HOW MANY STRIPES ARE ON THE TIGER IN MY DREAMS?

On imagery, indeterminacy, and introspective ignorance

Sascha Benjamin Fink
Philosophy–Neurosciences–Cognition Program
Otto-von-Guericke-University Magdeburg
sfink@ovgu.de
NYU-Version Sept. 2019

Abstract of the project: This is a précis of some parts of an ongoing book project on phenomenal imagery, its indeterminacy and our limited access to our own experiences, which I hope to finish in a year. I specifically want to encourage your input on section 2.1, 2.2, 3.1, 3.2.3 and (in 3.5) W1, W4, and W5. (Other sections are included for illustration or to show the scope of the project. Parts starting with "XXX" can be skipped, but are included to illustrate some points.

I close my eyes and see a flock of birds. The vision lasts a second or perhaps less; I don't know how many birds I saw. Were they a definite or an indefinite number? This problem involves the question of the existence of God. If God exists, the number is definite, because how many birds I saw is known by God. If God does not exist, the number is indefinite, because nobody was able to take count. In this case, I saw fewer than ten birds (let's say) and more than one; but I did not see nine, eight, seven, six, five, four, three, or two birds. I saw a number between ten and one, but not nine, eight, seven, six, five, etc. That number, as a whole number, is inconceivable; ergo, God exists.

Jorge Luis Borges (1976) Argumentum Ornithologicum from the book *Dreamtigers*

1. Introduction

In contemporary empirical consciousness studies, a wide range of findings have emerged where the interpretation of the data is unclear: There are doubts concerning how an experience during a dream, in the visual periphery, or in an unattended area of the phenomenal "field" is like — or even: whether there is something there feeling like anything. Concerning such cases, three kinds of interpretation are commonly proposed: The first holds that phenomenal imagery need not be fully accessible to the introspecting subject; a second holds that, at least in some critical cases, there is no phenomenal imagery; a third holds that phenomenal imagery can be, in some sense, indeterminate. It is specific this last claim and in some sense of "indeterminate" that I take issue with here.

Each interpretation upholds some specific commonly held intuition, but it does so at the price of another. I illustrate this by constructing an inconsistent triad of statement which I will call the "Striped Tiger Problem" (sec. 2.1). The form generalises and arises for a wide range of issues in philosophy of mind and psychology (see sec. 2.3). It can be characterised for specific cases by questions like How many stripes were on the tiger in my dream last night? I illustrate how the three interpretations connect to this problem in sec. 2.2.

Which of the three interpretations should we prefer? I focus on the position which rejects phenomenal imagery being determinate. I distinguish between different kinds of indeterminacy in sec. 3. In the most commonly proposed version, the position is tightly associated with Representationism. According to Representationism, the phenomenal character of some experience is determined by the content of that experience, such that features of the content fully account for all phenomenal features. However, I deem this the weakest position vis-à-vis Striped Tiger Problems for several reasons. By abduction, we ought to prefer one of the other interpretations. However, I hold that none of the other two is a reasonably generalises. Instead, we must decide on a case-by-case-basis which interpretation accounts best for all the data at hand.

2. The Problem

2.1. How many stripes are on the tiger in my dreams? Sometimes, when we dream, we may find it difficult to report on some features of our dream imagery. For example, say that one night I dream the following: Through the underbrush of a pink jungle, a tiger

¹ I introduce the problem with the example of dream imagery for two reasons: First, because dreams are paradigmatic instances of phenomenal imagery (pace Malcolm, 1959; Dennett, 1976). Second, because they are largely decoupled from sensory input and behavioural output. They are therefore non-public: The dreamer can hide what she is dreaming of (and how she is dreaming of something) from others. (But see Horikawa et al. (2013) for a proposal on how one might make aspects of dreaming publicly available by technical means. However, this method only predicts dream reports, not dream imagery.) It is this non-publicity which leads to a methodological or epistemic asymmetry: While the dreamer may know about her dreams by a first-person method (e.g. introspection), everyone else has to rely on behavioural or physiological data or on her reports to infer (consciously or unconsciously, intuitively or guided by theory) the content and character of that dream. Dreams therefore allow us to focus on the phenomenal aspects of imagistic experiences without confusing them with their perceptual or behavioural aspects.

approached me. Calmly extending my hand, I stroked its flank in the middle of a glade. Then, I wake up. After waking, I can ask myself a question such as:

 (Q_t) How many stripes were on the tiger I just dreamt of?

Very likely, I will find it hard to answer. However, I would likely assent to the following two statements:

- (1_t) This night, it felt like seeing a striped tiger in my dreams.
- (2_t) I hesitate to commit to any specific number of stripes, or even whether the number was even or odd.

The first thesis, 1_t , amounts to a commitment to phenomenal imagery: I felt as if I was watching a tiger, not as if I was thinking of a tiger. I might write down what I saw in my dream diary, but the closest representation to what this experience was like would be drawing a picture rather than writing a report — Dalí is closer to what dreams are like than Sōseki. The second thesis, 2_t , is merely a description of my behaviour.

Even though I may feel hesitant to commit to any specific number of stripes, my hesitance is not uniform when it comes to rejecting specific numbers: I can confidently state that I know that my dream-tiger did not have exactly 3 stripes or 300. My credence distribution — to what degree I believe that my dream-tiger had n stripes for any n, if I were rational — may explain why: For any number of stripes, my credence is below 0.5, but for some it approaches or is straightforwardly 0 (see fig. 1 for an illustration). Knowing that p entails believing that p. If believing that p means having a credence of larger than 0.5 in p, then I do not know that my dream-tiger had n stripes, because I do not believe that my dream-tiger had n stripes for any n. For no whole number, my credence is larger than 0.5. But conversely, if my credence for not-p is 1 minus my credence in p, then there are some whole numbers where I am certain that my dream tiger did not have that many stripes, e.g. any number below 3 or above 300. Here, my credence in the dream-tiger having that many stripes is straightforwardly 0. Because I do not believe that my tiger had n stripes for any n, I do not know that it had n stripes. I seem to lack first-order knowledge about my own experiences.

My lack of knowledge does not only seem to apply to ignorance of propositional (I know that my dream tiger had n stripes) or interrogative (I know how many stripes my dream tiger had) knowledge, but ignorance concerning knowing-how as well: I do not only fail to know that my dream-tiger had n stripes (for any n), but at the same time I also do not know how to draw the tiger in my dreams — and, at least in my case, this is not so because I do not know how to draw tigers. I simply do not know whether I succeeded in drawing the tiger in my dreams adequately, whether I depicted it just like it was presented to me or simply drew an icon standing in for that tiger in my dreams, but not resembling it. Thus, my hesitance is one concerning reporting as well as a hesitancy concerning actions.

Maybe this is too quick: Even though I would not ascribe myself knowledge, my own hesitance is compatible with the possibility that I always give the right answer if someone forces me. But due to my privileged access to (and the privacy of) my dream imagery, nobody else could verify whether my forced answer actually is the right one. Even if I were always correct, I would not ascribe myself knowledge — and nobody else is in the




FIGURE 1. A plausible credence distribution a subject S might have for the number of stripes on the tiger in S's dreams: a normal distribution around 60 roughly \pm 20. ($\mu = 60, \sigma = 2\sqrt{10}$).

position to do so either. If there were a way to verify, and if you told me that I was right in guessing n stripes, I myself might call it a lucky guess. Strictly speaking, my hesitance — a type of behavioural evidence — is only evidence against higher-order knowledge: I do not know whether I know. To deny me first-order knowledge, we may rely on the KK-rule in an argument of the following form:

- (3_t) My hesitance to assent to any number of stripes (or even: whether that number is even or odd) counts as behavioural evidence for a lack of higher-order knowledge about the number of stripes of my dream tiger: I do not know whether I know that my dream-tiger had n stripes (for any n). So I do not ascribe myself knowledge of its number of stripes. I reject that I know that I know that it had n stripes.
- (4_t) **KK-rule:** If S knows that p, then S knows that S knows that p.
- (5_t) I do not know how many stripes my dream tiger had. (From 3_t and 4_t by modus tollens)

A problem arises if we combine 1_t – 5_t with two common assumption, one about imagistic representations and the other about first-person access to phenomenal imagery.

Concerning imagistic representation, they are commonly seen as determinate concerning their lower-order properties (see e.g. Kosslyn 1980). Why? Even if we do not know what is represented in an image or if it is represented correctly, the image *qua image* is a concrete particular because images represent by structural isomorphism with the object they represent: It must look like the thing it represents to some degree in order to represent

² We cannot exclude malign forms of luck (Pritchard, 2005: 146 & 175) in this scenario, namely *veritic* and *reflective* epistemic luck. They truly endanger one's claim to know.

that thing at all (and not something else).³ But only concreta can have looks. Therefore, a phenomenal image is a concretum.⁴ At least, it *presents* itself as a concrete particular (even in cases where it does not *represent* a concrete particular). As concrete particulars, they are determinate in their non-dispositional, lower-order properties at a moment in time (see fig. 2).

(I) Images are *qua images* concrete particulars and thereby determined in their low-order properties.

If the determinacy-claim for images is correct, then my phenomenal imagery of a tiger in my dream must have had a determinate number of stripes. Then, we ought to expect one true answer for Q_t .

Concerning first-person access, it is commonly held that necessarily, if something was going on in the mind of a person, then that person knows about it.

(O) All phenomenal features of an experience are known to the experiencer.

Such a kind of auto-omniscience cum infallibility has been defended or assumed quite often in philosophy (for historical examples, see Alston, 1971). The omniscience-claim can be motivated by etymology: The latin root of consciousness (conscientia) means "knowing with", which also holds for other European languages.⁵ So if phenomenality is a type

⁵See, for example, Christian Wolff (1983: Vol. I, Chap. 3, §194) for the German equivalent to consciousness, *Bewusstsein* as "das bewusste Sein", the known being. Other European languages trace their lexical analogues back to either the latin or germanic root, such as *Bevidsthed* in Danish, *Coscienza* in Italian, *Consciencia* in Spanish. The Swedish *Medvetance* is a direct translation of the meaning of *conscientia* (literally: "with knowledge"), and the same appears to hold for the japanese *ishiki*, where the kanji (意识) suggests itself to be read as "idea that is known". In all of these cases, these words suggest that something mental comes "with knowledge" or "is known".

³ One must worry about fallacious reasoning here which Block (1983a) pointed out. I try to address these separately in a different chapter. Here, I say that what speaks for phenomenal imagery is its spatial and temporal distribution, that they represent non-symbolically, that there is no grammar, etc. These do not prove imagery, but are not prone to the same kind of introspective failures pointed out by Block.

⁴Claims that lend themselves to some form of determinacy of pictorial representations can be found in e.g. Richard Wollheim (1987), who argued for the twofoldness of pictorial appreciation: We can focus on the content of the picture (what is shown) or the surface of the picture (how it is shown). The same holds for experiences themselves: We can focus on what is represented in them or on how they represent. If we focus on how, we focus on their imagistic surface-quality. If we think that an image i represent some a, we presume an isomorphism between the representing i and the represented a: A picture of the Eiffel Tower must share some structural similarity to the Eiffel Tower in order to be recognised as a picture of the Eiffel Tower. Clearly, we can be wrong here, e.g. if we have no idea what the structure of the Eiffel Tower is or if there was no Eiffel Tower. But even then, in order to take something as a pictorial representation, we take it as a concrete particular that shares structures with what it depicts. These isomorphisms can be between monadic properties (having the same rust-color) or n-adic structures and relations. The same holds for iconic representations (like the walking man on a pedestrian stoplight or the "rubbish bin"-icon in our visual operating systems), but with a twist: In iconic representations, a picture is still a concrete particular and presents itself thusly, but it represents a paradigmatic (in the case of the rubbish bin) or prototypical instance (in the case of the pedestrian stop light) of something. Imagistic (or pictorial) and iconic representations therefore share a way in which they present themselves (as concrete particulars) in which their ability to represent relies.




FIGURE 2. Even if I do not know how many stripes Tony the Tiger has, and even if this image of Tony the Tiger is silent about his total number of stripes (because we cannot see his lower portion or the back of his head such that stripe 2 and stripe 13 might or might not be parts of the same stripe), I can count the stripes that this image of Tony the Tiger presents. This image *presents* Tony with 13 stripes, even though it is indeterminate with how many stripes Tony is *represented* by the image.

of consciousness, then phenomenal aspects must be known. The omniscience-claim can also be defended by epistemic authority: The experiencer knows best what happens in her mind, so if a subject does not ascribe herself knowledge, then there is nothing to be known. (Otherwise, we would need to claim that we know better than the experiencer what happens in her experience, because we believe that there *is* something to be known.) One might also defend it by Verificationism (Cohen and Dennett, 2011): The existence of unknown phenomenal experiences cannot be verified, therefore we should reject this possibility; ergo, all phenomenal experiences are known. If all phenomenal features of experiences are non-public, then at least the experiencer must know.

If we accept all these claims, they lead us to the following paradoxical rendition of my situation:

- (P1_t) I had phenomenal imagery as of a tiger in my dream and by being imagistic that phenomenal image must have had a determinate number of stripes. (From 1_t and I)
- $(P2_t)$ I do not know the number of stripes on that dream-tiger-image. (From 2_t - 5_t)
- $(P3_t)$ I know all phenomenal features of my experiences. (From O)

The problem generalises if we see the number of stripes as a stand-in for some arbitrary lower-order property F that can present itself imagistically and can be ascribed to some specific part a (or succession or distribution of parts) of the "phenomenal field" of a person S (e.g. an arbitrary section of the visual field, e.g. an area which I would take as a tiger in the case of perceiving a tiger, an arbitrary pitch-range in hearing, a succession of flavours, etc.) which is taken to be an imagistic presentation. P1_t says that phenomenal imagery is such that it either has or does not have an imagistically presentable lower-order property—and collections of stripes are such a property. Concerning some a, S can then ask: Is a (e.g. the tiger in my dream) F or not (e.g. having an even number of stripes)?

- (P1) DETERMINATE IMAGERY REALISM: (i) S, the experiencer, has a phenomenal image a and (ii) if S has a, then (for some imagistically presentable property F) a is either F or not.
- (P2) HESITANCE: It is possible that S, the experiencer of a phenomenal image a, does not know⁶ whether (S knows whether)⁷ a is F or not.⁸
- (P3) Omniscience: Necessarily, a phenomenal image a is F (or not-F) if and only if S, the experiencer of a, knows that a is F (or not-F).

P1, P2, and P3 form an inconsistent triad. If there is one true answer to Q_t (P1) and if I do not know this answer (P2), then I am not omniscient (reject P3). If I do not know the answer to Q_t (P2) and if I am supposed to be omniscient about the realm that Q_t asks about (P3), then there cannot be an answer to Q_t (reject P1). If there is a true answer to Q_t (P1) and if I am omniscient about what Q_t asks about (P3), then I know the answer to Q_t (reject P2). Somehow, we cannot hold all the premises simultaneously, even though they seem intuitive.

What are our option to resolve such "Striped Tiger Problems"?

 $^{^6\}mathrm{A}$ weaker reading would be "does not ascribe herself knowledge".

⁷This depends on whether we accept or reject the KK-rule. The hesitancy claim can be turned into a case of anti-luminosity, if we reject the KK-rule (Williamson 2002: ch. 4).

⁸Hesitancy in reporting as well as hesitancy in acting are both intended to be captured by P2, so the "know whether" is supposed to capture forms of knowing-that, knowing-how-many, and classical anti-intellectualist knowing-how.

 $^{^9}$ Compare Alston (1971: 230) for this rendition and for references to *loci classici*. The only more than cosmetic alteration made to Alston's rendition is that P3 here is missing the asymmetry-claim "[...] while no one else is so related to such propositions [(ascribing mental states to himself)]". This is done so because in the current context, asymmetry does not matter, and the omniscience of one person about x does not conceptually entail that others cannot be omniscient about x.

- 2.2. Three Responses: None, Not Known, Neither-Nor. There are several possible responses to such Striped Tiger Problems, marked by rejecting at least one of P1–P3. Let us start from the back:
- (Reject P3) Not Known: There can be phenomenal facts (e.g. facts about visual phenomenal imagery of tigers and the number of stripes they present) that are not known to the introspector. In principle, we can be wrong in our judgements about our own consciousness (a failure of infallibility), or some phenomenal facts might elude us (a failure of omniscience). So there might be phenomenal features of mental imagery where the experiencing subject fails to know. If we reject P3, my inability to answer Q_t rests on a lack of knowledge.
- (Reject P2) No! Our hesitance to assent to presented options, our unwillingness to assert some answer, the low credence of subjects in their forced assertions (and the divided opinion in the scientific community about what is going on) is strong empirical evidence for P2: We, as experiencers, do not know whether we know. Rejecting P2 means going against evidence.

P1 is a conjunct of an existential claim about something (there is something like phenomenal imagery) and a claim about how that thing essentially is (necessarily, if a is an image and F is an imagistically presentable property, then a is F or not). Therefore, there are two ways to negate it: Reject the first conjunct or the second.

(Reject P1a) None: We reject the existential claim concerning phenomenal imagery: S has no phenomenal imagery (at least, in that case). Thus, my apparent inability to answer Q_t is due to the inexistence of the imagistic phenomenal experience of a striped tiger in my dreams. The question Q_t is loaded. Assenting to any given number

 $^{^{10}}$ These positions match those presented by Fantl and Howell (2003), minus Alterationism, which states that the number of stripes on the tiger might change over time. Alterationism is not a permissible answer if we make our question more precise: How many stripes does the phenomenal tiger image of a subject S have at the exact moment t_i ? (There might be still some complications if phenomenal moments are slightly extended, but let us bracket this for the moment.) An alteration needs at least two moments between which a switch can take place, and is therefore not an adequate answer to this more precise question limited to just one moment. Granted that if we ask how many stripes a tiger has in moments t_i to t_{i+n} , then there might be no one true answer due to constant alterations. But Alterationism still suggests a form of nonomniscience concerning one's experience, as the change between images is apparently not noticed. If we would experience the change and could access the number of stripes in all instances, then we could easily answer e.g. "32 to 64 stripes with a median of 44." Thus, Alterationism does not arise for time-slices, can be reduced to other options, and additionally suggests that we reject Omniscience. I will therefore not treat it separately here.

Another possible answer might be Paradoxism, which claims that phenomenal imagery can both exhibit and simultaneously lack a property F (see Crane, 1988)): The effect of the Waterfall Illusion, says Crane, is that I see something both as stationary and moving. I lack the space to discuss Paradoxism in details here, but I believe that some of my responses to Indeterminism will affect this position as well.

¹¹ See e.g. Williamson (1990, 2002); Schwitzgebel (2011).

¹² One theory compatible with this would be the cassette theory by Dennett (1976) where dreams are merely false memories, but not sensations during sleep.

would mean accepting the existential presupposition. The right answer would be: to challenge it.¹³

(Reject P1b) NEITHER-NOR: We accept the existential claim, but reject that phenomenal facts are determinate: S can have a phenomenal image a even though a is neither F nor not-F. Phenomenal imagery can be indeterminate, and my apparent inability to answer Q_t is simply due to the nature of the phenomenal fact itself. I got it right: I know that there is nothing to know and that is what my hesitancy expresses. 14

While Phenomenal Imagery Eliminativists (None) and Indeterminists (Neither-Nor) prefer an ontological strategy in order to save forms of epistemic privileges for the experiencing subject, Phenomenal Imagery Epistemicists (Not Known) prefer to treat the problem as a failure of our epistemic access in order to save a more standard ontology for phenomenal facts.

I mapped out a problem stemming from partially reflected intuitions, the Striped Tiger Problem, illustrate that it generalises and how we may respond to it. But is this really an interesting class of problems? I think so.

2.3. **XXX** Tigers in the Wild: Experimental and Philosophical Examples. Are there any cases of this kind beside me and my weird obsession with the tiger in my dreams? Yes, as I illustrate below.

I think that the problem crosscuts some interesting debates. Cases which can be reconstructed in this fashion arise not only for dream experiences. Additionally, they do not arise only for highly abstract properties like the even or uneven cardinality of sets of features (e.g. stripes on a tiger) but also concerning allegedly primitive "qualia" (e.g. redness) in some conditions like experiencing in the visual periphery.

Importantly, they arise both for perceptual as well as aperceptual cases like dreams or hallucinations. The problem is therefore not about the specificity, accuracy, precision or acuity of how we get to know the world through our sensory organs, or whether percepts are like photographs.¹⁵ It could be that our perceptual experiences are immensely imprecise, but the phenomenal image in these cases is determined about the whole visual field, like a blurry image is determined in how it presents itself even though we cannot make out whether it depicts a lumberjack or Sasquatch. Some answer available in the case of

 $^{^{13}}$ For illustration: This resembles a Russellian explanation (2001) of why I hesitate to answer "Is the present king of France bald?" Answering either way is problematic, as non-existent kings are neither bald nor hairy. The right way to answer is to say: "There is no present king of France!" Q_t resembles loaded questions like "Have you stopped beating your grandparents?" The acceptable answers suggested by the question are yes and no, but giving either forces one to admit that one once did beat one's grandparents. Similarly, Q_t suggests as answers natural numbers, but either forces one to admit that there are phenomenal images of tigers. We might overlook this presupposition and the right way to challenge it and thereby hesitate to give any answer.

¹⁴ This answer resembles the Intuitionist's explanation of why I hesitate to answer "Is it the case that every even number larger that 2 is the sum of two primes?", namely by rejecting tertium non datur for the fact in question (until there is a way to prove it in the mathematical case.

¹⁵ For this, see Block (1983b).

perception (e.g. sensory motor contingency answers) do not easily apply to the aperceptual cases. Additionally, the issue does not seem to be easily settled by empirical evidence. Instead, the three responses None, Not Known and Neither-Nor seem to be general ways to interpret data that suggests that individuals fail to know something about their minds. (In at least some of these cases, positions corresponding to None, Not Known or Neither-Nor can be attributed to participants in the debate.) Here are a few examples with the corresponding interpretations.

XXX The following is only a very schematic illustration.

- (Ex1) **Dream coloration.** Question: Is my dream imagery coloured or in black and white (Schwitzgebel 2002, 2003)? Description: People differ in their answers to whether their dreams are coloured or not. Some contradict themselves, and if given the option, up to 20% of people answer neither coloured nor black-and-white. NOT KNOWN: People have only limited access to their dream imagery (Rosen, ming). None: People have no dream imagery (Malcolm, 1959; Dennett, 1976). Neither-Nor: People's dream imagery can be neither coloured nor black-and-white Schwitzgebel (2011).
- (Ex2) Phenomenality in the visual periphery. Question: Do I have concrete but unnoticed phenomenal imagery of the perturbations in my visual periphery in crowding experiments (Freeman and Simoncelli, 2011)? Description: With eye-tracking, it is possible to change a stimulus in real-time depending on where a subject looks at any given moment. We can present a subject with a visual stimulus that looks natural at the fovea (the center of one's visual field) but is systematically disturbed in the periphery. Subjects then report seeing a whole scene, but apparently do not notice the ongoing perturbations in the periphery. This has been shown with texts and reading as well as with photographs and illustrations (see figure 3 for an example). Not Known: The phenomenal character of experiences the periphery is determinate, yet we lack epistemic access to it. None: We lack phenomenal imagery in the periphery. Neither-Nor: The phenomenal character of experiences in the periphery is neither of this nor of that.
- (Ex3) Inattentional Blindness. Question: Did I have a phenomenal image as of a gorilla in that inattentional blindness video (Simons and Chabris, 1999)?¹⁷ Description: Subjects are shown a video of two teams passing basketballs, with the instruction to count the passes of one of the teams. In the middle of the video, for 9s (5s in another condition) a person in a gorilla costume walks through the passing players and thumps her chest, while the players pass around her Yet, from 204 subjects, only about 50% report or remember this strange occurrence. Not Known: There either was or wasn't a gorilla in the subject's phenomenal imagery corresponding to the video, but the subject cannot access whether there was a gorilla present in her visual experience. None: There was no gorilla present in

 $^{^{16}}$ A meta-study covering the last 80+ years as well as cross-cultural data can be found in Schwitzgebel (2011).

¹⁷The video can be seen here: https://www.youtube.com/watch?v=vJG698U2Mvo.



FIGURE 3. A stimulus used in peripheral crowding experiments. The point of fixation is indicated by where those lines indicated outside the picture would intersect. If I am presented with the second stimulus in an experimental setting, do I have phenomenal imagery as of a human shape in the periphery?

experience. Neither-Nor: Neither was there a gorilla present nor was there not a gorilla present in experience.

- (Ex4) Change Blindness. Question: Before I notice what changes, did I have changing phenomenal imagery in change blindness experiments (Rensink et al., 1997; Levin and Simons, 1997; Simons and Levin, 1997)? Description: Subjects are presented with a visual stimulus, followed by grey for roughly 0.5 seconds, followed by the same visual stimulus with a slight alteration (see figure 4 for an illustration). This cycle repeats indefinitely. Even though one has the experience of a whole scene, the detection of what changes between the presentations may take up to 50 seconds. Not Known: We do not immediately access all features of our experiences, including the color of the vendor's trousers (Block, 2007). None: There is no change of the vendor's trousers present in our experience until we notice this change, i.e. we have an "illusion of seeing" (Naccache and Dehaene, 2007). Neither-Nor: We have phenomenal imagery of the vendor's trouser which is neither grey nor blue until we notice the change.
- (Ex5) The Sperling Experiment. Question: Did I have a phenomenal imagery of all the letters I am unable to name in the Sperling experiment before being cued to repeat one line (Sperling, 1960)? Description: Sperling presented subjects for half a second with a 3×4 grid of letters. In the first experiment, subjects were unable to repeat more than about 4 letters of the whole grid they were presented with. In the second experiment, subjects were given an auditory cue after the presentation of the grid, which indicated by pitch (low = lowest line, mid = middle line, high = highest line) which line they ought to repeat. In this case, subjects were nearly faultless in stating which letters were in that line. NOT KNOWN:

Subjects cannot access all information which is present in the phenomenal imagery before the cue. The cue somehow enhances selective access (Block, 2007). None: Subjects have no experience as of letters before the cue. The letter detection is unconscious and access is selectively enhanced by the cue. Neither-Nor: Neither do subjects experience this or that letter before the cue, but only letters. Access to the unconscious information which letter was presented is selectively enhanced by the cue (Brown, 2012; Stazicker, 2011).

Interpretations corresponding to None, Not Known or Neither-Nor can be applied to the card experiment by Bruner and Postman (1949), the speckled hen (Chisholm, 1942; Tye, 2009, 2010), color experiences in the visual periphery (Dennett, 1991), experiences before they are altered by cognitive penetration, the specificity of Charles-Bonnet-hallucinations or psychedelic highs, experiences outside of the focus of attention more generally, whether coins look round (Schwitzgebel, 2011; Moore, 1953; Peacocke, 1983), and so on. (I will devote a chapter to spelling this out.)

Such a wide range of examples suggests that Striped Tiger Problems might be at the heart of consciousness studies. Because they rest on an inconsistent triad, they demand an answer. But, as long as we cannot make conscious experience intersubjectively accessible, no empirical data seems to settle which of the premises we ought to reject: If a subject hesitates to give an answer, no evidence seems decisive for what the nature of phenomenal character might have been in these cases determinate, indeterminate, or inexistent. All of the three responses considered here seem to account for (at least some of) the data. Whatever response we choose rather corresponds to a decision guided by which of the intuitions behind P1–P3 we deem weakest.

But are all three equally strong? I do not think so. I deem NEITHER-NOR dialectically weak, specifically in the context of a search for neural correlates of consciousness. I first clarify this position and suggest proponents in the next section before presenting some of its shortcomings. (Not every single one of them might be convincing, but in total they may be sufficient to lead to a rejection.)

3. On Neither-Nor

3.1. Kinds of indeterminacy. Let me clarify what indeterminacy is not. First, it is not the same as vagueness. Although my tiger-image might be indeterminate whether it has 42 stripes or not, there is nothing vague about the predicate having 42 stripes. Having 42 stripes or being the letter A (as in the Sperling experiment) are not vague predicates like being bald or being a heap. There are no borderline cases for having 42 stripes or being the letter A. This does not mean that vagueness might not play an important role in some cases of phenomenal experience and introspective access. But issues of indeterminacy are not automatically issues of vagueness. It is also not the same as contradiction: being neither even nor uneven is not the same as being even and uneven. The number π is neither even nor uneven, but it certainly is not both even and uneven; snails might be hermaphrodites

¹⁸For example, it might be vague whether speck counts as a stripe or not. But by stipulation that is or is not a stripe, we would arrive at a definite answer to Q_t because we know whether to count it or not.

and therefore both male and female, but they are not neither male nor female. (Their shells are neither male nor female, but not both male and female.)¹⁹

So, what might Phenomenal Indeterminism (Neither-Nor) be? There are at least three kinds of indeterminacy we ought to distinguish:

- (ind₁) **indeterminate:** If a is indeterminate concerning property F, then neither is a F nor is a not-F. Plausible Reasons might be:
 - (ind_{1a}) a is the content of a representation, the intentional object (Brentano, 1995: 68). For example: In "The Hound of Baskerville", Sherlock Holmes is indeterminate concerning his blood type.
 - (ind_{1b}) Ascribing F to a is a category mistake. For example: The first prime number with seven digits is indeterminate concerning its blood type.²⁰
 - (ind_{1c}) a does not exist. For example: The present queen of France is indeterminate concerning her blood type (Strawson, 1950).
 - In all of theses cases, indeterminacy is interpreted as an ontological claim about a: In ind_{1a} and ind_{1b} it is combined with an existence claim and a claim about the nature of a, and in ind_{1c} with rejecting that a exists.
- (ind₂) **indetermined:** If a is indetermined concerning F at t_i , then we do not know at t_i whether a is F or not, but we assume that either a is F or a is not-F and that we might know which at a later moment t_{i+n} . For example: In the period after the white smoke has left the chimney of the Sistine Chapel, but before we hear the Habemus Papam! from the balcony of St. Peter's Basilica, the blood type of the next pope is indetermined to us (even if we knew the blood type of each of the cardinals). Here, indeterminacy is interpreted mainly as a claim about a person's specific epistemic situation.

 $^{^{19}}$ If we are thinking in logical terms, assenting to indeterminate phenomenal facts suggests that one might try to capture these facts in a kind of intuitionistic logic, where tertium non datur does not hold and some propositions may lack a truth value. If phenomenal facts are thought of as contradictory, then they might be better modelled in a kind of paraconsistent logic, where $Fa \land \neg Fa$ can be true. If one believes that phenomenological statements like "The tiger in my dreams had an even number of stripes" can lack a truth value — as Phenomenal Indeterminists suggest — one needs the right logic as a background assumption. One's assent to Phenomenal Indeterminism or Paradoxism might then also depend on one's willingness to reject bivalence in order to defend Indeterminism or Paradoxism for phenomenological statements. For many, like me, this might already be a bit much.

²⁰The reason why this might be seen as a case of indeterminacy could have to do, again, with presuppositions: The question "What is the blood type of the first prime number with seven digits?" puts the presupposition that the first prime number with seven digits has a blood type into the common ground (Stalnaker, 2002). If this presupposition is not challenged, we are baffled at what the right answer might be. This is similar to the question "Have you stopped hitting your grandparents?" You can answer yes or no, but then you thereby accept the implicature that you hit your grandparents. The best response, however, would be to challenge the presupposition. But we do not always employ this strategy. Our subjects in the examples fail to challenge the false presupposition triggered by the question, sometimes due to the experimental procedure prohibiting it.

(ind₃) **indeterminable:** If a is indeterminable concerning F, then we cannot know whether a is F or not, but we assume that either a is F or a is not-F. For example, Gaius Julius Caesar's blood type is indeterminable (given that time travel is impossible). Here, indeterminacy is interpreted as a claim about our epistemic access in general.

Which of these is the kind of indeterminacy inherent in the Neither-Nor-response? Either ind_{1a} or ind_{1b} seem to fit best, for ind_{1c} is equivalent to None, while ind_2 as well as ind_3 are epistemic claims wherefore they correspond to Not Known.

NOT KNOWN has been defended by Ned Block (1997b,a, 2003, 2007) or (arguably) Paul Churchland (1985), while None appears to be held by (arguably) Daniel Dennett (1996, 1991, 2006, 1976, 1993) and explicitly by George Rey (1983, 1995), Michael ? and Keith Frankish (2016). What about Neither-Nor? Who are the proponents for this flavour of phenomenal indeterminacy?

- 3.2. **XXX** Who defends Neither-Nor? Here, I analyse positions in the history of philosophy and analytical debate concerning how they fit NEITHER-NOR. This section is massively abbreviated here. I welcome feedback on 3.2.3.
- 3.2.1. XXX Non-analytic philosophy. Husserl mentions indeterminacy in Ding und Raum (1973: 57f)(1907/1973: 57f) and in a few passages in Über Wahrnehmung (2005b) as well as Mittelbare und Unmittelbare Auffassung (2005a). But Husserl distinguishes between appearance and apprehension, where the latter concerns the cognitive grasp or understanding we have of the appearances. Thus, for Husserl, indeterminacy is not an aspect of phenomenal imagery or character itself, but in how we understand it what we take this experience to be an experience of.²¹ Husserl then does not defend Neither-Nor for phenomenal imagery, but only for apprehension of phenomenal imagery.

Merleau-Ponty (2002: 6) talks of "indeterminate vision" in his *Phenomenology of Perception* and talks of indeterminacy as "a positive phenomenon". But in his example, the Müller-Lyer-Illusion, he focuses on the lines outside of our mind, not the lines as we experience them. He also claims that either line in the illusion does not "belong to the same universe as the other", which appears to be false as we clearly see both lines in one experience and can empirically test influences on their comparative appearance. He further conflates indeterminacy, context-shift, ambiguity and contradiction, speaking of all of them in the context of the Müller-Lyer. Better not to count Merleau-Ponty as a defender of Neither-Nor.

3.2.2. XXX Contemporary analytic philosophy. In the current analytical debate, Neither-Nor is mainly associated with Representationism: In his discussion of the speckled hen,

²¹ Michael Madary suggested that Husserl might have changed his position on the appearance/apprehension distinction by the 1920s. At that time Husserl wrote of "determinable indeterminacy", which seems to map to what I call "indetermined". That Husserl is not a proponent of phenomenal indeterminacy, at least in his early writings, is supported by passages where he writes that the appearances, but not the apprehensions of these appearances, are fully determined (especially Husserl, 1973: §6).

Michael Tye (2009, 2010) claims that we can experience specks on a hen without experiencing each speck. For Tye, the fact that our experiences are indeterminate is evidence for Representationism. Similarly, Richard Brown (2012) argues against Block's idea of phenomenological overflow, especially in the Sperling experiments. On the question what the experience was like before being asked to recite one of the presented lines of letters, Brown says this pre-cue experience is "fuzzy": We experienced letters, but not specific letters like we experience something but not that thing when the optical situation is not perfect, e.g. when we do not wear our glasses. He ties this to the HOT-Theory, which identifies consciousness with the content of a higher-order thought. Another example comes from Eric Schwitzgebel (2011: 12f,40f), who asks why there is such divergence in reports concerning dream coloration — and why some people even say that their dreams are neither coloured nor black-and-white. He draws the analogy of colours in dreams to colours mentioned in novels ("She arrived in a wildly colored dress.") or black-and-white pictures (e.g., a blackand-white photograph of coloured bell peppers). In the novel case, the representation is indeterminate vis-à-vis the colour the dress had, but it is explicitly represented as *coloured*; in the case of the black-and-white photo, it is obvious that the bell pepper is coloured all bell peppers are — but it is indeterminate vis-à-vis the colour the photo represents. Schwitzgebel also defends this as a possibility for waking imagery.

3.2.3. Grush's third format. A more complex case is Rick Grush (2007) on Generic Phenomenology. In a response to a target article by Block (2007), who defends the distinction between access and phenomenal consciousness, Grush suggests an alternative interpretation for cases like the Sperling experiment: generic phenomenology, which is the idea that sometimes we are presented only with genera, e.g. not a specific letter, but simply the abstract feature being a letter. Grush claims that phenomenology is a specific type of format for a mental representation which is neither picture-like nor sentence-like nor a combination of the two. This new format "presents its content as affording answers to queries of detail" and it "presents its content as clear without presenting the details." To some degree, Grush suggests that the indeterminacy of such experiences rests on the content of the phenomenal experience. But in addition, the presentation itself allows for some degree of indeterminacy due to being in a special format.

Is there such a third format that is neither-language-nor-image-like? What is a format supposed to be? The format of a representation ought to be discernible from content and vehicle. The vehicle of a representation is the physical object at a specific spatiotemporal location that does the representing, e.g., some pressure waves in the case of a spoken word, some marks on a blackboard in the case of some written sentence, some muscles being moved in the case of a sentence being spoken in sign language. The content of a representation is what is represented, e.g., ■ in the case of the words "red square" or a specific imaginary tiger in figure 2. Grush suggests that images and sentences are formats. A format is then determined by certain features that govern how the vehicle represents the content. Sentence-like formats represent by using symbols (e.g. words or graphemes) which are arranged in a conventional manner (a grammar) by which more complex representations can be built compositionally. Picture-like representations have no clear distinct symbols

or identifiable grammar, but represent instead by mirroring their content to some degree. They thereby fulfil criteria for being analogue representations (Haugeland, 1981; Lewis, 1971; Jackson, 1960).

How does the format of a representation relate to vehicle and content? First, the content might remain the same while the format changes. For example, some photographs mirror the distribution of people from a certain vantage point, but a description of that scene does not mirror the distribution — it describes this distribution; the vibration of a loudspeaker's membrane mirrors the frequency of the sound waves produced by my voice, but a transcript of what I said does not — it may in addition describe my voice, but does not mirror it. So a photograph and its description or a recording and a transcript can have the same content, but they differ in format. Second, in some cases, the same vehicle can represent via a pictorial or a sentential format, for example in the case of hieroglyphs: If we take a certain physical object (see fig. 5) as a picture, it represents a duck; if we take it as hieroglyphic word, the very same physical object represents son, as in the honorific title "Son of Ra". The same vehicle then has a superposition of two different contents, duck and son, in virtue of the two formats it can represent in. So the same content can be represented by different vehicles in different formats, the same vehicle can represent different contents in different formats. Therefore, it is best to see format as not grounded in either vehicle or content.



FIGURE 4. The hieroglyph for the royal title "Son of Ra" from an obelisk of Thutmosis III at the temple of Karnak. The first symbol is in the shape of a pigtailed duck and stands for the word "son". (Adapted from a photograph by Gerard Ducher under a creative commons license.)

So could there be more than these two formats, images and languages, as Grush claims? It seems to boil down to the question: Does a vehicle represent by similarity or not? Or better: by a salient isomorphism or not? Here, the answer is either yes or no. 23

If yes, the vehicle represents pictorially: If we take the hieroglyph to represent by similarity, then the way it appears must be somehow isomorphic to the way its represented object would appear (at least from some vantage point). The isomorphism might be between lower-order properties (e.g., the way a printed colour in a catalogue mirrors the colour of the bucket of paint you might order from it) or the isomorphism might be between higher-order properties (e.g., the way the angles and length-proportions of a model mirror the building it is a model of). But even these higher-order properties always are instantiated by some lower-order properties which appear as well: A model must have some dimensions in order to represent the proportions of a specific building by isomorphism.

If no, the vehicle represents language-like: If a representation does not represent by similarities — like the words "black square" represent ■ or the way smoke represents fire —, then they do so by exploiting regularities (Grice, 1957). These might be conventional, as in the case of "black square". Or they might be natural, as in the case of smoke representing fire. But some regularity — may it be social, communicative, deliberate, or natural – must be exploited.²⁴

It seems then that there are only two formats: those that exploit similarities in order to represent and those that exploit regularities. (In some cases, these might be mixed, as in the case of maps, but either we take a sign on a map to represent by similarity or by regularity.) Those that exploit similarities are imagistic, those that exploit regularities linguistic. Representations are either pictorial or language-like — specific pictures or statements are only examples of these kinds of representation. Thus, Grush has not done enough to establish a third type of format that is neither-language-nor-image-like. We simply do not know what this would be.

But even if there were such a third format, this alone does not immediately entail anything about the appearance of a representation: The appearance of a hieroglyph does not alter if we switch from taking it to represent by similarity or by convention. Formats do not directly constrain appearance, as Grush seems to assume. But exactly this is what is at issue in the case of these experiences: How do they appear?

²²Or, more precisely, as we always ignore certain features of vehicles e.g. the weight of a photograph or the exact size of a written font does the vehicle appear similar to what it represents or not?

²³Here, it might look as if make the introspective mistake pointed out by Block (1983a: 509f), because it looks as if I presume a similarity between objects and image. I do not intend to do this. Instead, I think there is a relation between me and different ways in which I represent the object. In pictorial representations, I somewhat co-represent the image as providing me with structural isomorphism to the represented object; I have a stance toward these representations to preserve object-features in its presentation. I have no such stance towards non-pictorial representations.

²⁴ Obviously, similarities are regularities as well. But, in contrast to the ones mentioned here, similarities are regularities that need not be established over time, but can be detected in time slices.

The vehicles of a representation always appear in a specific way: The *ink-on-paper*-vehicle of the words "black square" still appear in a very specific way, namely as distributions of black on white. And the same goes for vehicles representing pictorially. For example, if " \blacksquare " represents a pool filled with ink seen from space, it still appears in a specific way.

Nearly all vehicles of symbolic representations can be seen as an imagistic representation with some phantasy. (Doesn't "8" look like a little bit like a snowman missing a head? Doesn't "Q" look a little bit like a cat from behind ducking her head?) Then, there must be a stance we take towards vehicles that determines whether we take them as imagistic or not. But the phenomenal character of a representation can be seen as somewhat independent of its specific format.

In effect, it is the stance we take to our mental representation that marks them out as being of one format or another. If we take the imagistic stance, we treat the representation as preserving structural features of the object represented. If we take the linguistic stance, we lack this expectation. Our natural stance towards our phenomenal representation is the imagistic stance, which explains why phenomenal experience is often seen as being transparent (Metzinger, 2003; Martin, 2002).

In conclusion, even if Grush's "third format" can be made sense of, it still does not follow that the format determines character in the way he presumes. At best, we can interpret his claim as a type of content-indeterminacy $(ind1_a)$.

3.3. **Summary.** Among contemporary Phenomenal Indeterminists, there appears to be unison: Tye, Schwitzgebel, Grush, and Brown explicitly assert that phenomenal experience can be indeterminate, and that it is phenomenal content — what presents itself in experience — which can be indeterminate.

In all cases of indeterminacy, the experience still represents a generic or higher-order property \mathfrak{F} without representing any of the specific subsumable properties F_1, F_2, \ldots, F_n , e.g. having specks but not having 1, 2, 3,..., n specks, being coloured but not being green, blue, red, yellow, ..., being a letter but not being A, B, C,..., Z. This matches Husserl's dictum that "complete indeterminateness is nonsense": Complete indeterminateness would mean that we cannot ascribe any property, not even highest-order ones. Our examples are not of this kind: There is always some higher-order \mathfrak{F} which we do ascribe to the phenomenal image; but we are uncertain which \mathfrak{F} -subsumable lower-order property F_1, F_2, \ldots applies to the specific case in question. More so: A higher-order property must apply in order to render the question meaningful: We must agree that the tiger was striped, for otherwise, the question "How many stripes did it have?" does not make any sense.

These contemporary philosophers agree that it is the *content* of the experience is indeterminate, and all state that at least a higher-order property \mathfrak{F} is experienced. For a claim of *content*-indeterminateness to become a claim of *phenomenal* indeterminacy, these proponents of indeterminacy need to assume that phenomenal character is grounded in (or identical to) some phenomenal content.²⁵ I think that this position has severe disadvantages vis-à-vis Striped Tiger Problems.

²⁵Supervenience does not suffice. See Block (2015).

- 3.4. Opposing Neither-Nor. The arguments are here merely sketched. A chapter will present them in full. I specifically welcome feedback on W1, I clarified the notion of Phenomenal Indeterminism (Neither-Nor) in contrast to Phenomenal Eliminativism (None) and Introspective Failure (Not-Known) in the preceding section. Given writings on indeterminacy of phenomenal experiences, it appears to be best construed as reliant on some form of Representationism. However, I believe that the Neither-Nor-option, understood this way, has several dialectical weaknesses with respect to the inconsistent triad. Here, I present a few of these weaknesses:
- (W1) No Explanation: If our experiences actually were indeterminate, the answer to Q_t would be easy: "The tiger in my dreams has no clear number of stripes." This would be the answer an Indeterminist should expect. But this is not the answer we give in these cases. Instead, we hesitate and show all behavioural signs of not knowing the answer. So, somehow, the Indeterminist must give an additional reason why we hesitate to answer. It seems we have to make further assumptions in order to explain our behaviour adequately, e.g. that the indeterminacy of experiences is introspectively inaccessible. Without such an additional assumption, Phenomenal Indeterminism does not explain the main fact leading to Striped Tiger Problems sufficiently.
- (W2) Content-Character Confusion: The Indeterminists discussed before focus on the content of the experience (what is represented in experience). Yet, in some cases (dream coloration, the number of stripes on my tiger, or change blindness, etc.) we are interested in the character of experience (how something is represented in experience). We are not obviously interested whether my dream is about a tiger of 32 stripes or not, but whether the phenomenal image of the tiger in my dreams does present it as having 32 stripes or not. But content and character are neither prima facie nor a priori the same. For example, we can give different answer to what I dreamt of vs. how I dreamt of something, or we will get different groupings if we sort our mental states by content-similarity rather than character-similarity. ²⁶ That phenomenal content can be indeterminate does not directly entail that phenomenal character can be indeterminate. Content Indeterminists thus do not directly answer what is at issue in a Striped Tiger Problem, namely whether phenomenal character can be indeterminate.
- (W3) **Dependency on Strong Representationism:** Content-indeterminateness is only an answer to phenomenal-character-indeterminacy, if we accept a strong form of Representationism. Weak forms of Representationism claim that character supervenes on content. But because supervenience is compatible with forms of Dualism, the indeterminacy of content then does not entail the indeterminacy of character. Additionally, such a Supervenience Representationism merely claims that a change in character requires a change in content. From this, it does not follow

²⁶Content-wise, the experience of caramel is more similar to the experience of honey than to the experience of motor-oil. But character-wise, a visual experience of caramel is more similar to the experience of motor-oil than to the experience of honey a fact exploited by advertisement agencies when filming caramel.

- that character is indeterminate, only that it can change. Something stronger is needed in order to validate the inference from content-indeterminacy to character-indeterminacy, e.g., an identity or a grounding claim. This, however, weakens the claim of phenomenal indeterminacy because its acceptance is now reliant on the acceptance of Strong Representationism, for which one has to argue independently.
- (W4) **Trivialisation:** If we do reduce character to content, indeterminacy becomes trivial because content-indeterminacy is ubiquitous. The sentence "The cat is on the mat" is indeterminate concerning whether the cat is a Siamese or a Maine Coon, or whether the mat is filled with feathers or made of straw. Every representation short of representing exactly one possible world is content-indeterminate concerning something. Thus, content-indeterminacy is ubiquitous. Yet, phenomenal indeterminacy was supposed to be a *special* feature of some *but not all* of our experiences, e.g., of anomalous conscious states like dreaming, or of parts of our experiences like the visual periphery, or of experiences in uncommon experimental conditions. If we identify character with content, then phenomenal indeterminacy is present in every single perception and in all its parts. It becomes so ubiquitous that we might actually wonder why we need to refer to such weird circumstances to notice it in the first place.
- (W5) Incompatibility with Non-Symbolic Representational Formats: Brown, Grush, Schwitzgebel, and Tye assent to the presentation of higher-order properties without subsumable lower-order properties, i.e., a plentitude of specks but not a specific number of specks, a colour but not a specific colour, a grid of letters but not of specific letters. Thus, something abstract is presented without something concrete. But this is incompatible with pictorial (or even quasi-pictorial) formats (Kosslyn 1980: 31). In symbolic formats, the representation of higher-order properties \mathfrak{F} without lower-order properties F_1, F_2, \ldots doing the representing is possible: "speckled" does not mean having 97 specks, "coloured" does not mean brown, blue, amber, or vermillion, "letter" does not mean A, B, or Q. But nothing like that is possible in pictures. Because pictorial representations work non-symbolically, ²⁷ a higher-order \mathfrak{F} needs to be represented by some concrete and subsumable lowerorder property F_i : A picture of a speckled hen presents itself with a countable number of specks; a black-and-white picture of a bell pepper uses a determinate grey; a grid of fuzzy letter-shapes is has a specific black-white-distribution (see fig. 6). So if phenomenal imagery is pictorial, then it is presentation-determinate even if it is content-indeterminate.
- (W6) Incompatibility with the Identity Theory: If a and b are identical, a and b share all properties. For phenomeno-neural identification, this has been questioned by proponents of the Identity Theory: Place (1956) argued that just because there is nothing green in my brain when I experience a green afterimage, this does not entail that the green afterimage can not be identical with a brain state. For lower-order

 $^{^{27}}$ Certainly, some pictures can be symbols or icons, but they do not represent pictorially in that case, but iconically or symbolically. See above.



FIGURE 5. Examples of the compatibility between content-indeterminacy and presentation-determinacy: A speckled hen; a coloured pepper; two fuzzy letters.

properties and the cognitive act of identification, this may hold. But there are limits to identification despite discriminability. For example, we cannot identify what differs in its highly abstract properties or its ontological category: if a is a temporally extended process (e.g. playing chess) and b is a state (e.g., the rook being on a4), then $a \neq b$ and a fortiori: $\Box(a \neq b)$. The indeterminate/determinate-distinction is of such a kind: it is highly abstract and the Phenomenal Indeterminists see it as somewhat ontological. If so, it is questionable whether Phenomenal Indeterminacy is compatible with Identity Theory, because a neural state or event is never indeterminate. Instead, it is — as a macro-physical event — fully determinate. Therefore, if we accept the character-indeterminacy of phenomenal experiences, it seems that we cannot easily identify phenomenal experiences with neural events. (Anybody who rejects the Identity Theory for other reasons may not care about this.)

- (W7) Questionable Motivation: The triad that constitutes a Striped Tiger Problem is such that P1 must be rejected if we want to save omniscience. But at least in scientific contexts, it seems odd to claim that we have absolute epistemic access to a section of reality or a class of natural phenomena. Consciousness would be weirdly unique in this way. Instead, we should in general doubt claims that anyone has perfect epistemic access to any section of reality, including consciousness.²⁹
- (W8) Giving up Too Much: If we reject Omniscience (P3), then there is no need to reject a further premise (like P1), for the problem simply does not arise anymore. But an Indeterminist must give up both the determinacy claim in P1 and P3 (omniscience), for otherwise subjects would address the indeterminacy in question in their answers to Q_t (see W1).

Even if not all weaknesses W1–W9 are fully convincing, acceptance of any of the short-comings pointed out here should weaken one's confidence in the claim that phenomenal

 $^{^{28}\}mathrm{See}$ e.g. Tononi (2004, 2008) for how to determinately individuate neural states.

²⁹This is also more on par with the anti-authoritarian view of scientific inquiry

experience are character-indeterminate. Then, Phenomenal Indeterminacy is less convincing and in a dialectically weaker position than its competitors. I therefore propose to give up Neither-Nor in favour of the other two options.

The rest of the book argues that neither None nor Not Known are suitable general strategies. Instead, we have to determine on a case-by-case basis if we want to ascribe ignorance or a lack of phenomenal features to a person. For this, I present a model of introspective access that distinguishes different sources of error. I then present in a case study (a comparison between psychiatric reasoning concerning Charles-Bonnet-Syndrome and Anton-Babinski-Syndrome) how this model may be applied to justify one's assessment of another person's consciousness.

References

Alston, W. 1971. Varieties of privileged access. American Philosophical Quarterly, 8(3):223-241.

Block, N. 1983a. Mental pictures and cognitive science. Philosophical Review, 92(4):499-542.

Block, N. 1983b. The photographic fallacy in the debate about mental imagery. $No\hat{u}s$, 17(4):651-62.

Block, N. 1997a. Begging the question against phenomenal consciousness. In Block, N., Flanagan, O., and Güzeldere, G., editors, *The Nature of Consciousness: philosophical debates*, pages 175–180. Cambridge, MA: MIT Press.

Block, N. 1997b. On a confusion about the function of consciousness. In Block, N., Flanagan, O., and Güzeldere, G., editors, *The Nature of Consciousness: philosophical debates*, pages 375–415. Cambridge, MA: MIT Press.

Block, N. 2003. Mental paint. In Hahn, M. and Ramberg, B., editors, Reflections and Replies: Essays on the Philosophy of Tyler Burge. Cambridge, MA: MIT Press.

Block, N. 2007. Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30:481–548.

Block, N. 2015. The puzzle of perceptual precision. In Windt, J. and Metzinger, T., editors, *Open Mind*, pages 1–52. Borges, J. L. 1976. Argumentum ornithologicum. In *Dreamtigers (El Hacedor)*. University of Texas Press.

Brentano, F. 1995. Psychology from an Empirical Standpoint. London: Routledge.

Brown, R. 2012. The myth of phenomenological overflow. Consciousness and Cognition, 21(2):599-602.

Bruner, J. S. and Postman, L. 1949. On the perception of incongruity: A paradigm. Journal of Personality, pages 206–223.

Chisholm, R. 1942. The problem of the speckled hen. Mind, 51(204):368-373.

Churchland, P. M. 1985. Reduction, qualia, and the direct introspection of brain states. *The Journal of Philosophy*, 82(1):8–28.

Cohen, M. A. and Dennett, D. C. 2011. Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, 15(8):358–364.

Crane, T. 1988. The waterfall illusion. Analysis, 48(3):142-147.

Dennett, D. C. 1969/1996. The nature of images and the introspective trap. In *Consciousness and Content*. London: Routledge.

Dennett, D. C. 1976. Are dreams experiences? The Philosophical Review, 85(2):151-171.

Dennett, D. C. 1991. Consciousness Explained. Boston: Little, Brown and Company.

Dennett, D. C. 1993. Quining qualia. In *Readings in Philosophy and Cognitive Science*, pages 381–414. Cambridge, MA: MIT Press.

Dennett, D. C. 2006. Sweet Dreams: Philosophical Obstacles to a Science of Consciousness. Cambridge, MA: MIT Press.

Fantl, J. and Howell, R. 2003. Sensations, swatches, and speckled hens. *Pacific Philosophical Quarterly*, 84:371–383. Frankish, K. 2016. Illusionism as a theory of consciousness. *Journal of Consciousness Studies*, 23(11-12):11–39.

Freeman, J. and Simoncelli, E. P. 2011. Metamers of the ventral stream. Nature Neuroscience, 14(9):1195–1201.

Grice, H. P. 1957. Meaning. Philosophical Review, 66(3):377-388.

Grush, R. 2007. A plug for generic phenomenology. Brain and Behavioral Sciences, 30(5/6):504-505.

Haugeland, J. 1981. Analog and analog. Philosophical Topics, 12(1):213-225.

Horikawa, T., Tamaki, M., Miyawaki, Y., and Kamitani, Y. 2013. Neural decoding of visual imagery during sleep. Science, 340(6132):639-642.

Husserl, E. 1907/1973. Ding und Raum, volume XVI of Husserliana. Den Haag: Martinus Nijhoff.

Husserl, E. 2005a. Über mittelbare und unmittelbare Auffassung. In Vongehr, T. and Giuliani, R., editors, Wahrnehmung und Aufmerksamkeit. Texte aus dem Nachlass (1893-1912), volume XXXVIII of Husserliana. New York: Springer.

Husserl, E. 2005b. Über Wahrnehmung. In Vongehr, T. and Giuliani, R., editors, Wahrnehmung und Aufmerksamkeit. Texte aus dem Nachlass (1893-1912), volume XXXVIII of Husserliana. New York: Springer.

Jackson, A. S. 1960. Analog Computation. New York: McGraw-Hill.

Levin, D. T. and Simons, D. J. 1997. Failure to detect changes to attended objects in motion pictures. *Psychonomic Bulletin & Review*, 4(4):501–506.

Lewis, D. 1971. Analog and digital. Nôus, 5(3):321-327.

Malcolm, N. 1959. Dreaming. London: Routledge & Kegan Paul.

Martin, M. 2002. The transparency of experience. Mind & Language, 17(4):376-425.

Merleau-Ponty, M. 1945/2002. Phenomenology of Perception. London: Routledge & Kegan Paul.

Metzinger, T. 2003. Phenomenal transparency and cognitive self-reference. Phenomenology and the Cognitive Sciences, 2:353-393.

Moore, G. E. 1953. Sense Data, chapter II, pages 28-40. London: George Allen & Unwin.

Naccache, L. and Dehaene, S. 2007. Reportability and illusions of phenomenality in the light of the global neuronal workspace model. Behavioral and Brain Sciences, 30(5/6):518–520.

Peacocke, C. 1983. Sense and Content. Oxford: Oxford University Press.

Place, U. T. 1956. Is consciousness a brain process? British Journal of Psychology, 47:44–50.

Pritchard, D. 2005. Epistemic Luck. Oxford: Oxford University Press.

Rensink, R. A., O'Regan, J. K., and Clark, J. J. 1997. To see or not to see: The need for attention to perceive changes in scenes. Psychological Science, 8(5):368–373.

Rey, G. 1983. A reason for doubting the existence of consciousness. In Consciousness and Self-Regulation, volume 3. New York: Plenum.

Rey, G. 1995. Towards a projectivist account of conscious experience. In Metzinger, T., editor, *Conscious Experience*, pages 123–142. Paderborn: Schöningh.

Rosen, M. G. forthcoming. What I make up when I wake up: anti-experience views and narrative fabrication of dreams. Frontiers in Consciousness Research, Special Issue: Contrasting Dreaming and Wakefulness.

Russell, B. 1905/2001. On denoting. In Martinich, A. and Sosa, D., editors, *Analytic Philosophy: An anthology*, pages 32–41. Oxford: Blackwell.

Schwitzgebel, E. 2011. $Perplexities\ of\ Consciousness$. Cambridge, MA: MIT Press.

Simons, D. J. and Chabris, C. F. 1999. Gorillas in our midst: Sustained inattentional blindness for dynamic events. Perception, 28(9):1059–1074.

Simons, D. J. and Levin, D. T. 1997. Change blindness. Trends in Cognitive Science, 1:261-267.

Sperling, G. 1960. The information available in brief visual presentations. *Psychological Monographs: General and Applied*, 74(498):1–29.

Stalnaker, R. 2002. Common ground. Linguistics and Philosophy, 25:701-721.

Stazicker, J. 2011. Attention, visual consciousness and indeterminacy. Mind & Language, 26(2):156-184.

Strawson, P. F. 1950. On referring. Mind, 59(235):320-344.

Tononi, G. 2004. An information integration theory of consciousness. BMC Neuroscience, 5(42).

Tononi, G. 2008. Consciousness as integrated information: a provisional manifesto. Biological Bulletin, 215:216–242.

Tye, M. 2009. A new look at the speckled hen. Analysis, 69(2):258-263.

Tye, M. 2010. Up close with the speckled hen. Analysis, 70(2):283-286.

Williamson, T. 1990. Identity and Discrimination. Oxford: Blackwell.

Williamson, T. 2002. Knowledge and Its Limits. Oxford: Oxford University Press.

Wolff, C. 1751/1983. Der Vernünfftigen Gedancken von Gott, der Welt und der Seele des Menschen, auch allen Dingen überhaupt anderer Theil, bestehend in ausführlichen Anmerckungen. Hildesheim: Georg Olms Verlag.

Wollheim, R. 1987. Painting as an Art. Thames & Hudson Ltd.