & T. Metzinger (Eds.), *Open mind: Philosophy and the mind sciences in the 21th century* (pp. 1–14). Cambridge, MA: MIT Press.

Fisch, L., Privman, E., Ramot, M., Harel, M., Nir, Y., Kipervasser, S., Andelman, F., Neufeld, M. Y., Kramer, U., Fried, I., & Malach, R. (2009). Neural "ignition": Enhanced activation linked to perceptual awareness in human ventral stream visual cortex. *Neuron, 64*(4), 562–574. https://doi.org/10.1016/j.neuron.2009.11.001

Gross, S., & Flombaum, J. (2017). Does Perceptual Consciousness Overflow Cognitive Access? The Challenge from Probabilistic, Hierarchical Processes. *Mind & Language, 32*(3), 358–391. https://doi.org/10.1111/mila.12144

Hansen, T., Pracejus, L., & Gegenfurtner, K. R. (2009). Color perception in the intermediate periphery of the visual field. *Journal of vision, 9*(4), 1–12. https://doi.org/10.1167/9.4.26

Hautus, M. J., Macmillan, N. A., & Creelman, C. D. (2021). *Detection Theory: A User's Guide*. Taylor & Francis.

Henderson, J. M., & Hollingworth, A. (1999). The Role of Fixation Position in Detecting Scene Changes Across Saccades. *Psychological Science, 10*(5). https://doi.org/10.1111/1467-9280.00183

Himmelberg, M. M., Winawer, J., & Carrasco, M. (2023). Polar angle asymmetries in visual perception and neural architecture. *Trends in Neurosciences, S0166-2236(23)00068-1.*. https://doi.org/10.1016/j.tins.2023.03.006

James, W. (1890). *The principles of psychology, Volume I*. Henry Holt and Co. https://doi.org/10.1037/10538-000

Jin, M., & Glickfeld, L. L. (2019). Contribution of sensory encoding to measured bias. *The Journal of Neuroscience, 39*(26), 5115–5127. https://doi.org/10.1523/JNEUROSCI.0076-19.2019

Kang, Y. H. R., Petzschner, F. H., Wolpert, D. M., & Shadlen, M. N. (2017). Piercing of Consciousness as a Threshold-Crossing Operation. *Current biology : CB, 27*(15), 2285–2295.e6. https://doi.org/10.1016/j.cub.2017.06.047

Knotts, J. D., Odegaard, B., Lau, H., & Rosenthal, D. (2019). Subjective inflation: Phenomenology's get-rich-quick scheme. *Current Opinion in Psychology, 29*, 49–55. https://doi.org/10.1016/j.copsyc.2018.11.006

Ko, Y., & Lau, H. A. (2012). A detection theoretic explanation of blindsight suggests a link between conscious perception and metacognition. *Philosophical Transactions of the Royal Society B: Biological Sciences, 367*, 1401–1411. https://doi.org/10.1098/rstb.2011.0380

Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: Two distinct brain processes. *Trends in Cognitive Sciences, 11*(1), 16–22. https://doi.org/10.1016/j.tics.2006.10.012

Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences, 14*(7), 301–307. https://doi.org/10.1016/j.tics.2010.04.006

Lamme, V. A. F. (2010). How neuroscience will change our view on consciousness. *Cognitive Neuroscience, 1*(3), 204–220. https://doi.org/10.1080/17588921003731586

Lamme, V. A. F. (2022). Behavioural and Neural Evidence for Conscious Sensation in Animals: An Inescapable Avenue Towards Biopsychism? *Journal of Consciousness Studies, 29*(3–4), 78–103. https://doi.org/10.53765/20512201.29.3.078

Lau, H. (2022). *In Consciousness We Trust*. Oxford University Press.

Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in cognitive sciences, 15*(8), 365–373. https://doi.org/10.1016/j.tics.2011.05.009

Lee, J. L., Denison, R., & Ma, W. J. (2023). Challenging the fixed-criterion model of perceptual decision-making. *Neuroscience of Consciousness, 2023*(1). https://doi.org/10.1093/nc/niad010

Li, F. F., VanRullen, R., Koch, C., & Perona, P. (2002). Rapid natural scene categorization in the near absence of attention. *Proceedings of the National Academy of Sciences, 99*(14), 9596–9601. https://doi.org/10.1073/pnas.0922775

Li, M. K., Lau, H., & Odegaard, B. (2018). An investigation of detection biases in the unattended periphery during simulated driving. *Attention, Perception & Psychophysics, 80*(6), 1325–1332. https://doi.org/10.3758/s13414-018-1554-3

Luo, T. Z., & Maunsell, J. H. R. (2018). Attentional Changes in Either Criterion or Sensitivity Are Associated with Robust Modulations in Lateral Prefrontal Cortex. *Neuron, 97*(6), 1382–1393.e7. https://doi.org/10.1016/j.neuron.2018.02.007

Matthews, J., Schröder, P., Kaunitz, L., van Boxtel, J. J. A., & Tsuchiya, N. (2018). Conscious access in the near absence of attention: Critical extensions on the dual-task paradigm. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 373*(1755), 20170352. https://doi.org/10.1098/rstb.2017.0352

Merker, B. (2007). Consciousness without a cerebral cortex: A challenge for neuroscience and medicine. *The Behavioral and Brain Sciences, 30*(1), 63–81. https://doi.org/10.1017/S0140525X07000891. discussion 81–134.

Michel, M., & Lau, H. (2021). Is Blindsight Possible Under Signal Detection Theory? Comment on Phillips (2021). *Psychological Review, 128*(3), 585–591. https://doi.org/10.1037/rev0000266

Morales, J. (2021). Introspection Is Signal Detection. *The British Journal for the Philosophy of Science*. https://doi.org/10.1086/715184

Naccache, L. (2018). Why and how access consciousness can account for phenomenal consciousness. *Philosophical Transactions of the Royal Society B: Biological Sciences, 373*(1755), 20170357. https://doi.org/10.1098/rstb.2017.0357

Noy, N., Bickel, S., Zion-Golumbic, E., Harel, M., Golan, T., Davidesco, I., Schevon, C. A., McKhann, G. M., Goodman, R. R., Schroeder, C. E., Mehta, A. D., & Malach, R. (2015). Ignition's glow: Ultra-fast spread of global cortical activity accompanying local "ignitions" in visual cortex during conscious visual perception. *Consciousness and cognition, 35*, 206–224. https://doi.org/10.1016/j.concog.2015.03.006

Odegaard, B., Chang, M. Y., Lau, H., & Cheung, S.-H. (2018). Inflation versus filling-in: Why we feel we see more than we actually do in peripheral vision. *Philosophical Transactions of the Royal Society B: Biological Sciences, 373*(1755), 20170345. https://doi.org/10.1098/rstb.2017.0345

O'Regan, J. K., Deubel, H., Clark, J. J., & Rensink, R. A. (2000). Picture changes during blinks: Looking without seeing and seeing without looking. *Visual Cognition, 7*, 191–211. https://doi.org/10.1080/135062800394766

O'Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change-blindness as a result of 'mudsplashes'. *Nature, 398*(6722). https://doi.org/10.1038/17953

Ossmy, O., Moran, R., Pfeffer, T., Tsetsos, K., Usher, M., & Donner, T. H. (2013). The Timescale of Perceptual Evidence Integration Can Be Adapted to the Environment. *Current Biology, 23*(11), 981–986. https://doi.org/10.1016/j.cub.2013.04.039

Overgaard, M., & Sandberg, K. (2021). The Perceptual Awareness Scale—recent controversies and debates. *Neuroscience of Consciousness, 2021(1)*. https://doi.org/10.1093/nc/niab044

Pereira, M., Perrin, D., & Faivre, N. (2022). A leaky evidence accumulation process for perceptual experience. *Trends in Cognitive Sciences, 26*(6), 451–461. https://doi.org/10.1016/j.tics.2022.03.003

Peters, M. A., & Lau, H. (2015). Human observers have optimal introspective access to perceptual processes even for visually masked stimuli. *eLife, 4*, Article e09651. https://doi.org/10.7554/eLife.09651

Peters, M. A., Ro, T., & Lau, H. (2016). Who's afraid of response bias? *Neuroscience of consciousness, 2016*(1). https://doi.org/10.1093/nc/niw001

Phillips, I. B. (2011). Perception and Iconic Memory: What Sperling Doesn't Show. *Mind and Language, 26*(4), 381–411. https://doi.org/10.1111/j.1468-0017.2011.01422.x

Phillips, I. B. (2021). Blindsight is qualitatively degraded conscious vision. *Psychological Review, 128*(3), 558–584. https://doi.org/10.1037/rev0000254

Pringle, H. L., Irwin, D. E., Kramer, A. F., & Atchley, P. (2001). The role of attentional breadth in perceptual change detection. *Psychonomic Bulletin & Review, 8*(1), 89–95. https://doi.org/10.3758/BF03196143

Rahnev, D., Maniscalco, B., Graves, T., Huang, E., de Lange, F. P., & Lau, H. (2011). Attention induces conservative subjective biases in visual perception. *Nature Neuroscience, 14*(12), 1513–1515. https://doi.org/10.1038/nn.2948

Rahnev, D., Block, N., Denison, R. N., & Jehee, J. (2021, May 5). Is perception probabilistic? Clarifying the definitions. doi: 10.31234/osf.io/f8v5r.

Reddy, L., Wilken, P., & Koch, C. (2004). Face-gender discrimination is possible in the near-absence of attention. *Journal of Vision, 4*(2), 106–117. https://doi.org/10.1167/4.2.4

Reddy, L., Reddy, L., & Koch, C. (2006). Face identification in the near-absence of focal attention. *Vision Research, 46*(15), 2336–2343. https://doi.org/10.1016/j.visres.2006.01.020

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To See or not to See: The Need for Attention to Perceive Changes in Scenes. *Psychological Science, 8*(5). https://doi.org/10.1111/j.1467-9280.1997.tb00427.x

Rolfs, M., Dambacher, M., & Cavanagh, P. (2013). Visual adaptation of the perception of causality. *Current Biology, 23*(3), 250–254. https://doi.org/10.1016/j.cub.2012.12.017

Rosenholtz, R. (2016). Capabilities and Limitations of Peripheral Vision. *Annual Review of Vision Science, 2*(1), 437–457. https://doi.org/10.1146/annurev-vision-082114-035733

Sánchez-Fuenzalida, N., Van Gaal, S., Fleming, S. M., Haaf, J. M., & Fahrenfort, J. J. (2023). Predictions and rewards affect decision-making but not subjective experience. *Proceedings of the National Academy of Sciences, 120*(44). https://doi.org/10.1073/pnas.2220749120

Seth, A. K., & Bayne, T. (2022). Theories of consciousness. *Nature reviews. Neuroscience, 23*(7), 439–452. https://doi.org/10.1038/s41583-022-00587-4

Shadlen, M. N., & Kiani, R. (2011). Consciousness as a Decision to Engage. In S. Dehaene, & Y. Christen (Eds.), *Characterizing Consciousness: From Cognition to the Clinic?* (pp. 27–46). Springer. https://doi.org/10.1007/978-3-642-18015-6_2.

Sligte, I. G., Scholte, H. S., & Lamme, V. A. F. (2008). Are There Multiple Visual Short-Term Memory Stores? *PLOS ONE, 3*(2), e1699. https://doi.org/10.1371/journal.pone.0001699

Solovey, G., Graney, G. G., & Lau, H. (2015). A decisional account of subjective inflation of visual perception at the periphery. *Attention, Perception & Psychophysics, 77*(1), 258–271. https://doi.org/10.3758/s13414-014-0769-1

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied, 74*(11), 1–29. https://doi.org/10.1037/h0093759

Stazicker, J. (2011). Attention, visual consciousness and indeterminacy. *Mind & Language, 26*(2), 156–184. https://doi.org/10.1111/j.1468-0017.2011.01414.x

VanRullen, R. (2013). Visual attention: A rhythmic process? *Current Biology, 23*(24), R1110–R1112. https://doi.org/10.1016/j.cub.2013.11.006

Tononi, G., Boly, M., Massimini, M., & Koch, C. (2016). Integrated information theory: From consciousness to its physical substrate. *Nature Reviews Neuroscience, 17*(7), 450–461. https://doi.org/10.1038/nrn.2016.44

Van Rullen, R. (2018). Attention Cycles. *Neuron, 99*(4), 632–634. https://doi.org/10.1016/j.neuron.2018.08.006. PMID: 30138586.

Vishne, G., Gerber, E. M., Knight, R. T., & Deouell, L. Y. (2022). *Representation of sustained visual experience by time-invariant distributed neural patterns* (p. 2022.08.02.502469). bioRxiv. doi: 10.1101/2022.08.02.502469.