



Neural phase: a new problem for the modal account of epistemic luck

Adam Michael Bricker¹ 

Received: 1 August 2019 / Accepted: 10 December 2019 / Published online: 16 December 2019
© The Author(s) 2019

Abstract

One of the most widely recognised intuitions about knowledge is that knowing precludes believing truly as a matter of luck. On Pritchard’s highly influential modal account of epistemic luck, luckily true beliefs are, roughly, those for which there are many close possible worlds in which the same belief formed in the same way is false. My aim is to introduce a new challenge to this account. Starting from the observation—as documented by a number of recent EEG studies—that our capacity to detect visual stimuli fluctuates with the phase of our neural oscillations, I argue that there can be very close possible worlds in which an actual-world detectable stimulus is undetectable. However, this doesn’t diminish our willingness to attribute knowledge in the case that the stimulus is detectable, even when undetectability would result in the same belief formed in the same way being false. As I will argue at length, the modal account appears unable to accommodate this result.

Keywords Veritic luck · Knowledge · Epistemology · EEG · Visual detection · Perception

1 Introduction

The purpose of this paper is to introduce a specific class of counterexample to the modal account of epistemic luck. While still categorically stipulative, this type of counterexample originates from the somewhat recent discovery in cognitive neuroscience that visual detection thresholds vary systematically with regular, rapid neural oscillations. That is, the intensity required for certain visual stimuli to be detectable is not fixed, even probabilistically, but rather is constantly fluctuating. Using this observation as a starting point, I argue that this allows us to describe cases in which S’s visual detection threshold, while low enough to facilitate conscious perception of a given stimulus in the actual world, is insufficiently low for conscious perception

✉ Adam Michael Bricker
adam.michael.bricker@gmail.com

¹ History of Sciences and Ideas, University of Oulu, Oulu, Finland

in a wide range of close possible worlds. Crucially, however, it seems that we are still willing to say that, when S's perceptual threshold is low, her belief that no stimulus was presented can constitute knowledge, despite there being many close possible worlds—those in which her perceptual threshold was too high—in which she forms this same belief in the same way, and it is false. As the modal account of epistemic luck is inconsistent with such knowledge, we can understand the nature of the challenge posed: The modal account appears to label certain instances of genuine knowledge as unknown, luckily true belief.

The seriousness of this challenge will hinge on the plausibility that (1) the beliefs in question actually constitute knowledge and (2) the modal account in fact labels such beliefs as luckily true. Over the course of this paper, I hope to make clear that we have good reason to suppose both (1) and (2), and therefore that the modal account indeed gets these neural phase cases wrong. Crucially, the point here isn't that there is no species of epistemic luck incompatible with knowledge,¹ but rather that the modal account sometimes struggles to distinguish between knowledge-compatible and knowledge-incompatible varieties of epistemic luck.

In arguing along these lines, this paper is ordered in the following way: First, I provide a brief introduction to epistemic luck generally and the modal account specifically (Sect. 2). Next, I introduce the case of neural phase and provide an initial assessment of how it might pose a problem for the modal account of epistemic luck (Sect. 3). After this, I provide a more thorough argument that the neural phase case indeed poses a serious challenge for the modal account (Sect. 4). Finally, I offer something of a diagnosis of where the modal account goes astray (Sect. 5) and discuss a few potential objections (Sect. 6).

2 The modal account of epistemic luck

The purpose of this section is to provide a brief overview of the modal account of epistemic luck. In doing so, I'll first begin with an example of the sort of belief we judge to not constitute knowledge because it is true as a matter of luck. After this, I'll provide an overview of Pritchard's modal account of this epistemic luck, along with an assortment of preliminary distinctions. Consider the following:

The Night Train

The train is scheduled to pass through at 9:43 every night, but occasionally it is late. Nightly Alex watches from her window as it cuts across the countryside, its lights piercing the darkness, too distant to make a sound. It is 9:43 now. Looking to the window, Alex sees nothing, and on this basis comes to believe that the train is running late tonight. As it happens, Alex is correct. However, unbeknownst to her, the heavy screen Alex installed in the window earlier that day reduces the intensity of incoming light, so that the light of the train would

¹ To use Pritchard's parlance, this isn't an argument against an "anti-luck condition" on knowledge.

be below her perceptual threshold. Even if the train were on time, tonight she couldn't have seen it through the window.

In this case, we judge that Alex doesn't know that the train was late, despite believing truly that it was. The reason Alex's belief fails to constitute knowledge is simple: It was true as a matter of luck. This intuition, that beliefs true as a matter of luck cannot be known, is one of the most widely recognised in epistemology, and it is the driving force behind a multitude of epistemological counterexamples. It is the thread that ties together Gettier's (1963) original examples, Chisholm's (1966) sheep-in-the-field case, Goldman's (1976) fake barn cases, and a myriad of subsequent examples (e.g. the "Backward Clock," Williams and Sinhababu 2015). Regardless of how one might fill in the details of *what* exactly constitutes epistemic luck, it is clear that there is some species of luck that is incompatible with knowing.

The dominant epistemological account of the details of this luck is Pritchard's modal theory. Distinguishing it from "benign" varieties of epistemic luck (2005, p. 145), Pritchard categorises knowledge-undermining "veritic" luck in the following way:

[T]he key type of epistemic luck that is relevant here is that which concerns the truth of the belief in question, what we will call 'veritic' epistemic luck: It is a matter of luck that the agent's belief is true. In terms of our account of luck, this demands that the agent's belief is true in the actual world, but that in a wide class of nearby possible worlds in which the relevant initial conditions are the same as in the actual world—and this will mean, in the basic case, that the agent at the very least forms the same belief in the same way as in the actual world...—the belief is false (2005, p. 146).

Understanding veritic luck in this way, we can easily explain why Alex's belief that the train is late, although true, fails to qualify as knowledge: There is a wide class of close possible worlds in which Alex forms the same belief in the same way, but the belief is false. In this case, these are those worlds in which the train is on time. As the heavy screen impedes her ability to visually detect the train, she would fail to see it even if it rolled through precisely on time. Accordingly, she would then come to believe falsely that it was late. In this manner, the modal account is generally quite good at picking out such beliefs that are true as a matter of luck. However, over the course of this paper, I hope to demonstrate that it is not always successful in doing so.

Before moving on, there are a few important preliminary points to address. First, it is important to distinguish between the modal account of *epistemic* luck—the target of this paper—and the modal account of luck, with the latter seeking to provide a more general account of the concept of luck itself via similar modal framework (see Pritchard 2014). This distinction is key because we might still maintain that the modal account provides a good framework for understanding *epistemic* luck, even if we reject it as a general theory of luck. Indeed, we might note that key problems identified for the general modal account of luck arise from cases of non-epistemic luck (e.g. Lackey 2008; Carter and Peterson 2017). As the modal account was developed first and foremost to handle cases of *epistemic* luck, it is reasonable to

think it might do that better than accounting for all cases of luck. With this in mind, we might understand the significance of the claim that the modal account, in some cases, struggles even to capture *epistemic* luck. This is presumably what it should be best at.

Next, in fairness to Pritchard, it is worth mentioning that the above only represents the core or “canonical” (Carter and Peterson 2017, p. 2175) formulation of the modal account. Pritchard has continued to develop his modal account over a variety of dimensions (see e.g. 2014), with highlights of this development including the conceptual integration of risk and luck (2015), a distinction between intervening and environmental varieties of veritic luck (Carter and Pritchard 2015), and a refinement with respect to basing and reasons (Bondy and Pritchard 2018). However, because the core modal formulation has remained more-or-less unchanged, and this is what I’m interested in challenging here, developments beyond this core aren’t especially relevant for the purposes of this paper.

It is also important to note that the type of example I’ll discuss here—being a simple perceptual belief—falls squarely under the purview of the “basic case.” Pritchard also discusses cases in which there is more to satisfying the “relevant initial conditions” than simply being the same belief formed in the same way, for example when a belief continues to be believed for a different reason than its original formation (2005, p. 155). However, as we’ll see in the next section, the cases I’m interested in involve no such complications. Accordingly, for our purposes, the question of whether the relevant initial conditions are met reduces to the question of whether the same belief is formed according to the same belief-forming process.²

Finally, I again want to stress exactly which part of the modal formulation I intend to challenge. I do not dispute that veritic luck—the luck that one’s belief is true—is incompatible with knowledge. Rather, I will argue that the modal formulation presented by Pritchard is limited in its ability to capture this, in some instances labeling benign epistemic luck as veritic. Put another way, there can be many close possible worlds, which satisfy the relevant initial conditions, in which S’s (actual-world true) belief is false without this meaning that S’s belief is true as a matter of luck.

3 The challenge from neural phase

In this section, I introduce the type of counterexample that forms the core of my challenge to the modal account of epistemic luck. In order to do this, I will first provide a brief introduction to the empirical foundation of this challenge—the observation that visual detection thresholds fluctuate with neural phase. I will then use this as the basis for describing a case in which it seems we might judge that S can know that no stimulus was presented, despite there being many close possible worlds in which S falsely believes the same. Note that a more thorough evaluation is reserved for the next section.

² Note that I’ll return to this point in Sect. 6.




Fig. 1 A 10 s segment of EEG data displaying strong alpha-band activity (Bricker 2019). Note that each channel corresponds with a specific site on the scalp. Data has been preprocessed (re-reference to mean, 1–40 Hz band-pass filter, ICA for EOG artifacts) using MNE (Gramfort et al. 2013, 2014)

In order to understand neural phase, it might be helpful to begin with the measurement of neural activity itself. Although there are number of different ways in which neural activity might be observed, electroencephalography (EEG) is the preferred means to measure millisecond-scale activation across the entire brain. By recording changes in electrical potential between different sites on the scalp, EEG can provide measurements of brain activity with extremely high temporal resolution.³ One of the most visually striking features of an EEG recording (e.g. see Fig. 1) is the often regular, wave-like nature of the electrophysiological signal it observes. As this illustrates, one key feature of neural activity is that it displays regular, repeating fluctuations at a variety of frequencies. These fluctuations, known as “neural oscillations,” are roughly sinusoidal, with varying amplitudes on the microvolt scale (as measured at the scalp). Although they inevitably vary in amplitude and frequency composition, neural oscillations are strongest around the 10 Hz range (dubbed the “alpha band”). While the exact shape of alpha-band oscillations changes with a number of neurocognitive factors, this means that roughly every 100 ms, the alpha band completes a full oscillation.

One of the more remarkable features of neural oscillations is that our cognitive capacities can vary according to what part in the cycle (i.e. what phase) our neural oscillations happen to be. As reported by a number of recent studies (Busch et al. 2009; Mathewson et al. 2009; Sherman et al. 2016; Harris et al. 2018), our visual detection thresholds notably vary with alpha-band phase. That is, we are able to detect weaker visual stimuli when the stimulus is presented in phase with alpha oscillations (i.e. at 0) than when it is presented out of phase (i.e. π). Critically, these “periodic fluctuations of visual perceptual abilities” (Busch et al. 2009, p. 7875) mean that the same stimulus might be detectable when presented in phase

³ A representative sample rate for an EEG study is 500 Hz, and even ordinary recording equipment is capable of scanning much faster. We might contrast this with fMRI, which will typically only sample a new volume every 2.5 s (0.4 Hz) or so.

but undetectable when presented out of phase. If you like, you might think of this as meaning, to a certain approximation, that ten times a second our perceptual capacities alternate between a higher and lower detection power.

As nice as this approximation might be, unfortunately there are two complications I need to add before continuing. First, neural phase is not a binary, strictly in or strictly out, but rather continuously variable. While a stimulus might happen to present when the alpha band is at 0 or π , it might also fall at $\pi/2$ or $\pi/4$ or anything else between $-\pi$ and π . This means that there will be a continuously variable distribution of detection power as a function of phase, which peaks at 0 degrees and falls off as it approaches $\pm\pi$ (e.g. see Busch et al. 2009, p. 7873). As the Busch et al. (2009) findings illustrate, there appears to be a non-linear relationship between detection power and phase.⁴ This non-linearity is especially important, as the example I'll present exploits the characteristics of a non-linear detectability function. Next, because neural processing is inherently noisy, we cannot treat detectability as binary either. The fundamental lesson of signal detection theory is that a high detection power never *guarantees* that a stimulus will be perceived—regardless of external conditions—only a high probability of detection (for more see Winer and Snodgrass 2015). In short, the more in-phase a stimulus presentation, the more likely it is to be consciously perceived.

With all this in mind, we might now explore how these fluctuations of visual detection threshold with neural phase might pose a challenge for the modal account of epistemic luck. In order to do this, let's imagine someone for whom stimulus presentation can be systematically correlated with neural phase:

The EEG Experiment

Cameron is a participant in an EEG study. In some trials a stimulus—in this case a rapid flash of light—is presented, and in other trials no stimulus is presented. Whether or not a stimulus is presented is simply random, and Cameron is instructed to indicate whether the stimulus was presented for each trial. Stimulus intensity is held constant between trials.

For the intensity of stimulus used in the experiment, Cameron's visual system displays the following, non-linear relationship between alpha-band phase and detectability: For most of the phase range, Cameron's detection power is quite high. Her hit rates⁵ and correct-rejection rates⁶ are both greater than .95 for the vast majority of alpha phases. These are called, perhaps a bit misleadingly, "in-phase" trials. However, when alpha phase is in the $\pi \pm .005\pi$ range, her hit rate plummets to less than .05. These are dubbed the "out-of-phase" trials.

Stimulus presentation is randomized over all possible alpha-band phases, so for every trial all phases are equally probable. Accordingly, there will be far more in-phase trials than out-of-phase trials.

⁴ The Harris et al. (2018) results show something roughly similar.

⁵ The probability that she perceives a stimulus given that a stimulus was presented.

⁶ The probability that she doesn't perceive a stimulus given that no stimulus was presented.

Now let's consider an in-phase trial for which no stimulus is presented. At least initially, I think we might judge that Cameron can know that no stimulus was presented. After all, for an in-phase trial, it is well within her perceptual capacity to detect whether a stimulus has been presented.⁷ Accordingly, it doesn't seem like a matter of luck that she believes truly⁸ that no stimulus was presented. Her true belief is formed on the basis of a reliable perceptual process, exercised within its parameters for reliability. While I will discuss this matter in more detail in the next section, it is at least initially plausible that Cameron, in an in-phase trial, can know that no stimulus was presented.

Next, however, we might also note that there seem to be many close possible worlds—worlds in which Cameron forms the same belief in the same way—in which her belief is false. These are all those worlds in which a stimulus was presented out of phase. Definitively establishing modal closeness is of course notoriously tricky. However, in this case things are streamlined by two key points. First, the example is set up in such a way that Cameron's beliefs are always formed in the same way, i.e. on the basis of her ordinary perception. Second, it is posited that, for any given trial, experimental condition (in-phase vs. out-of-phase) and stimulus presentation (stimulus vs. no stimulus) are both random. While out-of-phase conditions are far less *probable* than in-phase conditions, the fact that experimental condition is randomized ensures that out-of-phase worlds are no more modally distant than in-phase worlds.⁹ In short, these stipulations seem to ensure that very little needs to change between an in-phase/no stimulus world and an out-of-phase/stimulus world. Thus, when Cameron, in an in-phase/no-stimulus world, forms the (true) belief that no stimulus was presented, there are many close possible out-of-phase/stimulus worlds in which she forms that same (false) belief in the same way. While I will also explore this point in more detail in the following section, such modal closeness between conditions seems at least initially plausible.¹⁰

In this way, we can understand how the phase-dependency of visual detection thresholds allow us to describe a problem for Pritchard's modal account of luck. On

⁷ One might worry that, because very brief stimuli presentation—on the order of tenths of milliseconds—would be required so that stimuli might fall within the out-of-phase range rather than overlapping it, it is not reasonable to stipulate that Cameron can reliably detect such stimuli in any conditions. However, even stimuli as brief in duration as .1 ms can still be “easily perceivable” given the right intensity conditions (Yeşilyurt et al. 2009, p. 450).

⁸ To be clear, it does appear that there is still some sort of luck involved, insofar as it was lucky that she *had the perceptual capacity to know*. This however isn't incompatible with knowledge. I discuss this more in section four.

⁹ One of the hallmarks of randomly possible events is that very little needs to change in the actual world for them to occur, regardless of how probable they are. The lottery is the paradigm example of this. Despite a lottery win being extremely improbable, because it can occur randomly, there are very close possible worlds in which any lottery ticket is a winner. See Pritchard (2015), §4 for a good discussion to this point.

¹⁰ Some might find it surprising that here I maintain that there are wide class of possible worlds in which Cameron's belief is false, despite it being highly probable that it is true. However, this apparent tension can be resolved by noting that because in this case modal ordering so completely decouples from probability, it's not clear we have any reason to think that a high likelihood of truth translates to truth in most possible worlds.

the modal account, Cameron's belief might be labelled as luckily true, and therefore incompatible with knowledge. However, this is at odds with our judgement that Cameron can indeed know that no stimulus was presented (for in-phase trials in which no stimulus is presented). In the next section, I will take a closer look at whether we want to say that (1) Cameron really knows and (2) the modal account really labels the beliefs as luckily true, arguing that we have compelling reason to accept both of these claims.

4 A closer look at knowing and initial conditions

In the previous section, I described a case in which it seems plausible—at least initially—that the modal account of epistemic luck mislabels an instance of perceptual knowledge as instead luckily true and therefore unknown. The extent to which this case poses a genuine challenge to the modal account hangs on two basic points: (1) whether the case described is in fact one of knowledge and (2) whether the modal account in fact labels the case described as involving knowledge-undermining veritic luck. In this section, I'll take a closer look at each question in turn, arguing that we have quite good reason to concur with the preliminary assessment that the case is indeed (1) knowledge and (2) mislabelled as luckily true by the modal account. In this manner, we might understand that the systematic variation of visual detection threshold with the phase of neural oscillation forms the basis of a serious problem for the modal account of epistemic luck.

Let's begin with the question of whether Cameron really does have knowledge in the above *EEG Experiment* case. In order to do this, it might be helpful to re-enumerate all the epistemically relevant details of her belief. For some given trial...

- (1) BELIEF CONTENT: Cameron believes that no stimulus was presented.
- (2) TRUTH VALUE: No stimulus was presented.
- (3) BELIEF-FORMING PROCESS: Cameron forms her belief according to simple, ordinary visual perception. She didn't see a stimulus, so she believes that none was presented.
- (4) ALPHA-BAND PHASE/DETECTABILITY: The phase condition is "**in phase**" (i.e. outside the $\pi \pm .005\pi$ range), so Cameron's visual detection threshold is significantly lower than the stimulus intensity if one is presented. In these conditions, there is a high ($> .95$) probability both that Cameron will perceive a stimulus when one is presented and not perceive a stimulus when none is presented.

Additionally, for all trials in the experiment...

- (5) PHASE CONDITION BASE RATES: Because all phases are equally probable, and in-phase trials occupy a much wider phase range than out-of-phase trials, the base rate of an in-phase trial is much higher than that for an out-of-phase trial. That is, for any given trial, the probability that it is in phase will be .995, and the probability that it is out of phase will be .005.

- (6) **STIMULUS PRESENTATION:** Whether a stimulus is presented is random. If a stimulus is presented, its timing is randomized within a 1 s window.¹¹

With all this in mind, we might begin to address the question of whether Cameron's belief constitutes knowledge. Immediately, we might note, per (1)–(3), that it is a true belief formed according to ordinary perception. While this itself might be insufficient for knowing—after all, ordinary perception isn't always reliable—closer inspection reveals that her ordinary perception here is reliable over two critical domains: the conditions of the specific trial and the conditions of the overall experiment. First, for all in-phase trials, Cameron's beliefs will form according to ordinary visual perception of an easily detectable stimulus, per (4). Thus, it shouldn't be controversial that she is employing a reliable belief-forming process for in-phase trials. Next, while the matter of reliability across the entire experiment is a bit more complicated, it is still clear that Cameron is employing a reliable belief-forming process, even when we allow for out-of-phase stimulus presentation. This is due to the fact that, per (5), the base rate for out-of-phase trials is quite low. Thus, even though Cameron is unreliable at detecting a stimulus when it is presented out of phase, that is sufficiently rare an occurrence that Cameron's simple strategy of forming beliefs based on her perceptions is still highly reliable for the experiment generally. In short, when Cameron believes that no stimulus was presented, she does so on the basis of her highly reliable, ordinary visual perception. This is a strong confirmation that her belief does in fact constitute knowledge.¹²

To stress this point a bit further, it might be helpful to do away with the neurocognitive framing. While this framing might be necessary to pose the challenge to the modal account of epistemic luck, its unfamiliarity and technicality is likely something of an impediment for understanding why Cameron has knowledge. Accordingly, let's consider a parallel retelling of the *EEG Experiment* example, but along the lines of the *Night Train* scenario presented in section one. In such a case, the in-phase condition is comparable to no screen being in the window, and an out-of-phase condition is comparable to the screen being in the window. The stimulus now is the train itself, which we might imagine is randomly late. And finally, in a slight departure from what was presented in section one, we might posit that Alex isn't responsible for placing the screen on the window, but instead some third party will randomly install it on the window (and subsequently remove it), so that for any given night there is a .5% chance that the screen is on the window. Put in this way, it seems something of a stretch to deny that, when there is no screen in the window, Alex can know that there is no train. While we should be wary of making any conclusions

¹¹ I should note that neither of these are invented for the purpose of ticking the necessary theoretical boxes, but are instead are quite common in neurocognitive research. The latter, known as "jittering," is especially prevalent in fMRI studies (less so for EEG).

¹² There is a worry that while her belief is reliable, it isn't safe, which might indicate she actually doesn't know. However, as I discuss more in Sect. 6, due to the theoretical proximity of safety to the modal account of epistemic luck, a shortcoming in the latter is likely to also be present in the former. Thus, it would be a mistake to employ safety in the defense of the modal account of luck as if it's some neutral arbiter.

solely on the basis of such a “reframing,” it is telling that when we reframe the *EEG Experiment* example in this way, presenting the same epistemically relevant features in more accessible terms, we clearly judge that such features result in knowledge. When taken together with the above considerations of reliability, we have strong reason to conclude that Cameron indeed knows that no stimulus was presented, given (1)–(5).

I now want to turn to the question of whether the modal account indeed labels Cameron’s belief, as described per (1)–(6), as a case of veritic luck. Here again I think we have strong reason to suppose that it does, on the basis that we might describe the requisite nearby possible worlds in which Cameron’s same belief, formed in the same way, is false. These, I maintain, are all those possible worlds in which we hold (1), (3), (5), and (6) fixed, but instead replace (2) and (4) with the following:

- (2*) TRUTH VALUE: A stimulus was presented.
- (4*) ALPHA-BAND PHASE/DETECTABILITY: The stimulus was presented “**out of phase**” (i.e. within the $\pi \pm .005\pi$ range), so Cameron’s visual detection threshold is significantly higher than the stimulus intensity. In these conditions, there is a low ($< .05$) probability that Cameron will perceive a stimulus when one is presented.

That is, all possible worlds in which Cameron was presented with a stimulus out-of-phase with her alpha-band oscillations, and accordingly fails to perceive the stimulus, are close possible worlds in which she forms—via the same belief forming process described in (3)—the false belief that no stimulus was presented. First, as the example is designed so that in-phase/no-stimulus worlds are very close to out-of-phase/stimulus worlds, this part shouldn’t be especially controversial. Per (6), both whether a stimulus is presented and what phase condition that stimulus falls under, if presented, will be random for every trial. As mentioned above, because very little needs to change between worlds for random events to occur or not occur, this means that in-phase/no-stimulus and out-of-phase/stimulus worlds will be very close. Next, moving to initial conditions, the example is also designed so that it isn’t controversial that in the vast majority of out-of-phase/stimulus worlds, Cameron will form the same belief as in the in-phase/no-stimulus worlds, i.e. that no stimulus was presented, and thus believe falsely. After all, the vast majority of the time that a stimulus is presented in phase, Cameron won’t perceive it.

Finally, there is the question of the other relevant initial condition, the belief-forming process employed in out-of-phase/stimulus worlds, and whether it is the same as that employed in in-phase/no-stimulus worlds. This is a bit trickier. Although initially we might observe that Cameron is simply using ordinary visual perception in both out-of-phase/stimulus worlds and in-phase/no-stimulus worlds, this meets with the following objection: Cameron’s neural systems are in different phase configurations for in-phase and out-of-phase conditions. Therefore, some might be tempted to just say that we might individuate belief-forming processes according to neural phase in order to avoid any challenge to the modal account. The problem with such a proposal, however, is that neural phase is far too fine-grained

to serve as a reasonable basis for the individuation of belief-forming processes. Not only does it change rapidly, but it is constantly fluctuating in a number of different frequency bands, including the alpha (around 10 Hz), but also the beta (around 15–30 Hz), gamma (around 40 Hz), and a number of others. Even assuming we might simplify the phase-states of each band to two (in or out), for a total of 8 possible alpha–beta–gamma phase configurations, individuating according to these would have two obviously suspect implications: (1) Our latent perceptual belief-forming processes are never stable, and instead fluctuate tens—perhaps hundreds—of times a second. Accordingly, (2) most belief-forming processes previously considered identical due to cognitive-level features would in fact be distinct, and determining whether two different epistemic agents, or one agent at two different points in time, formed a given belief or beliefs according to the same process is impossible without appeal to neural measures (i.e. EEG).

These implications would require a fundamental break with the way we talk and think about belief-forming processes. Let's consider a simple example of ordinary visual perception: I'm looking for cheese in the fridge. I first check the top shelf. I don't see any cheese, so I come to believe (b1) that there is no cheese on the top shelf. I then check the middle shelf. I again don't see any cheese, so I come to believe (b2) that there is no cheese on the middle shelf. Provided I do so under ordinary conditions, we clearly judge that I use the same belief-forming process in forming b1 and b2. Suggesting that I likely used different processes seems to fundamentally miss what we have in mind when we talk about belief-forming processes. Further, imagine now that my partner looks for cheese on the bottom shelf. Because she doesn't find any, she comes to believe (b3) that there is no cheese on the bottom shelf. Again, while we judge that she uses the same process to form b3 as I did to form b1 and b2, if we individuate according to neural phase, this would likely not be the case. Accordingly, because it would require a non-trivial re-evaluation the very concept of the belief-forming process, I don't think that individuating belief formation according to neural phase is an attractive option.

Next, let's consider a second strategy for maintaining that different belief-forming processes are employed in out-of-phase/stimulus worlds and in-phase/no-stimulus worlds: individuating via reliability. That is, if we say that the belief-forming processes employed by Cameron in these two conditions are distinct *in virtue of the process being unreliable in out-of-phase worlds and reliable in in-phase worlds*, this might allow us to circumvent the challenge to the modal account. The problem with such a proposal is that it too would require a significant reevaluation of the concept of the belief-forming process. Philosophers tend to think of the same belief-forming processes as being capable of assuming, for example, different degrees of reliability at different points in time (see e.g. Frise 2018; Tolly 2018). Thus, individuation via reliability simply presents as implausible.

Finally, one might try individuating via *cognitive capacity*.¹³ Because the belief-forming processes employed by Cameron in out-of-phase/stimulus worlds and in-phase/no-stimulus worlds are operating at different cognitive capacities, specifically visual detection thresholds, those processes are distinct, or so the objection goes. I will grant that, on its surface, this doesn't seem like the worst way to individuate belief-forming processes, and there might even be a certain intuitive appeal to this proposal. Indeed, as I'll discuss more in the next section, I think the reason that a counterexample like *EEG Experiment* hasn't been raised sooner is that we don't intuitively think that cognitive capacities can vary within a single belief-forming process. However, the problem with using cognitive capacity to individuate belief-forming processes is that doing so causes its own serious problem for the modal account of epistemic luck, notably in a much more straightforward way than allowing for the converse. In order to illustrate this problem, consider the following sheep-in-the-field-esque case:

Greg in the Field

Greg is standing in a field about 50 meters away from a sheep-sized dog, immediately behind which is a sheep. From Greg's perspective, the dog's body perfectly occludes the sheep. Ordinarily, Greg's perceptual capacity to discriminate a sheep from a sheep-sized dog at 50 m is quite low—say a hit rate and correct-rejection rate both around .55. However, a curse has been put on Greg such that, when he is in this exact field/sheep/sheep-sized dog configuration, his discriminatory capacity is much stronger—say a hit rate and correct-rejection rate both in excess of .99.

Thus, when Greg looks at the dog in his present situation, he has an accurate perceptual capacity to discriminate it from a sheep. However, due to neural noise (i.e. the < 1% chance of a miss), in this case Greg happens to misperceive the dog to be a sheep, despite the accuracy of his perception. On this basis, he comes to believe that there is a sheep in the field.

In this case, clearly Greg doesn't know that there is a sheep in the field. His belief is true as a matter of luck. However, if we individuate Greg's belief-forming processes based on cognitive capacity, this means that how Greg forms a belief in his present configuration will be different than how he will form it when there isn't a sheep hiding behind the dog. Crucially, this results in there being no close possible worlds in which he falsely believes that there is a sheep in the field *according to the same process in the example*. Thus, his belief doesn't satisfy the requirements for veritic luck according to the modal account. This of course isn't a problem if we just say that all those close possible worlds in which Greg falsely believes there to be a sheep in the field—still inferred from the same misperception, but now with a much lower visual

¹³ Note that here I'm using "cognitive capacity" here to mean something like "the processing capacity of a cognitive system." On this usage, visual perception isn't a cognitive capacity as such, but rather a cognitive system that has different processing capacities (accuracy, latency, etc.) for different individuals and contexts.

capacity—employ the same belief-forming process as the example world. Thus, it appears that individuating belief-forming processes via cognitive capacity poses its own problem for the modal account.

In conclusion, because there appears to be no plausible means by which we might argue that Cameron is employing different belief-forming processes in the different phase conditions, we have compelling reason to accept our initial assessment that she is using the same process across all conditions—simple, ordinary visual perception. Thus, the modal account of epistemic luck labels her in-phase/no-stimulus beliefs as true as a matter of luck and therefore unknown. However, as we also have compelling reason to concur with our initial judgement that these beliefs constitute knowledge, we can understand how this case exposes a serious shortcoming in the modal account of epistemic luck.

5 Diagnosing the problem

In the preceding sections, I argued that the fluctuation of visual detection threshold with neural phase poses a serious challenge to the modal account of epistemic luck. While the purpose of this paper is not to offer a positive account of epistemic luck, in this section I will take a preliminary step in this direction by offering some like a diagnosis of the modal account's shortcomings in the case of neural phase. In brief, the modal account confuses the capacity luck (knowledge-compatible) at work in the *EEG Experiment* case with veritic luck (knowledge-incompatible), likely due to an implicit assumption regarding the sorts of changes in cognitive capacity that can easily occur.

To begin, I want to start with the observation that, clearly, there is some variety of epistemic luck at play in the *EEG Experiment* example. I maintain that we might easily understand *capacity* luck—the luck that someone is capable of gaining knowledge (see Pritchard 2005, p. 134)—to be the operative variety of epistemic luck in the Cameron case. When Cameron forms her beliefs on the basis of a reliable perceptual capacity (i.e. in in-phase trials), it isn't a matter of luck that her belief is true. However, it is a matter of luck that she has that reliable perceptual capacity in the first place. Had she instead been in an out-of-phase trial, an easy possibility, she wouldn't have had the perceptual capacity to reliably detect the target stimulus. This, I think, is the natural way in which to understand our intuition that there is some luck involved in when Cameron forms true beliefs in in-phase trials—she is lucky to have the perceptual capacity to reliably form those true beliefs in the first place. Given that capacity luck is uncontroversially compatible with knowing, we might conceptualize the modal account's struggle with the case of neural phase as the misidentification of capacity luck as veritic luck.

I now want to say something about *why* the modal account confuses capacity luck for veritic luck in the *EEG Experiment* case. I'm not sure that the answer to this is especially profound, but it's worth highlighting anyway. The problem with the

modal account of epistemic luck is just that it rests on the implicit assumption that there cannot be easy, endogenic¹⁴ changes in capacity to the cognitive systems we rely on for knowledge. Provided that this assumption holds, there would then never be instances of capacity luck which satisfy the “same belief-forming process” initial condition, and thus there would be no risk of the modal account labelling capacity luck as veritic luck.

As indicated above, this seems like quite an intuitive assumption. Indeed, I think that part of what makes the empirical findings discussed in Sect. 3 so interesting is that we of course don't intuitively think that our capacity to detect visual stimuli is rapidly fluctuating. The problem is that, once we recognize that cognitive systems like visual perception can and do in fact undergo easy, endogenic, fine-grained changes, the modal account is faced with a dilemma. On the one hand, maintaining that such changes cannot occur within a single belief-forming process leads to the problem highlighted by the above *Greg in the Field* example—obvious instances of veritic luck will no longer be labelled as such by the modal account. However, allowing for changes in cognitive capacity within the same belief-forming process results in problem highlighted by the *EEG Experiment* case—instances of knowledge involving capacity luck will mistakenly be labelled as veritic luck. Thus, I think the best diagnosis I can offer here is just that the modal account of epistemic luck made the mistake of implicitly assuming that there cannot be these specific sorts of changes in knowledge-generating cognitive capacities.

6 Three further worries: safety, initial conditions, and the anti-luck intuition

To close out this paper, I want to address three additional worries that might feasibly be raised against the preceding account: (1) As Cameron's beliefs in the *EEG Experiment* case will never be safe, might we not use this to as a basis to deny that they can constitute knowledge? (2) Might we not add some further initial conditions to the modal account so that it can accommodate neural phase cases? (3) Rather than reject the modal account, why not instead say that we were mistaken to take all veritic luck to be knowledge-incompatible? I'll briefly discuss each objection in turn, arguing that none of these concerns significantly threaten my argument.

When attempting to determine whether a tricky case of true belief actually constitutes knowledge, it is often helpful to appeal to some epistemic condition that generally covaries with knowledge. This is exactly what I did in section three with reliability. By demonstrating that Cameron's belief-forming process was reliable, I was able to bolster the claim that her true beliefs indeed constitute knowledge. Beyond reliability, the other conventional choice for such a condition would have been safety—roughly, S knows that p only if S's belief that p couldn't have easily been

¹⁴ I.e. arising from internal, biological processes. Of course, there can be easy exogenic changes to perceptual capacity, e.g. those caused by psychotropics or TBI, but we seem much less willing to say that belief-forming processes are preserved across such changes.

false. However, had I opted for safety instead of reliability, I would have run into a problem. As mentioned at the time, Cameron's beliefs will never be safe. Understanding "p couldn't have been easily false" in the usual "there are no close possible worlds in which p is false" way, the example is designed so that, when Cameron believes that p truly, she always could have easily believed that p falsely. Accordingly, one might then object to my account that, because her beliefs will never be safe, this counts against the assessment that these beliefs can constitute knowledge, at a minimum cancelling out any support offered by reliability.

In many cases, this sort of objection could present a serious challenge. However, here we might note that, unlike reliability, safety does not represent an independent gauge by which to assess the modal account of epistemic luck. Instead, the safety condition and the modal account are inexorably linked, and Pritchard makes clear in *Epistemic Luck* that he understands the safety condition to "eliminate veritic epistemic luck" (2005, p. 145). Accordingly, it simply comes with the territory that a counterexample to the modal account of epistemic luck will likely be incompatible with safety, so we cannot treat a violation of safety as an independent problem for my argument.

Next, one might wonder whether there is some initial condition we could add to the modal account, beyond the same belief and belief-forming process, which might allow it to accommodate cases involving neural phase. The most obvious suggestions would be either (1) same neural phase or (2) same relevant cognitive capacity. However, the problem with these is that if we add either as general initial conditions for veritic luck, the result is an account of luck that is far too strict. Let's first consider that for S's belief that p to be true as a matter of luck, this requires that there are a wide class of worlds in which S's belief that p, formed according to the same process and *in the same neural phase*, is false. But clearly this doesn't capture the full range of scenarios in which we might say someone's belief is true as a matter of luck. Imagine that we describe the initial *Night Train* example from section one such that Alex's true belief that the train is late was formed in some phase configuration C, but that for some reason this phase configuration is tied to her belief being true—she will never be in C when she believes falsely.¹⁵ Accordingly, there will not be any close possible worlds in which she believes the same belief falsely *in phase C*, and thus her belief would no longer qualify as veritic luck under the present proposal. Nevertheless, because her belief in the actual world was clearly still true as a matter of luck, we can easily observe that neural phase doesn't function as a viable initial condition.

We might then say something similar if we take relevant cognitive capacity as an initial condition: For S's belief that p to be true as a matter of luck, this requires that there are a wide class of worlds in which S's belief that p, formed according to the same process *operating at the same capacity*, is false. In order to see how this also results in too strict an account of epistemic luck, let's return to the *Greg in the Field* example from Sect. 4, with an additional stipulation:

¹⁵ It doesn't matter why. Say a wizard did it.

Greg's curse can only be bestowed given that a confluence of events, none of which can easily occur, occur. The point is, there are no close possible worlds in which Greg doesn't have this curse.

Recall now that, in this example, it is clear that Greg's belief is true as a matter of luck. Moreover, while there will be close possible worlds in which he forms this same belief in the same way, and is false, in all of these worlds Greg will happen to have a much lower perceptual capacity. In this way, we can understand that adding "same relevant cognitive capacity" to the modal account results in something that excludes obvious cases of veritic luck. Accordingly, it doesn't seem like adding further initial conditions is an especially viable alternative to my argument. To be fair, I cannot rule out the possibility that perhaps a more sophisticated initial condition might be capable of doing work the above cannot. However, until such a condition is described, I think that we can put this worry on the backburner.

Finally, one might wonder why I have decided to use the peculiarities of neural phase to motivate a case against the modal account of epistemic luck, rather than to argue that not all veritic luck is incompatible with knowing. My thinking here is that it seems clear that the operative luck in the Cameron case is capacity luck, and not veritic luck. However, I do recognize that if someone were to disagree with this assessment, then the natural conclusion is that some veritic luck is compatible with knowledge.¹⁶ Nevertheless, this would of course still constitute a challenge to the modal account, which at its core seeks to isolate a knowledge-incompatible species of epistemic luck.

7 Conclusion

In this paper, I described a new problem for the modal account of epistemic luck, which emerged from recent empirical findings from cognitive neuroscience. Because the human capacity to visually detect stimuli fluctuates with neural phase, we can describe cases in which a low visual detection threshold could have easily been much higher. While this doesn't mean that we cannot gain knowledge in low-threshold conditions, the modal account responds as if it does, mislabelling this capacity luck as veritic luck. This error on the part of the modal account seems to derive from an implicit assumption that there cannot be this sort of easy change in the capacities of the cognitive systems we rely on for knowledge.

Beyond anything specific about the limitations of the modal account of epistemic luck, there is a more general observation to be made about the relationship between epistemological theories and the empirical sciences—especially cognitive neuroscience. The modal account of luck, as is standard practice in epistemology, is designed to handle human cognitive processes *as imagined by epistemologists*. This means that, when the actual workings of human cognition inevitably diverge, often in strange and

¹⁶ In fact, you might even take this a step further and, following Carter and Pritchard's distinction between intervening and environmental varieties of veritic luck (2015), say that environmental—but not intervening—veritic luck is compatible with knowledge.

surprising ways, from the epistemological imagination, epistemological theories will likely be ill-equipped to handle these divergences. This perhaps is the more important lesson of this paper—if we wish for epistemology to produce theories capable of accommodating the peculiarities of human cognitive systems and the neural processes they supervene upon, then it is a mistake to isolate epistemological production from cognitive neuroscience.

Acknowledgements Open access funding provided by University of Oulu including Oulu University Hospital. I am singularly grateful to Henry Railo for pointing me in the direction of this neural phase phenomenon, if not for which this paper wouldn't exist.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bondy, P., & Pritchard, D. (2018). Propositional epistemic luck, epistemic risk, and epistemic justification. *Synthese*, 195(9), 3811–3820.
- Bricker, A. M. (2019). Unpublished EEG data.
- Busch, N., Dubois, J., & Vanrullen, R. (2009). The phase of ongoing EEG oscillations predicts visual perception. *The Journal of Neuroscience*, 29(24), 7869–7876.
- Carter, J., & Peterson, A. (2017). The modal account of luck revisited. *Synthese*, 194(6), 2175–2184.
- Carter, J., & Pritchard, D. (2015). Knowledge-how and epistemic luck. *Noûs*, 49(3), 440–453.
- Chisholm, R. (1966). *Theory of knowledge*. Englewood Cliffs, NJ: Prentice-Hall.
- Frise, M. (2018). The reliability problem for reliabilism. *Philosophical Studies*, 175(4), 923–945.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23(6), 121–123.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy*, 73(20), 771–791.
- Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., et al. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, 7(7), 267.
- Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., et al. (2014). MNE software for processing MEG and EEG data. *NeuroImage*, 86, 446–460.
- Harris, A., Dux, P., & Mattingley, J. (2018). Detecting unattended stimuli depends on the phase of prestimulus neural oscillations. *The Journal of Neuroscience*, 38(12), 3092–3101.
- Lackey, J. (2008). What luck is not. *Australasian Journal of Philosophy*, 86(2), 255–267.
- Mathewson, K., Gratton, G., Fabiani, M., Beck, D., & Ro, T. (2009). To see or not to see: Prestimulus alpha phase predicts visual awareness. *The Journal of Neuroscience*, 29(9), 2725–2732.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Clarendon.
- Pritchard, D. (2014). The modal account of luck. *Metaphilosophy*, 45(4–5), 594–619.
- Pritchard, D. (2015). Risk. *Metaphilosophy*, 46(3), 436–461.
- Sherman, S., Kanai, R., Seth, A., & VanRullen, R. (2016). Rhythmic influence of top-down perceptual priors in the phase of prestimulus occipital alpha oscillations. *Journal of Cognitive Neuroscience*, 28(9), 318–330.
- Tolly, J. (2018). Does reliabilism have a temporality problem? *Philosophical Studies*, 176(8), 1–18.
- Williams, J., & Sinhababu, N. (2015). The backward clock, truth-tracking, and safety. *Journal of Philosophy*, 112(1), 46–55.
- Winer, E., & Snodgrass, M. (2015). Signal Detection Theory. In M. Matthen (Ed.), *The Oxford handbook of philosophy of perception*. Oxford: Oxford University Press.

Yeşilyurt, B., Whittingstall, K., Uğurbil, K., Logothetis, N., & Uludağ, K. (2009). Relationship of the BOLD signal with VEP for ultrashort duration visual stimuli (0.1 to 5 ms) in humans. *Journal of Cerebral Blood Flow and Metabolism*, 30(2), 449–458.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.