



Safety and dream scepticism in Sosa's epistemology

J. Adam Carter¹ · Robert Cowan¹

Received: 12 November 2023 / Accepted: 19 March 2024
© The Author(s) 2024

Abstract

A common objection to Sosa's epistemology is that it countenances, in an objectionable way, *unsafe knowledge*. This objection, under closer inspection, turns out to be in far worse shape than Sosa's critics have realised. Sosa and his defenders have offered two central response types to the idea that allowing unsafe knowledge is problematic: one response type adverts to the *animal/reflective* knowledge distinction that is characteristic of bi-level virtue epistemology. The other less-discussed response type appeals to the threat of *dream scepticism*, and in particular, to the idea that many of our everyday perceptual beliefs are unsafe through the nearness of the dream possibility. The latter *dreaming response* to the safety objection to Sosa's virtue epistemology has largely flown under the radar in contemporary discussions of safety and knowledge. We think that, suitably articulated in view of research in the philosophy and science of dreaming, it has much more going for it than has been appreciated. This paper further develops, beyond what Sosa does himself, the dreaming argument in response to those who think safety (as traditionally understood) is a condition on knowledge and who object to Sosa's account on the grounds that it fails this condition. The payoffs of further developing this argument will be not only a better understanding of the importance of insights about dreaming against safety as a condition on knowledge, but also some reason to think a weaker safety condition, one that is relativised to SSS (i.e., skill/shape/situation) conditions for competence exercise, gets better results all things considered as an anti-luck codicil on knowledge.

Keywords Ernest Sosa · Scepticism · Virtue Epistemology · Safety · Dreams

✉ J. Adam Carter
adam.carter@glasgow.ac.uk
Robert Cowan
robert.cowan@glasgow.ac.uk

¹ University of Glasgow, Glasgow, UK

1 Introduction

When epistemologists think of Sosa's virtue epistemology in connection with dreaming and dream scepticism, they will likely have in mind what has been a much-discussed line of argument that Sosa (2007, Ch. 1) advances in Ch. 1 of his *A Virtue Epistemology*, and which exploits a nonstandard approach to the nature of dreams – viz., the *imagination model of dreaming*¹ – in an attempt to defuse the sceptical import of *dream scepticism*.²

To bring this line of argument into focus, compare Descartes' two sceptical hypotheses introduced in *Meditation One*: the hypothesis that I am now dreaming and so forming perceptual beliefs that are false rather than true (*the dream sceptical hypothesis*); and the hypothesis that I am presently the victim of an elaborate illusion orchestrated by an evil genius, also forming beliefs that are false rather than true (*the evil genius hypothesis*). Which scenario should give us the most concern in so far as we hope to avoid making any substantial concessions to the sceptic?

We might be tempted to think, as Descartes encouraged us to, that it's the latter. After all, if the dreaming hypothesis obtained while you thought you were looking at a fire, then this would seem to imperil at most (would-be) items of perceptual knowledge (e.g., if you believed you were looking at a fire, you'd be wrong) but it wouldn't (not obviously, anyway) imperil any logical or mathematical beliefs you might form while dreaming³; whereas, if the evil genius hypothesis obtained, then (*ex hypothesi*) even logical and mathematical propositions might *seem* necessarily true, clearly and distinctly so, when they were false, something not typical of even our most bizarre dreams. What is more, as Descartes (in *Meditation Six*) as well as Austin (1962)⁴ had thought, there would plausibly be various tell-tale signs that we could exploit while dreaming to tell we are dreaming, but there would be *ex hypothesi* be no such signs to exploit were the evil genius hypothesis to obtain.

Even so, a point of widespread agreement is that the *remoteness* of the obtaining of a given sceptical hypothesis is an important factor that bears, *ceteris paribus*, on its capacity to imperil the knowledge we have in the actual world.⁵ Compare (i) a belief that *could* have been incorrect with (ii) a belief that *could easily* have been incorrect. The fact that the evil genius sceptical hypothesis is such that, if it obtained, many of our beliefs would be false, gives us *at most* reason to think, of any given belief

¹ Defenders include, e.g., Ichikawa (2009); McGinn (2004); Thompson (2014).

² This argument first appears in his "Dreams and Philosophy" (2005), reprinted as Chap. 1 in *A Virtue Epistemology: Volume 1* (2007). For some representative critical discussion of this line of argument in epistemology, see e.g., Ballantyne and Evans (2010); Brown (2009; Ichikawa, (2009, 2008); Bueno (2009) and Carter (Forthcoming).

³ As Descartes puts it: "For whether I am awake or asleep, two and three added together are five, and a square has no more than four sides. It seems impossible that such transparent truths should incur any suspicion of being false." (Descartes [1641] 1984).

⁴ Consider Austin's (1962) example: "I may have the experience [...] of dreaming that I am being presented to the Pope. Could it be seriously suggested that having this dream is 'qualitatively indistinguishable' from actually being presented to the Pope? Quite obviously not. (1962, 48)

⁵ This idea has gained support from a range of different camps, including from contextualist epistemologists (DeRose, 1992) including those who embrace a 'relevant alternatives' contextualism (Lewis, 1996), as well as from invariantist epistemologists who accept forms of fallibilism (e.g., Brown, 2018).

we have, that (i) it *could* be incorrect (viz., were that hypothesis to obtain, the belief would be false), but *not* (ii) that it could *easily* be incorrect.

By contrast – and here is where the dreaming sceptical hypothesis is of special interest – the fact that *that* hypothesis is such that, if it obtained (right now, as we form perceptual beliefs), many of our beliefs would be false, plausibly *does* give us reason to think *not only* that these beliefs (i) *could* be incorrect, but that they (ii) *could easily* be incorrect. After all, we already know that sleeping and dreaming are ordinary features of the actual world. Dreaming is something normally functioning brains do all the time, sometimes for a given individual multiple times a day. The very idea that we are dreaming right now doesn't look at all like a *remote* possibility; that is, it's *not* something that would be happening right now only if the world was much different from ours. Here's Sosa (2007):

[...] the dream possibility is *too close for comfort*. If while dreaming we have real beliefs based on real phenomenal experiences, then a normal perceptual judgment could always be matched by a subjectively similar, similarly based judgment, made while one dreams. Too easily, then, we might right now be dreaming when we form perceptual beliefs (2007, 3, our italics).

Setting aside for now just how close is 'too close' – we'll return to this – Sosa's most influential theoretical innovation related to dreaming was to show that *even once the* dreaming hypothesis is granted to be a *non-remote* hypothesis, it would have 'sceptical teeth' at all only if we were, in fact, *forming false beliefs while* dreaming – a scenario that itself is predicated upon the Orthodox Theory of Dreaming:

ORTHODOX THEORY OF DREAMING: Dreams are partly composed of occurrent beliefs (typically false) formed by the sleeping subject.

A direct entailment of the Orthodox Theory is that subjects form beliefs while they are dreaming. But Sosa takes it that we have good independent reasons to think that the Orthodox Theory of Dreaming is mistaken, and that we should accept instead the Imagination Theory of Dreaming:

IMAGINATION THEORY OF DREAMING: Dreams are partly composed of cognitive imaginings formed by the sleeping subject.⁶

Sosa takes this claim to entail that, while dreaming, subjects *do not* form beliefs any more than they climb mountains. While dreaming, we merely *imagine* that we climb mountains, and we likewise merely imagine, while dreaming, that we form beliefs.

This point matters (*vis-à-vis* dream scepticism) because if Sosa is right and the Imagination Theory of Dreaming is correct, then it looks like no matter how eas-

⁶ Note that this formulation of the Imagination Theory of Dreaming lines up with the kind of position Sosa appealed to under this description, and which takes just the cognitive aspects of dreams to be imaginative. Later, we consider a more robust version of the thesis which allows more aspects of dreaming to be imaginative.

ily we might be dreaming, it will never be the case that *too easily, we might right now be dreaming when we form perceptual beliefs*. If, necessarily, we're not forming beliefs while dreaming (as per the Imagination Theory of Dreaming), then it follows we're not forming false beliefs while dreaming; dreaming possibilities are thus misdescribed as *error possibilities*. Armed with the Imagination Theory of Dreaming, we have straightforward ammunition against dream scepticism.

Importantly for what follows, even though Sosa takes it that the above line of argument suffices to defuse the sceptical import of the dreaming sceptical hypothesis, he doesn't want to stake too much on it. The Imagination Theory of Dreaming is, after all, controversial⁷; Sosa wants to show how his preferred virtue theoretic account of knowledge could in principle be vindicated against the threat of dream scepticism *even if we assume the Orthodox Theory of Dreaming* he rejects, where the modal nearness of the dreaming hypothesis then becomes a more serious epistemological challenge.

2 Sosa's dreaming reduction against safety

As we will later explain, we think Sosa's appeal to the Imagination Theory of Dreaming (at least, on the most plausible characterisation of what this theory is committed to, in substantive detail) anyway doesn't work.

For now, though, we want to zero in on the line of argument that will be our central reference point in the remainder of what follows: this is Sosa's *Dreaming Reductio Against Safety* as a necessary condition on knowledge, which is an argument that sets aside entirely any reliance on the Imagination Theory of Dreaming.

The argument purports to show (in dilemmatic form, for now) that that the nearness of the dreaming possibility forces us to give up either (i) a suitably anti-sceptical position about the scope of our perceptual knowledge when awake; or (ii) a *safety condition* on knowledge. Since a sceptical stance about perceptual knowledge carries an enormous theoretical cost, we should, as the argument holds, reject safety as necessary for knowledge.

A few brief remarks about safety are in order. First, as a condition on knowledge, safety has a lot of backers in the contemporary literature, and for a good reason.⁸ A desiderata many philosophers accept on any account of knowledge is that it rules out beliefs that are just 'luckily true'. This includes not only standard 'Gettiered' beliefs, but also beliefs that are luckily true simply on account of the agent being situated in a modally risky environment – as in BARNs:

BARNs: Using his reliable perceptual faculties, Barney forms a true belief that the object in front of him is barn. Barney is indeed looking at a barn. Unbe-

⁷ For some criticisms of the Imagination Theory of Dreaming, see, e.g., Rosen (2021) and Whiteley (2020). We take up this point in more detail in § 3.

⁸ For some representative defences of safety, see, e.g., Pritchard (2005, 2007, 2009, 2008, 2015), Williamson (2000, 147), Sosa (1999, 146), Sainsbury (1997); Luper-Foy (1984).

knownst to Barney, however, most objects that look like barns in these parts are in fact barn façades.⁹

One sure-fire way to rule out Barney as possessing knowledge is to require that known beliefs be *sensitive*¹⁰ – viz., to be such that, where the target proposition false, the subject would not continue to hold the belief. A sensitivity condition secures this result (i.e., that Barney lacks knowledge) in BARNs, though, at the cost of giving what goes on in remote, far-off worlds relevance in the actual world.¹¹ Do you right now know your family members have not been replaced by identical looking imposters?¹² Sensitivity, apart from any of its virtues or other limitations, looks too strong.

In contrast with sensitivity, a *safety condition on knowledge* can also rule out knowledge in standard Gettier cases as well as in ‘environmental luck’ cases like BARNs but – and here’s an important advantage – it can do this *without* the cost of giving remote scenarios like the imposter-switching scenario any relevance whatsoever to whether we count as knowing mundane things we take ourselves to know in the actual world. To see why this is, consider the standard formulation of a safety condition (SCK) on knowledge due to Pritchard (2007), unpacked in terms of the belief’s being modally safe (SC):

SAFETY CONDITION ON KNOWLEDGE (SCK) S’s belief that p constitutes knowledge that p only if S’s belief that p is safe (as per SC).

SAFETY CONDITION (SC): S’s belief is safe if and only if in most nearby possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, and in all very close nearby possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, the belief continues to be true.¹³

SC (as it features in SCK) is a modal representation of the idea (glossed intuitively) that a belief is safe iff it *couldn’t easily have been wrong*, given how it was formed. SCK gets the result that Barney lacks knowledge (very easily, Barney would be looking at a façade, forming a false belief – he does so in very close nearby worlds); but it gets this result *without* also implying you’d know your family’s not replaced by imposters *only if you’d not be duped if they were*. This is because when you use perception and identify your family members by sight as you normally would, it’s not the case that *very easily* you’d be wrong. In most near-by worlds in which you form

⁹ This case, originally due to Gettier (1963), was introduced widely by Goldman (1979).

¹⁰ See, e.g., Nozick (1981).

¹¹ A related issue concerns the requirement that a sensitivity theorist deny that knowledge is closed across competent deduction; see, Dretske (1970).

¹² This delusion is characteristic of Capgras Syndrome.

¹³ This formulation, from Pritchard (2007) is an amended version of his Pritchard (2005). The details won’t concern us here, but the reason for this amendment in Pritchard (2007) was to attempt to ensure that a safety condition on knowledge would be able to rule out knowledge in lottery cases without being so strong as to rule out knowledge in more mundane cases, such as Sosa’s ‘rubbish chute’ case; see also for discussion here Pritchard (2012).

your belief in the same way, the belief is true, even if it's false in far-off worlds. For *safety*, what goes on in these far-off worlds simply doesn't matter.

The capacity of SCK to handle Gettier (and other epistemic luck cases) without the modal strength of sensitivity is one reason it is a popular 'anti-luck' condition on knowledge. Controversially, Sosa thinks not only sensitivity, *but also safety* (i.e., SCK), is too strong an anti-luck condition, and that we can get by with weaker. Necessary and sufficient for knowledge for Sosa is just *apt belief* – to a first approximation for now, belief that is *accurate* (true) because *adroit* (i.e., competent) – where an (epistemic) competence is a disposition seated in an agent to reliably believe truly in contextually determined conditions.¹⁴ The requirement that for knowledge the accuracy of a belief must be because of, or creditable to, the exercise of one's epistemic *competence* is good enough to rule out knowledge in standard Gettier cases – in those case the subject gets it right but not *because* of their competence. But – and this is a central point of objection to Sosa – BARNs seems like a case that *does feature apt belief*; accordingly, the knowledge-aptness equivalence seems to carry with it the seemingly objectionable¹⁵ cost of ruling in knowledge in 'environmental luck' cases like BARNs, where the target belief is clearly *unsafe* by the lights of SC.

Question: just how bad is this result for Sosa? Here is where the dialectical situation gets thorny. Sosa has offered two entirely different lines of response to the claim that it is a theoretical cost of his view that he gives up SKC by allowing knowledge in cases like BARNs. The first, and most discussed, is a kind of *mitigating* reply – which goes as follows: *even though* the view *attributes* apt belief and thus knowledge in cases where one's belief is unsafe by the lights of SC, there is at the same time a kind of knowledge Sosa *withholds* – reflective *knowledge*, which is to be distinguished from *mere* apt belief, or 'animal knowledge'. Here's Sosa:

Animal knowledge requires that one get it right through competence rather than just luck. Reflective knowledge goes beyond that by requiring not only apt attainment of truth but also apt attainment of aptness (2021, 169).

As the thought goes, Barney believes aptly, and so has a kind of low-grade, 'animal' knowledge (which requires mere apt belief). Why do we credit him with apt belief despite the apparent luckiness in play? Because he believes truly through the exercise of a competence in appropriate conditions. This is so even though the obtaining of those conditions (he could have been situated so as to be looking at a fake) is unsafe.¹⁶ What Barney doesn't do is *aptly attain apt belief*; he thus lacks reflective knowledge. Armed with the animal/reflective distinction, Sosa accordingly has a *mitigating* response to the objection that giving up SKC is costly: "At least there is

¹⁴ This is a simple characterisation, to be superseded by a more careful discussion in, e.g., § 6.

¹⁵ Though whether this is objectionable will be a point we engage with critically shortly. It is also worth noting that results from experimental philosophy tend to support the attribution of knowledge in fake barn cases (where the unsafety of the target belief is down to one's being in a modally risky environment), even though not in standard Gettier cases. See, e.g., Colaço et al. (2014); Turri et al. (2015); Turri (2016).

¹⁶ Barney thus believes aptly (accurate because adroit) but *that Barney believes aptly* is itself lucky; that Barney is in the right conditions to exercise his barn spotting competence isn't something we credit to any competence of his to, e.g., affirm there is a barn only if he would do so not just correctly but aptly.

still a *kind of knowledge* that must be withheld in BARNS, where the subject's belief fails SC."

The matter of whether it's bad at the end of the day for Sosa to count Barney as knowing (in virtue of having mere apt belief) in cases where safety fails has usually been adjudicated as a kind of metaepistemological dispute¹⁷ about whether attributing knowledge is fair enough as long as at least one kind of knowledge of human interest (reflective knowledge) is withheld.

Sosa's latest (2021) book, *Epistemic Explanations* (EE) importantly *reverses course* in such a way as to make this mitigating reply no longer available to him to appeal to. The mitigating reply is no longer available because EE introduces new theoretical machinery (the machinery of *default assumptions*) with reference to which, on Sosa's latest virtue epistemology, it is now claimed that Barney attains not merely animal but also *reflective knowledge*¹⁸, and thus, that *not just animal knowledge (apt belief) but also reflective knowledge does not require safety*.

With Sosa's previous mitigating strategy *now off the table entirely* as a way to respond to the objection that his view problematically allows unsafe knowledge, we can see that the earlier dreaming argument against safety is of special relevance. That argument is not a mitigating reply at all, but a straightforward *rebutting* reply – it says (in short) that we shouldn't accept SCK to begin with, and so it's therefore no problem for Sosa's virtue epistemology that it permits knowledge without safety.

We'll now make precise the details of this argument – a kind of dream-scepticism reductio against SCK – and in the remainder of the paper, we'll consider in depth a range of responses available to the proponent of SCK, and it will be shown that none ultimately holds water.

The unacceptable result that features in Sosa's dreaming reductio (which he takes to be implicit in endorsing SCK) is widespread scepticism about ordinary perceptual knowledge. We'll take 'perceptual beliefs' to be beliefs that are, in some loose sense, seemingly the result of perceptual engagement with the world. A paradigm example is your belief that, say, there is a computer in front of me (*computer* for simplicity). Proponents of SCK, as the reductio holds, are committed to regarding dreaming as a widespread threat to perceptual knowledge, and this is because the dreaming scenario renders beliefs like *computer* unsafe, given the following assumptions:

- (a) *The Orthodox Theory of Dreaming*: Dreams are partly composed of occurrent beliefs (typically false) formed by the sleeping subject. (Thus, while dreaming, subjects sometimes form beliefs like *computer*, which are typically false.)¹⁹
- (b) *Modal Proximity of Dreaming*: Either a sufficient proportion of nearby possible worlds to the actual world are those in which the subject is dreaming and continues to form the relevant (false) belief, or at least one very close nearby scenario

¹⁷ For discussion, see Pritchard (2012); Lackey (2009); Carter, Pritchard, and Turri (2018); Carter (Forthcoming, Chs. 1–2); Carter and Sosa (2022).

¹⁸ Note that this shift also applies to cases like that of Simone (see Sosa, 2010), who is also now (on the 2021 view) attributed along with Barney with reflective knowledge *despite* possessing unsafe belief.

¹⁹ It is worth noting that Sosa is minimally assuming that subjects undergo experiences while they are dreaming – see Malcolm (1959), Dennett (1976) for scepticism]. Proponents of a): Augustine, Descartes, Russell, Windt, Rosen, Hobson, many other contemporary dream theorists.

is one in which the subject is dreaming and continues to form the relevant (false) belief.

- (c) *Basis Similarity*: The false beliefs formed while dreaming are formed in the same way as relevant beliefs formed by waking subjects.

With these assumptions, Sosa's thought is that if, while dreaming, subjects form false beliefs based upon their experience in the same way in which they form perceptual beliefs in waking life, and if the dreaming scenario occupies a sufficient proportion of nearby possible worlds relative to the actual world, then the modal proximity of the dreaming scenario undermines the safety of ordinary perceptual beliefs, problematically so for a proponent of SCK.

Here is the structure of the reductio – where we can take the computer belief described as our illustrative case

Dreaming Reductio against SCK:

1. I know that there is a computer in front of me only if not easily would I be wrong when forming a belief “there is a computer in front of me” in the way I form this belief in the actual world (from SCK).
2. While dreaming, I form false beliefs based upon their experience in the same way as I form perceptual beliefs in waking life. (From *Orthodox Theory of Dreaming* and *Basis Similarity*)
3. When I form a belief “there is a computer in front of me”, I could easily be wrong; that is, either a sufficient proportion of nearby possible worlds to the actual world are those in which I am dreaming and continue to form the relevant (false) belief, or at least one very close nearby scenario is one in which I am dreaming and continue to form the relevant (false) belief. (From *Orthodox Theory of Dreaming* and *Modal Proximity of Dreaming*)
4. Therefore, it's not the case I know that there is a computer in front of me. (from 1 to 3).

Of course, you could plug any perceptual belief you like into the above reductio and reason from 1 to 4 given SCK in conjunction with the *Orthodox Theory of Dreaming*, *Modal Proximity of Dreaming*, and *Basis Similarity*.

Unless there is a mistake in Sosa's dreaming reductio, this is really a devastating result for SCK, and suggests why it is no objection to a theory that it doesn't accommodate it. Given the significance of the result, and given the fact that it's received little attention in discussions of safety and knowledge, we want to look closely at it. In what follows, we consider and reject three lines of response that the proponent of SCK might give.

Here is how the remaining three sections will be organised. § 3 considers how the friend of SCK could try to escape the reductio by rejecting (as noted in § 1 Sosa himself thought was a perfectly viable option) the *Orthodox Theory of Dreaming*; § 4, considers the strategy of rejecting *Modal Proximity of Dreaming*, and § 5 considers the strategy of rejecting *Basis Similarity*. None turns out to offer a satisfactory way out of the reductio. Once this point is established over §§ 3–5, § 6 then contextualises these results within Sosa's own epistemology, including in connection with what,

we'll see, is a weaker version of a safety that Sosa accepts, but which interfaces differently than SCK with the nearness of dreaming scenario.

3 Escape route for safety: reject the orthodox theory of dreaming?

Regardless of what one's theoretical commitments are about knowledge, one is free to deny – as a view about the nature of what dreaming involves – that while dreaming, subjects ever form beliefs at all (false or otherwise). Going this route requires adopting a substantive theory of dreaming that is at odds with the Orthodox Theory of Dreaming. But it offers a way for the friend of SCK to deny (at once) both 2 and 3 of the *reductio* which rely on this assumption.

Sosa's own sympathies to this route included signing on to (as noted in § 1) the Imagination Theory of Dreaming according to which – as he understood that view – at least the 'cognitive' (as opposed to the 'sensory' or 'affective') aspects of dreaming are essentially imaginative. More thoroughgoing versions of the view will hold that all aspects of dreams – sensory, cognitive, affective – are imaginative.²⁰ Henceforth the 'Imagination Theory' will refer to the more thoroughgoing view; this is for simplicity and nothing in our argument turns on this assumption.

To provide a first pass characterisation of the Imagination Theory (which we will shortly problematise), it is crucial to appreciate a distinction that we have so far glossed over, between *x* occurring *in* a subject's dream versus *x* occurring *while* a subject is dreaming (see Malcolm [1959] 2017 and Sosa, 2005). To illustrate, consider the following dream report:

PARTY: At a party with many people. Dance closely with L. Wonderful feeling. I explore many hidden rooms, find clothes and towels. Tempted to steal some but resist. Then go down to a lake where I see an old woman with grey hair who speaks German.²¹

In this subject's dream, *inter alia*, they attended a party, danced, explored, and had a chance encounter. But what goes on *in* a subject's dream – the content that they may report – is not typically a good guide to what goes on *while* the subject dreamt (and vice versa). While the subject dreamt, they were sleeping, and hence were not actually attending a party or dancing, etc. The Imagination Theory makes the further claim that, although subjects sometimes report forming beliefs, experiencing emotions, and undergoing sensory experiences *in* their dreams, it is nevertheless true that *while* subjects dream, they are imagining.

If we adopt the Imagination Theory as a theory of dream constitution then it may seem that it is never true that subjects form beliefs *while* dreaming, and so what goes on while dreaming is simply irrelevant to the assessment of the safety of our waking

²⁰ Ichikawa (2009) and Thompson (2014) more thoroughgoing versions.

²¹ Taken from the dream collection "Emma: 48 years of dreams: #1951-035 (5/16/51)" and accessed from www.dreambank.neton26/07/2023.

beliefs. Hence, the adoption of the Imagination Theory appears to afford a neat way out for the proponent of SCK.

But this move as we see it quickly reaches a problem; this is because, plausibly, the best version of the Imagination Theory is consistent with, and may even require, the claim that subjects sometimes form beliefs while dreaming.²²

To understand this, consider the case of *lucid* dreams. These are standardly defined as dreams wherein subjects believe or even know that they are dreaming while dreaming.²³ They are often accompanied by enhanced voluntary control over the dream, (though this is not thought to be definitive of the phenomenon), an ability to access waking memories, and to enact waking intentions.

While some proponents of the Imagination Theory (e.g., Ichikawa, 2009; Thompson, 2014) have appealed to lucid dreaming in support of the theory – appealing to the idea that imagining is essentially connected to the will and taking lucid dreaming to constitute evidence for dreams’ susceptibility to voluntary control – reflection reveals that they constitute a *prima facie* counterexample. The key idea here is that some dreamers are knowledgeable in dreams. In fact, some types of dreaming (lucid dreaming) are distinguished by involving the tokening of some knowledge that are not imaginative. The initial statement of the objection (which we will refine) proceeds as follows:

Lucid dreaming subjects know that they are dreaming while dreaming and can recollect waking beliefs. Imagining is never a state of knowing or believing (at least with respect to propositions about the actual world). Therefore, when lucid dreaming subjects know that they are dreaming while dreaming and can recollect waking beliefs, this involves the tokening of states that are non-imaginative. If the Imagination Theory is true, *no dreaming subjects will token states that are non-imaginative*. Therefore, The Imagination Theory of Dreaming is false.

It might seem that the obvious reply for the proponent of the Imagination Theory is to appeal to the *in* dreams vs. *while* dreaming distinction. The thought would be that lucid dreams are simply those where subjects have knowledge that they are dreaming *in their dream*, but that they lack genuine knowledge of dreaming *while* dreaming.

However, this deflationary move is highly unsatisfactory. The overwhelming consensus among dream theorists is that convergent data from lucid dream reports, experimental evidence (see eye movement lab studies from, e.g., LaBerge et al. (2018), which demonstrate communication in real time between lucid dreamers and waking subjects), and neurophysiology (reactivation of brain areas associated with metacognition), is best explained by the claim that lucid subjects *do* know (or at least believe) that they are dreaming *while* dreaming, have access to waking memories, etc.

It might be tempting then to think that the case of lucid dreaming simply undermines the truth of the Imagination Theory, in which case the proponent of SCK

²² Ichikawa (2008) argues that adopting the Imagination Theory leads to an even more radical form of scepticism. Our argument complements his.

²³ Some representative characterisations of lucid dreaming: “becoming aware of the fact that one is dreaming during ongoing sleep.” (Baird et al., 2019: 305); “the dreamer is aware that he or she is dreaming” (Erlacher & Stumbrys, 2020: 1); “Lucid dreams are a remarkable range of phenomena in which the subject realises they are dreaming.” (Rosen, 2012: 54).

couldn't appeal to it to avoid the challenge from dream scepticism. However, we think otherwise. Rather than undermining the Imagination Theory we think that lucid dreaming provides an opportunity for clarification of the account. We take it that the Imagination Theory, suitably qualified, is consistent with the tokening of beliefs and knowledge by a subject *while* dreaming so long as those states are *not instances of dreaming*. After all, the Imagination Theory is a theory of the *constitution of dreaming*. In the case of lucid dreaming, there are reasons to think that the beliefs and knowledge in question are plausibly not themselves *instances of dreaming* and so don't pose a threat to the Imagination Theory.

The reader might wonder whether it is ad hoc to posit the existence of states tokened while dreaming that are not themselves dreaming states. But there are independent reasons for thinking that there are such states. Consider, for instance, a case where a subject is having a dream which is interrupted by a perceptual experience of their alarm going off (see Macpherson 2011 and forthcoming for discussion of such cases). The thought is that the experience of the alarm – which is plausibly an instance of veridical perception – is tokened *while* the subject is dreaming but doesn't count as a dream state. Why? One reason is that it is an instance of perceptual sensitivity to the environment that is not integrated with the experience of the dream environment. In any case, the fact that this example is plausible and has nothing to do with lucid dreaming demonstrates that the move being made by the Imagination Theory is not ad hoc.

To clarify: our thought is that lucid dreaming demonstrates that the best version of the Imagination Theory is consistent with the tokening of beliefs and knowledge while dreaming, so long as those states are not instances of dreaming. Similar to the perception of the alarm, we think that the beliefs and knowledge distinctive of lucidity are states which either manifest or afford a sensitivity to the extra-dream environment, providing a perspective external to the dreamworld. Hence, they are not best classified as dream states.

As we see it, then, the proponent of SCK who appeals to the Imagination Theory of Dreaming faces a dilemma: If the beliefs and knowledge distinctive of lucid dreaming *are* instances of dreaming, then the Imagination Theory is false. If the beliefs and knowledge distinctive of lucid dreaming *aren't* instances of dreaming, then the proponent of SCK cannot appeal to the Imagination Theory to meet the challenge of dream scepticism.

We now consider two objections. First, the proponent of SCK might argue that the threat from dream scepticism arises in *nonlucid* dreams rather than lucid ones. So, even if lucid dreams involve the tokening of beliefs, this doesn't have any implications for nonlucid dreams, and therefore doesn't threaten their response to dream scepticism. A related version of this objection is to argue that the best version of the Imagination Theory should just be limited to nonlucid dreams rather than including lucid dreams (see Crowther (2018), O'Shaughnessy (2000), Soteriou (2020)).

In reply, we think that lucid dreams have ramifications for what any plausible version of Imagination Theory should say about nonlucid dreaming.

First, if the Imagination Theory allows for the existence of beliefs and knowledge tokened while dreaming in the case of lucid dreaming, this opens the possibility that there *could be such cases even in nonlucid dreams* even if the Imagination Theory is

true. The most plausible candidates are *false* awakenings (where subjects appear to believe that they are awake in their normal surroundings) and *prelucid* dreams (where subjects consider whether they are dreaming, apparently in response to beliefs about incongruities in their environment but come to believe that they are not dreaming). Interestingly, false awakenings have been found to co-occur alongside lucid dreaming and there is some limited evidence that they have a similar neurophysiology (Buzzi 2019, Voss and Hobson 2015).

Second, and importantly, the epistemic features of the beliefs formed by subjects while lucid dreaming may require that we posit beliefs with similar epistemic features formed *while subjects are nonlucidly dreaming*. To appreciate this point, consider the following dream report:

“From the top of a fairly flat unknown mountain I look out over a wide plain towards the horizon. The thought occurs to me that I do not know at all what time of day it is. I examine the position of the sun. It is almost vertically overhead in the sky and its usual brightness. This surprises me, as it occurs to me that it is already autumn and a short time ago it was much lower in the sky. I think “The sun is now standing vertically at the equator, so here it must be at an angle of about 45 degrees. Therefore if my shadow is not equal to my own height, I must be dreaming.” I look at it: it is about 30 cm long. It costs me a fair effort to regard the whole almost dazzlingly bright landscape with all its villages as an illusion.” (Green and McCreery 1994, 17, quoted in Windt & Metzinger 2007:232).

In this case, the dreaming subject appears to arrive at the realisation that they are dreaming by recognising an incongruity *and then inferring that they are dreaming*. If they come to know that they are dreaming by recognising some incongruity in their environment, then a natural way of characterising this knowledge is as a case of *inferential* knowledge. If that's right, then for the subject to be said to have knowledge that they are dreaming they must have formed beliefs that constitute knowledge while they were *nonlucidly* dreaming.

This line of thought is by no means decisive (e.g., the proponent of the Imagination Theory might argue that, contrary to appearances, these are cases of merely apparent psychological/epistemic inference), but we don't need it to be to make our point: that the proponent of SCK would be foolish to rest their hopes of avoiding dream scepticism on the adoption of the Imagination Theory. Not only is it dialectically unwise to appeal to the truth of the Imagination Theory, but there are reasonable grounds for thinking that the Imagination Theory is in fact compatible with the existence of beliefs formed while dreaming.

This brings us to the second objection to our argument. If there are problems with adopting the Imagination Theory to aid the proponent of SCK doesn't that just provide reason to adopt an *alternative* theory of dreaming? No. Adopting any of the plausible alternative theories of dreaming currently defended – e.g. the view that dreams are *sui generis* states (Windt 2015), or a pluralist view about dreaming such that only some dreams are imaginative (Rosen, 2012), the view that dreaming is a kind of intensified mind-wandering (Domhoff 2023) (a close cousin of the Imagina-

tion theory) - will not help matters.²⁴ First, such a move still suffers from the dialectical problems already highlighted. Second, and more importantly, the line of thought that we have presented regarding lucid dreaming has general applicability. That is, any theory of dreams which sets out to deny that subject sometimes form beliefs while dreaming will need to contend with the line of argument we have just sketched argument concerning the status of lucid dreaming and its relation to ordinary nonlucid cases.

In conclusion, the proponent of SCK cannot address the challenge of dream scepticism simply by tying themselves to a denial of the Orthodox Theory of Dreaming.

4 Escape route for safety: reject the modal proximity of dreaming?

There are other ways the SCK proponent could get out of the reductio even by accepting the Orthodox Theory of Dreaming. In outline form, they might attack instead *Modal Proximity of Dreaming*. Maybe we form beliefs in dreaming, but it's false that when perceiving when awake, we could *easily enough* be dreaming and forming false beliefs such that this possibility would render our perceptual beliefs unsafe.

Let's consider what this move involves in a bit more detail. With reference to SC, taking this strategy requires that one either deny that a sufficient proportion of nearby possible worlds are those in which the subject is dreaming and continues to form the belief, or deny that there is at least one *very close* nearby world in which the subject is dreaming and continues to form the belief.

For a true belief that *p* to meet SC, it must be the case that in *most nearby possible worlds* and in *all* very close nearby worlds in which *S* continues to form her belief that *p* in the same way as in the actual world, *S*'s belief that *p* continues to be true. So, for instance, for my *computer* belief to meet SC, it must be the case that in all or most alternative scenarios in which I form this belief in the same way as in the actual world, and which differ from the actual scenario in relatively minimal ways, my belief comes out as true. The challenge for the proponent SCK is that for many examples of true beliefs that *p*, including perhaps my *computer* belief, it seems that a range of nearby worlds in which I form this belief are ones in which I am asleep and dreaming. But if that's right, then it looks like ordinary perceptual beliefs such as my *computer* belief will fail to meet SC. Either most nearby worlds or at least one very close nearby world in which the subject forms her belief that *p* in the same way as in the actual world, are such that the belief that *p* is false. This point generalises to many (apparently) epistemically innocuous beliefs about our environment. Thus, an implausible scepticism looms for proponents of SCK.

The intuitive idea, and the one which we examine in this section, is that it wouldn't require much about the actual scenario (in which I form various beliefs about my environment) to change for me to find myself in the dreaming scenario. Sleeping

²⁴ The only theory that *might* help avoid dream scepticism is the one associated with Malcolm and Dennett which denies that dreams are experiences or thoughts tokened while a subject is sleeping. This is almost universally rejected so it's an even worse candidate for helping the proponent of SC-Safety! A pithy way of putting it: it might avoid an implausible epistemic scepticism only at the cost of a highly implausible scepticism about dreaming.

and dreaming are, after-all, natural products of normal biological functioning that can occur without (and can sometimes be thwarted by) conscious effort. Two comparisons are germane. First, consider again Barney in BARNS. Just as it wouldn't have required much to change about Barney's actual scenario in order for his belief that *there's a barn* to turn out to be false (instead of looking at the sole real barn, he finds himself looking at a facsimile), it wouldn't take too much about my current scenario to change in order for me to be dreaming and for my *computer* belief to be false. Second, both BARNS and the dreaming scenario stand in contrast to more 'hyperbolic' alternative scenarios in which I'm a brain-in-a-vat or being deceived by a malevolent demon. We can assume that *a lot* would need to change about our current scenario for us to find ourselves in these benighted worlds. Unlike fake barns and dreaming, these do *not* threaten the safety of our beliefs.

We consider and reject three broad options for the proponent of SCK to respond:

- i. *Remoteness of dreaming when forming beliefs about our environment*: The dreaming scenario is never, or is only rarely, a nearby world relative to the actual world in which we form ordinary beliefs about our environment.
- ii. *Remoteness of forming beliefs about the environment while dreaming*: The scenario in which I form belief in the target proposition about the environment while dreaming will rarely be a nearby world relative to the actual world.
- iii. *Insufficient non-remoteness of forming beliefs about our environment while dreaming*: The scenario in which I form belief in the target proposition about the environment while dreaming may often be a nearby scenario, but such worlds do not constitute a sufficient proportion and none of them are very close nearby worlds.

To understand option i, *Remoteness of dreaming when forming beliefs about our environment*, first note that it is plausible that there will be many waking scenarios where the dreaming scenario is *not* well characterised as a nearby world. For instance, Jay is just about to give their inaugural Professorial lecture. They are full of adrenaline and have consumed a large quantity of caffeine. We submit that rather a lot about Jay's actual scenario would have to change for them to find themselves in a dream scenario. Mostly this is because it would take a lot for them to fall asleep.

Despite this, we think that there are lots of circumstances where the dreaming scenario will not be so distant. For instance, John is a single father with two very young children. He gets very little sleep at night and works gruelling 9-hour shifts during the day. He is generally sleep deprived. We submit that for much of the large swathes of time during which John is awake very little would have to change for him to find himself sleeping and dreaming.

Most people, most of the time, do not find themselves in circumstances like Jay or John. Yet, it seems that often we will be in scenarios such that it would not take much for us to find ourselves in the dream scenario. Here is a cursory list of such scenarios: engaging in vigorous exercise, reading (especially the works of Derrida and Foucault) at night (or during the day!), watching a movie, long-distance travelling, listening to music (especially minimalist music), listening to audiobooks or long-form podcasts, consuming alcohol in a comfortable chair, and eating large quantities of starchy food.

Given this, we think that the proponent of SCK cannot credibly claim that the dreaming scenario is never or only rarely a nearby alternative to the actual world.

However, the reader will note that we've so far simply been referring to the nearness of the dream scenario to the actual scenario, when what is relevant to assessing safety is the nearness of the dream scenario in which *I continue to form the belief that p*. This brings us to option ii: *Remoteness of believing while dreaming*.

Even if it wouldn't take much for us to find ourselves dreaming, and even if we form beliefs while we dream, the proponent of SCK might claim that the scenario in which I form ordinary beliefs about my environment such as *computer* while dreaming *will not* constitute a nearby alternative to the actual scenario in which I form this belief.

To support this, they might ally themselves with the view of dreaming popularised by neuroscientist Allan Hobson which claims that dreams (at least those occurrent during REM sleep) aren't faithful replications of our engagement with ordinary perceptual environments. Instead, normal dreaming is distinguished by a combination of bizarre content – including but not limited to environmental distortions, unusual juxtapositions, sudden scene changes (see Hobson 2002, Domhoff, 2005, for discussion) – plus a substantial impairment of our ability to detect it as such. Indeed, Hobson is well-known for characterising dreams as a kind of psychosis or delirium (Hobson 1999). These claims are mostly grounded in neurophysiological evidence concerning REM sleep (which is highly but imperfectly correlated with dream reports): that the REM brain is distinguished by a pattern of activation and neurochemical modulation that is very different from the normal waking brain and is such that we might expect dream content and dreamer capacities to have these psychological features.

Assuming that something like Hobson's model is correct (it coheres well with the popular conception of dreaming and is suggested by the dream-inspired artworks such as those created by the Surrealists), the proponent of SCK might argue that the kinds of beliefs being formed by dreaming subjects are unlikely to bear any relation to the kinds of beliefs that they are likely to be tokening while awake. So, it will not be the case that a *nearby* possible world in which I form the *computer* belief will be one in which I am dreaming. Too much about the nature of ordinary dreaming would have to be different.

We make two points against this line. First, a consideration of multiple dream content analyses reveals that Hobson's claims about dream bizarreness are, at best, overstated. Here is Domhoff (2005) summarising the evidence:

"Contrary to his [Hobson's] belief that dreams are 'cognitive trash' (2002: 23) best characterised by the symptoms of delirium, including illogical cognition, unstable emotion, and dull intellectual functions, dreams are most often reasonable simulations of waking life that contain occasional unusual features in terms of setting, characters, and activity (Dorus et al 1971; Foulkes 1985; Hall & Van de Castle 1966; Snyder 1970)" (Domhoff, 2005: 14)²⁵

²⁵ The most striking example of faithful recreations of waking life are *false awakenings* where dreaming subjects are convinced that they have woken up, seemingly in a perfect replica of their waking environment. Note, however, that these appear to be infrequent.

Our second point is that dreaming about events in the previous day is very prevalent, occurring in about 60–75% of dream reports (Nielsen & Powell, 1992). Further, a recent study (Picard-Deland et al., 2023) suggests that dreams about the events of the previous day are even more prevalent during the early stages of a night's sleep and particularly during N1 and REM sleep (dream content from later in the sleep cycle tends to draw more upon memories of distant events). This is significant because N1 is the period immediately after we fall asleep, and thus has greatest temporal proximity to the time at which we are awake. Putting this nascent research into dream content together, we think it is entirely reasonable to think that the possible world in which I am forming the *computer* belief while dreaming could often constitute a *nearby* alternative to the actual world. Hence, option ii fails for the proponent of SCK.

This brings us to the final option, iii, *Insufficient non-remoteness of forming beliefs about our environment while dreaming*. The proponent of SCK can argue that even if it is often true that the scenario in which I form the *computer* belief while dreaming is a possible world nearby to the actual world, such possible worlds either do not constitute a sufficient proportion of nearby worlds or are never very close to the actual world. Therefore, the dreaming scenario doesn't threaten Safety.

To clarify things, consider again BARNS. Barney forms the true belief that there is a barn because, luckily for him, he happens to be looking at the *one* real barn in fake barn country. But in nearby worlds in which he forms the belief that *there is a barn in the field* he will be looking at a fake. Further, and crucially for present purposes, there is a clear sense in which *most* of the nearby possible worlds in which Barney forms a belief that *there is a barn in the field*, his belief will be false (because of all the barn facsimiles in his environment), and at least one very close nearby scenario will be one where his belief that *there is a barn in the field* will be false. Returning to the dream case: the thought is that while there may be multiple analogies between BARNS and the dream scenario (as we have argued above) we lack reason to think that there is a parallel between these cases with respect to the proportion of nearby worlds, or with respect to existence of at least one very close possible world, where the subject is dreaming and believes falsely.

Against this, we think that the burden of argument is squarely on the shoulders of the proponent of SCK who claims that the dream scenario in which the subject continues to form belief in the target proposition is *never* a very close possible world. Given the above arguments concerning the proximity of the dream scenario and the similarities in dream content to waking content, we think that this is going to be a tall order.

When we consider the proportion of nearby worlds more generally, it might be thought that the argumentative burden rests upon us to show that the dreaming scenario is such as to often undermine this condition for ordinary perceptual beliefs. Note that this needn't be because the dream scenario in which we form a false perceptual belief *itself* constitutes *most* nearby worlds, but because the proportion of such worlds is such as to render false the claim that *most* nearby worlds where the subject continues to form the perceptual belief are ones where the belief is true.

While we don't take ourselves to have provided compelling reasons in support of this, neither do we think that there are compelling reasons against. In any event, even

if the dreaming scenario isn't itself a showstopper when it comes to safety, it will nevertheless constitute a kind of standing background threat to the safety of many of our ordinary perceptual beliefs. Thus, the proponent of SCK is committed to thinking that dreams make ordinary perceptual knowledge less secure and its acquisition more difficult. That is a bad result, though perhaps not as bad as it initially seemed.²⁶

Further, we have shown that the proponent of SCK is committed to the following: for a typical person our confidence that a given perceptual belief of theirs constitutes knowledge will decrease relative to the ease with which they might fall asleep, independently of any connection to the condition of their waking capacities. That is, if it turns out that Jane can fall asleep more easily than James, then the proponent of SCK would seem to be committed to claim that Jane's beliefs are less epistemically secure than James', where this has *nothing* to do with the condition of their waking perceptual belief-forming mechanisms and everything to do with the fact that Jane finds it much easier to fall asleep. While not decisive, that seems like a strongly counterintuitive result.^{27,28}

²⁶ It's worth registering that a proponent of SCK might attempt to resist the idea that the dreaming scenario could easily obtain by reasoning from the premise that our perceptual beliefs are (often) sensitive (to various deception possibilities) when we are awake. For instance, when awake, if someone removed the computer from its position on the table and took it out of the room, we would not continue to believe it is there. Might this observation about the apparent sensitivity of our perceptual beliefs when awake suggest that not easily would we be mistaken about our belief "there is a computer" when awake? We think not; for even if it is true of us when awake that not easily would we think the computer is present in the nearest worlds where it's physically removed from the room, it might nonetheless be (as we have suggested) that very easily we might be dreaming. The unsafety of the target belief in the actual world is accordingly *compatible* with retaining when awake various discriminatory capacities, including some which feature sensitivity while awake. Thanks to Martin Smith for raising this point.

²⁷ In fairness to the proponent of SCK: they might argue that their theory gives us the correct verdicts in cases in which it seems relatively clear that believing *p* while dreaming constitutes most nearby worlds to the actual world in which the subject believes *p*. Here is a case: Jeb is the subject of recurrent false awakening dreams; in his dream, he wakes up and forms normal perceptual beliefs: that his husband is lying next to him, that there are clothes strewn across the floor, that the mirror is dusty, etc. Considering claim (a) above, we can assume that Jeb also forms these beliefs *while* dreaming. Let's assume that Jeb is subject to such dreams 4–5 times per week. Given all of this, when Jeb wakes up each morning and forms tokens of these types of perceptual beliefs, it is plausible that at least some of these perceptual beliefs will fail to meet the conditions for Safety. But, and in defence of SCK, that seems like the correct verdict! If Jeb really were plagued by false awakening dreams (note that these can come in cycles, such that it takes multiple false awakenings before the subject really wakes up), then this would plausibly cast an epistemic shadow over his perceptual knowledge immediately after waking.

²⁸ It might be thought that Modal Proximity can be challenged by appeal to a form of *semantic externalism*. The thought is that for dream beliefs such as "there is a computer in front of me", the indexical "there" is going to have a different extension in my dream than it would if I were awake and actually looking at a computer. If that's right, then the dream scenario would not constitute a nearby scenario in which I continue to form the *computer* belief that I form in waking life. Aside from having doubts about whether semantic externalism would have this application to dream beliefs, the more straightforward reply is to point out that not all dream beliefs will include indexicals in their content, e.g. "a computer exists", and that others include indexicals which would not plausibly shift extensions between dreaming and wakefulness, e.g. "I am looking at a computer". Thanks to Lilith Newton and Ernest Sosa for helpful discussion.

5 Escape route for safety: reject basis similarity?

Let's consider now a final way out of the dreaming *reductio* against SCK, which involves denying *Basis Similarity*, viz., to deny that the beliefs formed while dreaming are formed in the same way as those formed by waking subjects.

Assume, for now, that subjects form beliefs while dreaming, and that the dream scenario is relevantly nearby to the actual scenario. The proponent of SCK might attempt to resist the epistemic threat posed by the dreaming scenario by identifying a contrast between the *ways* in which subjects form beliefs in waking life and while dreaming. Recall that SC requires that we consider nearby possible worlds in which S continues to form her belief about the target proposition *in the same way* as in the actual case. If it can be persuasively shown that dreaming beliefs are formed in a different way from waking beliefs, then the proponent of SCK can avoid the challenge.

We here consider and reject three proposals:

- A. *External basis disanalogy*. Ordinary perceptual beliefs are formed on external bases in a way that dream beliefs are not.
- B. *Imaginative basis disanalogy*. Dream beliefs are formed on the basis of imaginings in a way that ordinary perceptual beliefs are not.
- C. *Neurophysiological state disanalogy*. Dream beliefs are formed in a distinctive neurophysiological state in a way that ordinary perceptual beliefs are not.

The first option, *External basis disanalogy*, requires that external bases such as extra-mental environmental features are implicated in the way that ordinary perceptual beliefs are formed. For instance, a description of the way in which I formed my *computer* belief will refer to the presence of the computer on my desk. This contrasts with my dream belief that there is a computer in front of me; the description of the way in which this belief is formed will *not* refer to the presence of the computer on my desk, even if the methods of belief formation are otherwise similar. Given this, beliefs formed while dreaming are irrelevant to the assessment of the safety of our ordinary perceptual beliefs.

We think that there are (at least) two serious problems with option A.

First, there is the worry that this kind of move deals with scepticism (including the hyperbolic kind) far too easily. The sceptical challenge is based upon the assumption that there are common bases in the good and bad cases, and hence there is a genuine concern that this doesn't take the challenge seriously enough.

Second, while an appeal to external bases would address the challenge from dreaming, it generates unwelcome results elsewhere for the proponent of SCK. To appreciate this, note that dreaming is not unique as a circumstance where beliefs are formed in a way that *does not* implicate external bases for the belief. This is also true of waking hallucinatory experiences. Thus, an implication of option A is to consider hallucinatory beliefs as irrelevant to assessing the safety of ordinary non-hallucinatory perceptual beliefs. But now consider an amended version of BARNs.

BARNs + LESION: Suppose that Carney forms their actual belief in a way that refers to the presence of the sole actual barn in their environment; but instead of

being otherwise surrounded by barn facsimiles, suppose that Carney has a rare brain lesion such that they would hallucinate a barn were they to be looking at any other barn-like structure than the sole real barn, and that there would be no tell-tale signs (i.e., defeaters) were the lesion to be activated in this way.

If the proponent of SCK excludes the dreaming scenario as relevant to the safety of ordinary perceptual beliefs, it seems that they are committed to excluding the nearby possibilities that Carney hallucinates a barn as relevant to the safety of their actual true belief that *there is a barn in the field*. Given their verdict about the original BARNs case, we assume that the proponent of SCK will want to say Carney lacks knowledge in this case [their belief is luckily true²⁹], but the proponent of SCK cannot explain this in the way that would seem most natural for them, i.e., that Carney's belief is unsafe.

Option B requires that we also adopt some version of the Imagination Theory of Dreams; specifically, a form of the view which characterises the sensory aspects of dreams as imaginings.³⁰ If we adopt this, then there will be a clear discrepancy between the ways in which dream beliefs and waking beliefs are formed: the former, but not the latter, are based on sensory imaginings. Hence, dream beliefs are of no relevance to the safety of waking perceptual beliefs.

We think that this move is dialectically problematic for the reasons highlighted in the previous section. To repeat: we think that it is highly unwise for proponents of SCK to make the plausibility of their view dependent upon the resolution of a highly contentious (and relatively nascent) debate in philosophy of mind in their favour.

Further, recall that our argument that a thoroughgoing version of the Imagination Theory (that which applies to all aspects of dreams) is in fact compatible with the tokening of beliefs while dreaming (so long as the beliefs aren't dream states). The point of interest is that this also applies to restricted versions of the Imagination Theory, including those which apply only to the sensory dimensions of dreaming. That is, even if the sensory aspects of dreams are imaginative, this is compatible with perceptual experiences being tokened while dreaming, so long as those perceptual states are nondream states. For instance, while dreaming about climbing Ben Nevis, a typical ambient sound such as the sound of my bedside alarm clock might cause me to have an illusory perceptual experience of a fire-engine's siren. Partly due to the connection to the external environment, and partly due to the lack of integration with the content of the dream, such an experience seems like a good candidate for a nondream state that is nevertheless tokened while dreaming. Suppose, then, that a subject formed a belief on the basis of *this* perceptual experience, e.g. that *there is a fire-engine nearby*. This would be an instance of a subject forming a belief in the same way as ordinary perceptual beliefs while dreaming. If as seems the case, this

²⁹ The unsafety of the target belief in this case is broadly analogous to the unsafety of the target belief in Russell's (1912) stopped clock case. The clock is functioning correctly as an indicator of the time when one looks at it just when the time matches the correct time, but not otherwise. Otherwise, one accepts the deliverance of the clock and is wrong. Compare: when Barney* looks at a real barn, his otherwise faulty perceptual faculties deliver a correct result just by luck.

³⁰ This is the view defended by McGinn; conversely, Sosa thinks that the cognitive aspects of dreaming are imaginative, and the sensory aspects are perceptual.

is a genuine possibility, the proponent of SCK cannot depend upon the truth of the restricted Imagination Theory to neutralise the epistemic challenge from dreams.

Finally, option C, *Neurophysiological state disanalogy*, claims that, because of the distinctive neurophysiology of dreaming, this makes a difference to the way in which dream beliefs are formed relative to ordinary waking perceptual beliefs. To understand this, we borrow from Lyons (2019) the distinction between *algorithms* and *parameters*. Very roughly, algorithms are procedures for mapping inputs to outputs, e.g., visual face recognition, while parameters affect the workings (and reliability) of these procedures, e.g., whether someone is intoxicated. Applied to the dream case, the thought is that, even if dream beliefs and waking beliefs are similar in terms of the algorithms implicated in the way they are formed, these procedures may be subject to very different parameters. If so, then the ways in which they are formed will be different, and dream beliefs won't be relevant to assessing the safety of waking perceptual beliefs.

To support this move, the proponent of SCK might return to Hobson's neurophysiological claims about dreaming. Even if we have doubts about the implications of this for the bizarreness of dream content, we might nevertheless take seriously Hobson's claims that the dreaming brain is in a similar state to that found in psychosis patients or those under the influence of powerful psychoactive drugs. In some more detail: the dreaming brain (at least in REM sleep) undergoes a distinctive neurochemical modulation. The normal waking brain is marked by a balance between aminergic neurotransmitters (such as noradrenaline and serotonin) and cholinergic neurotransmitters (such as acetylcholine), with the balance in favour of the former. While amines are "essential to the processes that enable us to direct attention, reason things through, and decide to act" (Clark 2007: 8), when cholines dominate "emotional and analogical reasoning begin to dominate, and critical control and judgment wane." (Ibid: 8). Crucially, the REM brain is "aminergically demodulated and, reciprocally, cholinergically hypermodulated" (Hobson 2009: 810).

This modulation complements a distinctive pattern of brain (de)activation in REM sleep. Positron emission tomography and fMRI findings evidence a shift in regional blood flow away from the dorsolateral prefrontal cortex (DLPFC) to subcortical limbic structures such as the amygdala. This is significant, because the DLPFC is the "executive brain" which helps us to "organise our thinking, critically assess our own gut responses and maintain at least a modicum of top-down control" (Clark 2007: 4) and is associated with control of willed actions (Jahanshahi & Frith, 1998). The limbic system, on the other hand, is linked to analogical and associative thinking, and emotion.

Putting this together, the proponent of SCK might argue that we have good empirical grounds for thinking that there are distinctive parameters in play while subjects are dreaming such that dream beliefs are formed in a different way from ordinary perceptual beliefs formed in waking life. If that's correct, then dreaming is of no relevance to SC.

We make three replies. First, there are dialectical worries like those raised earlier about reliance on the truth of the Imagination Theory of dreams. Here, the concern is that the proponent of SCK would be very unwise to rest the plausibility of a seriously controversial empirically grounded theory of REM sleep. To illustrate, consider that

other promising theories of the neurophysiology of dreaming – such as Domhoff’s neurocognitive theory – do not identify analogies between the dreaming brain and psychosis or that exemplified with psychoactive drug use [instead, on Domhoff’s view the dreaming is closely connected to the operations of the “default network” which is a network that is very often operative while we are awake]. Without going further into the details, the general point is that on competing neurophysiological theories of dreaming it is far from obvious that there are different parameters operative compared to when waking.³¹

Second, and related, even if we were to accept as true Hobson’s model of the REM brain, it is far from obvious that this is relevant to assessing SC vis-à-vis the dreaming scenario. Recall from the discussion of Sect. 4 that it will be dreams tokened in the NREM1 stage of sleep that are most relevant to assessing SC. Given that Hobson’s model concerns the REM stage brain and given that the NREM brain is different in many key aspects from the REM brain, the empirical findings have, at best, very limited relevance.

Finally, we think that a further modified version of BARNs + LESION reveals problems for the proponent of SCK who wants to take option C.

BARNs + LESION*: Suppose that Carney again forms their actual belief that *there is a barn in the field*. This time, instead of being otherwise surrounded by convincing barn facsimiles, suppose that Carney is in an environment where there are *obviously* fake barn facsimiles. But suppose, further, that Carney has a rare brain lesion such that were they to be looking at any of these obvious fakes they would find themselves in a psychological state like psychosis and would end up forming a belief that *there is a barn in the field*.

Given their verdict about the original BARNs case, we expect that the proponent of SCK will want to diagnose Carney’s actual belief is luckily true. But if the proponent of SCK takes option C to deal with dreaming then they are unable to do this, at least in terms of a failure of Safety; this is because they will be committed to claiming that, in nearby worlds, Carney forms his belief in a different way from the actual. To avoid this kind of problem in BARNs + LESION*, the proponent of SCK might amend how they define SC: specifically, they might make it more restrictive such that, when assessing safety, we need to consider nearby worlds in which the belief is formed in the same way *and relevantly similar ways*.³² Then, in BARNs + LESION*, the nearby worlds would be such that Carney forms their belief in a relevantly similar way to the actual. However, this will also apply, *mutatis mutandis*, with respect to the dream scenario, i.e., even if there is a difference in the way in which dream beliefs in nearby worlds are formed, these are likely going to be relevantly like the way deployed in the actual world. If that’s right, then C fails as a way of avoiding the epistemic challenge from dreaming.

³¹ Along the same lines: some evidence suggests that the REM brain is heterogeneous, e.g., Kubota et al. (2011) identify DLPFC *activation* during REM sleep.

³² As Williamson (2000) puts it, “If one knows, one could not easily have been wrong in a similar case” (2000, 147).

6 Dreaming and SSS-relative safety

Let's take stock. We began by reviewing Sosa's own well-known appeal to the Imagination Theory of Dreaming to sidestep the threat of dream scepticism, a threat that Sosa we think rightly takes (unlike, e.g., BIV-style hypotheses) to be *modally close* in a way that raises a serious problem for proponents of SCK, in so far as they want to uphold the scope of everyday perceptual knowledge we have.

Because Sosa doesn't want to fully stake the capacity of his own theory of knowledge (which allows apt belief to be sufficient for knowledge) on a contentious theory of the nature of dreams, he argues his view is untroubled by the nearness of the dreaming hypothesis *even if we insist on the Orthodox Theory of Dreaming*, where dreaming is compatible with *believing* falsely. As we noted, the *most discussed* version of this strategy (which assumes for the sake of argument the Orthodox Theory of Dreaming) is a *mitigating reply*—viz., the reply that *at least Sosa doesn't attribute reflective knowledge* in, e.g., fake barn cases where the protagonist's belief is unsafe due to the nearness of the dreaming scenario.

Moreover, we saw that, with this mitigating line now out the window entirely given Sosa's latest virtue epistemology in *EE* (which grants Barney not just apt belief (i.e., animal knowledge) when his belief is unsafe, but *reflective knowledge*), it is all the more critical to look at the credentials of his remaining reply – which says that if we accept that SC is necessary for knowledge (as per SCK), we end up forced (with just minimal assumptions) to the conclusion that we have much less perceptual knowledge than we take ourselves to have. So, SC is – controversially – too strong a condition on knowledge to begin with.

Much of this paper has looked carefully at the specific theses (at least one of which) the proponent of SCK would need to reject to escape Sosa's dreaming reductio. These, we saw, were the *Orthodox Theory of Dreaming*, the *Modal Proximity of Dreaming*, and *Basis Similarity*. We argued in each case why we think the proponent of SCK will ultimately have no luck taking any of these routes; we considered anticipated responses to our arguments and responded to each of them. At the end of the day, it looks like Sosa is right that considerations of dream scepticism should lead us to reject SCK, rather than to think that positing unsafe knowledge is any serious theoretical cost.

A few residual points remain. Recall that the kind of knowledge one has when one attains an apt belief is meant to be 'luck-resistant enough' to rule out knowledge in standard Gettier cases (even though not in environmental luck cases like BARNs). The reason is that aptness – by virtue of implicating the exercise of an epistemic competence in conditions that are proper for its exercise – requires a very specific kind of luck-resistant modal robustness – viz., a *kind* of safety that is simply overtly relativised to conditions *other than what safety is standardly overtly relativised to*, e.g., as in SC. We want to conclude by clarifying exactly (i) why aptness involves (a kind of) bona fide safety, and (ii) why *this* kind of safety is not threatened by the nearness of the dreaming scenario and thus immune to the kind of reductio that SCK invites, *despite* being strong enough to rule-out Gettier protagonists as knowing.

So in what sense does aptness implicate a kind of safety? Let's return to our example belief, *computer* (i.e., here is a computer in front of me). Suppose now that you

are drunk, drugged, and hallucinating in a pitch dark, computerless room, and you form this belief (“there is a computer in front of me”). Your belief here is *incompetent* in the way that disqualifies it from being apt, and so you’d fail to know even if you did happen to be looking at a computer. However, even when you’re drugged and in the dark, etc., the following conditional *might still apply to you* in a way it would not apply to, say, a child who had never seen a computer before:

If you were in appropriate shape and appropriately situated to identify a computer, not easily would you be wrong.

Conditionals like the above correspond with the possession of a computer spotting *skill* (alternatively, for Sosa, ‘innermost competence’) that a subject can retain even when *not in the proper shape and situation to exercise that skill*. Proper shape here will at least include, e.g., being sober, awake, alert and the proper situation will include, e.g., normal ambient light, etc. These (and not the conditions described above!) are the conditions under which we value reliable performance at spotting computers, and so they’re the ones in which computer-spotting ‘competence’ attributions will be indexed to.

When a belief is *apt* (and thus known), the belief’s accuracy must be *attributable* to adroitness – and so attributable to competence in a way that makes overt reference not just to the skill, but to the skill exercised in the (seat and shape) conditions appropriate for its exercise.

CORRECTNESS ATTRIBUTION PRINCIPLE: For any correct belief... that p, its correctness is attributable to a competence only if it derives from the exercise of that competence in conditions appropriate for its exercise, *where that exercise in those conditions would not too easily have issued a false belief...* (my italics).

The ‘conditions appropriate for its exercise’ just involves being in proper shape and properly situated to exercise a skill. Crucially, note the ‘not too easily’ locution that features in the Correctness Attribution Principle. If your correct (true) belief is attributable to competence in the way that is necessary and sufficient for aptness, it follows you’ve exercised a skill *in conditions in which you couldn’t easily be wrong*, which will be just the conditions that constitute the *shape* and *situation* appropriate for the exercise of that skill. Put another way, when your belief is apt, then (by the correctness attribution principle) it will follow that your belief must be *SSS-safe*:

SSS-RELATIVE SAFETY: A belief is SSS-relative safe just in case: In close worlds where S believes p from skill S exercised in the actual world along with the shape S_h and Situation S_i that correspond with skill S, p is true.³³

³³ This is modified version of a formulation captured by Greco (2020, 5152).

SSS-relative safety effectively ensures that knowers couldn't easily be wrong *in the conditions that matter* for the exercise of the complete competence implicated by aptness and thus by knowing.³⁴

Question: with SSS-RELATIVE SAFETY on the table, it's worth considering whether the aptness account of knowledge (by rejecting SC safety) really "gives up safety" in any theoretically significant sense, as opposed to simply giving up *one* among several ways of overtly relativising safety to conditions. We say 'overtly relativising' because, notice, even Pritchard rejects safety *simpliciter* (i.e., the thesis that known beliefs couldn't easily have been wrong simpliciter) in favour of an overtly relational formulation of safety. A proponent of SSS-safety thus accepts (like Pritchard) (i) the core idea at the heart of safety as a condition on knowledge—viz., the idea that knowers couldn't easily have been incorrect; and rejects (also, like Pritchard) a (ii) *non-overtly relational* characterisation of safety aimed at capturing that insight. These points allow us to situate the 'giving up safety' criticism of the aptness view in a different light: the rejection of SC by the proponent of Sosa's "knowledge=aptness" equivalence is *merely* a rejection of the *specific characterisation of the respect in which safety is overtly relational* that is captured in SC, and it is against a background of *agreement* with proponents of SC about (i) and (ii). What the proponent of the aptness theory holds is that the *right* overtly relational characterisation of safety should be with reference to near-by worlds where we hold fixed not the way the belief was formed (SC), but rather, something else—namely, the seat/shape/situation tuple that corresponds with the exercise of a complete competence.³⁵

The above point brings us to our final residual issue. Now that we see that it's at best misleading to say that a view that holds aptness to suffice for knowledge *gives up safety* simply by rejecting the SC-formulation of it, it's worth clarifying how we should answer an anticipated question: Can we run a dreaming reductio to SSS-safety, assuming the *Orthodox Theory of Dreaming*, the *Modal Proximity of Dreaming*, and *Basis Similarity*?

The answer here is 'no'. To appreciate why, it's worth clarifying the sense in which a belief can be SSS-safe and still 'due to luck' in the kind of way that the correctness of my belief *there is a computer in front of me* might be lucky in so far as I could easily be dreaming but am not. Here's Sosa:

If the act is due to a competence exercised in its appropriate conditions, its success may be due to luck in various ways. It may be just an accident that the agent retains his relevant competence, for example, or that the conditions

³⁴ SSS-relative safety is a necessary condition on knowledge for Sosa. However, it is not an *irredundant* necessary condition. In an account A of any target phenomenon X, condition C is an irredundant necessary condition for X in A if and only if the satisfaction of C is necessary for X's obtaining, and the satisfaction of C is not entailed by the satisfaction of any other condition, C' in A. While on the aptness account of knowledge, SSS-safety is necessary for apt belief – if a belief is not SSS-safe, it is not adroit, and therefore not apt – SSS-relative safety is not an irredundant necessary condition. The conditions for SSS-safety are, in the aptness account of knowledge, entailed by the satisfaction of the adroitness condition, viz., the condition that the belief manifest a complete SSS-competence. For discussion of redundant and irredundant necessary conditions on knowledge, see Williamson (2000, Ch. 1) and Cassam (2009).

³⁵ For more detailed discussion on these points, see Carter (Forthcoming, Chs. 1–2).

remain appropriate. Either way, the act fails to be safely successful, since it might too easily have failed, through lack of the required competence or conditions. It might still be apt, nevertheless, indeed attributably, creditably apt. (2007, 81).

When you are awake and aptly forming beliefs about your environment, this implies the obtaining of the SSS conditions for the relevant competence exercised. But it doesn't imply their *safe* obtaining. And when they do obtain, but unsafely so, this doesn't undermine your aptness, even if perhaps it does undermine the aptness of the belief's *apt attainment*.³⁶

The dreaming reductio against SC-safety is accordingly not repurposable by a critic of SSS-safety.

7 Concluding remarks

In conclusion, we've seen that although Sosa's latest line on fake barns in *Epistemic Explanations* (2021) leaves him with but a single card to play – the Dreaming Reductio Against Safety – against the charge that his account problematically countenances unsafe knowledge, this Dreaming Reductio ultimately holds up, and (despite the little attention this argument has received) it gives us a good reason at the end of the day to reject safety, at least the formulation of safety that is taken for granted in contemporary discussions of safety as a condition on knowledge. What the Dreaming Reductio *doesn't* suggest, though, is that a suitably anti-sceptical account of our perceptual knowledge needs to abandon safety, differently overtly relativised, viz., to SSS-conditions. If knowledge requires SSS-safety, rather than SC-safety, there is no analogous sort of dreaming reductio that is going to threaten the scope of our perceptual knowledge. Taken together, these points reveal a number of ways in which virtue epistemology, at least Sosa's AAA-centred brand of it, far from succumbing to

³⁶ It's worth noting briefly how SSS-safety interfaces with *basis similarity*. There are two main points to note here. Firstly, as a dialectical point, because a belief's being SSS-safe is compatible (unlike a belief's being SC safe) with one's easily having been dreaming and so believing falsely, the proponent of SSS-safety lacks any dialectical pressure in the first place to give up (e.g., as the proponent of SC-safety would need to do) any of the *Orthodox Theory of Dreaming*, the *Modal Proximity of Dreaming*, or *Basis Similarity* in order to avoid the sceptical result that we lack much perceptual knowledge whenever we easily would be dreaming. Even so, one might wonder, as a point of independent interest, whether the proponent of SSS-safety is committed to rejecting basis similarity nonetheless; after all, as the thought goes, the *shape/situation* pair that obtains when a perceptual belief is SSS safe will include that one is awake. At this point, it is helpful to distinguish the kind of wholesale appeal to external bases (which was, we suggested, the problematic option by which a proponent of SC-safety might attempt to reject basis similarity), and a way of at least partially preserving basis similarity that remains available to view that knowledge apt (and thus, SSS-safe) belief. In the latter case, we can conceive of (in Sosa's terminology) the seat or *innermost* competence (Sosa 2017, 191) that one exercises when awake as being common across the two cases whereby one is awake and dreaming. In this respect, then, the idea that knowledge is apt belief, and thereby belief that is SSS-safe, isn't committed to (despite being an externalist strategy) denying basis similarity in the same entirely non-concessionary way that would involve denying anything commonality whatsoever to the bases in the good case and bad case. More importantly, though, the proponent of SSS-safety is under *no dialectical pressure* (as the proponent of SC-safety was) to stake its anti-sceptical credentials on any such denial. Thanks to a referee at *Synthese* for prompting discussion on this point.

criticisms related to safety, can instead claim (contrary to received opinion) its treatment of safety as a theoretical advantage.

Acknowledgements Carter's research is supported by the Arts and Humanities Research Council's *Expanding Autonomy* (AH/W005077/1) and *Digital Knowledge* (AH/W008424/1) projects and the Leverhulme Trust's *A Virtue Epistemology of Trust* project (RPG-2019- 302). For helpful feedback on earlier versions of this paper, the authors would like to thank audiences at the COGITO/University of Seville Sosa Summer School in Epistemology, the Scepticism Research Group hosted by Joshua Thorpe and Angela O'Sullivan, as well as University of Glasgow staff and students who offered feedback at the Junior Honours Reading Party in Stirling. Thanks also to Ju Wang and to anonymous referees at *Synthese* whose feedback improved the paper.

Declarations

Conflict of interest The authors declare no competing interests.

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

- Austin, J. L. (1962). *Sense and sensibilia*. Oxford University Press.
- Baird, B., et al. (2019). The cognitive neuroscience of lucid dreaming. *Neuroscience and Biobehavioral Reviews*, 100, 305–323.
- Ballantyne, N., & Evans, I. (2010). Sosa's dream. *Philosophical Studies*, 148(2), 249–252. <https://doi.org/10.1007/s11098-008-9309-y>.
- Brown, J. (2009). Sosa on scepticism. In E. Sosa (Ed.), *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 143 (3): 397–405.
- Brown, J. A. (2018). *Fallibilism: Evidence and knowledge*. Oxford University Press.
- Bueno, O. (2009). Sosa on skepticism. *Metaphilosophy*, 40(2), 195–202. <https://doi.org/10.1111/j.1467-9973.2009.01574.x>.
- Carter, J. A., & Sosa, E. (2022). Metaepistemology. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy*. Spring 2022, Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2022/entries/metaepistemology/>.
- Carter, J. A., Pritchard, D. & Turri, J. (2018). The value of knowledge. *Stanford Encyclopedia of Philosophy*.
- Carter, J. A. (2023). *Stratified virtue epistemology: A defence*. Cambridge University Press.
- Cassam, Q. (2009). What is knowledge? *Royal Institute of Philosophy Supplements*, 64, 101–120.
- Clark, A. (2007). The twisted matrix: Dream, simulation, or hybrid? In C. Grau (Ed.) *Philosophical Essays on the Matrix*, Oxford University Press.
- Colaço, D., Buckwalter, W., Stich, S., & Machery, E. (2014). Epistemic intuitions in fake-barn thought experiments. *Episteme*, 11(2), 199–212.
- Crowther, T. (2018). Experience, dreaming and the phenomenology of wakeful consciousness. In F. Macpherson & F. Dorsch (Eds.), *Phenomenal Presence*. O.U.P.
- Dennett, D. (1976). Are dreams experiences? *The Philosophical Review*, 85(2), 151–171.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4), 913–929.

- Descartes, R. (1641) (1984). *The philosophical writings of descartes*. Vol. 2. Cambridge University Press.
- Domhoff, G. W. (2005). Refocusing the neurocognitive approach to dreams: A critique of the Hobson versus Solms debate. *Dreaming*, 15(1), 3–20. <https://doi.org/10.1037/1053-0797.15.1.3>.
- Dretske, F. I. (1970). Epistemic operators. *The Journal of Philosophy*, 67(24), 1007–1023.
- Erlacher, D., & Stumbrys, T. (2020). Wake up, work on dreams, back to bed and lucid dream: A sleep laboratory study. *Frontiers in Psychology*, 11(1383), 1–8.
- Ginet, C. (1975). *Knowledge, perception, and memory*. D. Reidel Publishing Co.
- Goldman, A. I. (1979). What is justified belief? In *Justification and knowledge* (pp. 1–23). Springer.
- Greco, J. (2020). Safety in Sosa. *Synthese*, 197(12), 5147–5157. <https://doi.org/10.1007/s11229-018-1863-z>.
- Hobson, A. (2009). REM sleep and dreaming: Towards a theory of protoconsciousness. *Nature Reviews Neuroscience*, 10, 803–813.
- Ichikawa, J. (2008). Scepticism and the imagination model of dreaming. *The Philosophical Quarterly*, 58(232), 519–527.
- Ichikawa, J. (2009). Dreaming and imagination. *Mind & Language*, 24(1), 103–121. <https://doi.org/10.1111/j.1468-0017.2008.01355.x>.
- Jahanshahi, M., & Frith, C. D. (1998). Willed action and its impairments. *Cognitive Neuropsychology*, 15(6–8), 483–533.
- Kubota, Y., Takasu, N. N., Horita, S., Kondo, M., Shimizu, M., Okada, T., Wakamura, T., & Toichi, M. (2011). Dorsolateral prefrontal cortical oxygenation during REM sleep in humans. *Brain Research*, 1389, 83–92.
- LaBerge, S., Baird, B., & Zimbardo, P. G. (2018). Smooth tracking of visual targets distinguishes lucid REM sleep dreaming and waking perception from imagination. *Nature Communications*, 9(1), 3298. <https://doi.org/10.1038/s41467-018-05547-0>.
- Lackey, J. (2009). Knowledge and credit. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 142(1), 27–42. <https://www.jstor.org/stable/27734349>.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Luper-Foy, S. (1984). The epistemic predicament: Knowledge, Nozickian tracking, and scepticism. *Australasian Journal of Philosophy*, 62(1), 26–49.
- Lyons, J. C. (2019). Algorithm and parameters: Solving the generality problem for reliabilism. *The Philosophical Review*, 128(4), 463–509. <https://doi.org/10.1215/00318108-7697876>.
- Malcolm, N. (1959) 2017. *Dreaming*. Routledge. [https://books.google.co.uk/books?hl=en&lr=&id=Ve80DwAAQBAJ&oi=fnd&pg=PT8&dq=Dreaming+\(1959\)+Normal+Malcolm&ots=qsEXscNsmx&sig=eHEBxQJ3wn2yhrJKOT69YVr7Qp0](https://books.google.co.uk/books?hl=en&lr=&id=Ve80DwAAQBAJ&oi=fnd&pg=PT8&dq=Dreaming+(1959)+Normal+Malcolm&ots=qsEXscNsmx&sig=eHEBxQJ3wn2yhrJKOT69YVr7Qp0).
- McGinn, C. (2004). *Mindsight: Image, dream, meaning*. Harvard University Press.
- Nielsen, T. A., & Powell, R. A. (1992). The day-residue and dream-lag effects: A literature review and limited replication of two temporal effects in dream formation. *Dreaming*, 2(2), 67–77. <https://doi.org/10.1037/h0094348>.
- Nozick, R. (1981). *Philosophical investigations*. Belknap Press of Harvard University.
- O'Shaughnessy, B. (2000). *Consciousness and the world*. Oxford University Press.
- Picard-Deland, et al. (2023). The memory sources of dreams: Serial awakenings across sleep stages and time of night. *Sleep*, 46(4), 1–13.
- Pritchard, D. (2005). *Epistemic Luck*. Clarendon.
- Pritchard, D. (2007). Anti-luck epistemology. *Synthese*, 158(3), 277–297.
- Pritchard, D. (2008). Sensitivity, safety, and antiluck epistemology. In J. Greco (Ed.), *The Oxford Handbook of skepticism*. Oxford University Press. <https://doi.org/10.1093/oxfordhdb/9780195183214.003.0021>.
- Pritchard, D. (2009). Safety-based epistemology: Whither now? *Journal of Philosophical Research*. https://doi.org/10.5840/jpr_2009_2.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *The Journal of Philosophy*, 109(3), 247–279.
- Pritchard, D. (2015). Anti-luck epistemology and the Gettier problem. *Philosophical Studies*, 172(1), 93–111.
- Rosen, M. (2012). *Dream pluralism: A philosophy of the dreaming mind*, PhD Thesis.
- Rosen, M. (2021). Dreaming of a stable world: Vision and action in sleep *Synthese*, 1–36.
- Russell, B. (1912). *The problems of philosophy*. Home University Library.
- Sainsbury, R. M. (1997). Easy possibilities. *Philosophy and Phenomenological Research*, 57(4), 907–919.
- Sosa, E. (1999). How to defeat opposition to moore. *Noûs*, (s13), 33. <https://doi.org/10.1111/0029-4624.33.s13.7>.
- Sosa, E. (2005). Dreams and philosophy. *Proceedings and Addresses of the American Philosophical Association* 79 (2): 7–18. <https://www.jstor.org/stable/30046211>.

- Sosa, E. (2007). *A Virtue Epistemology: Apt belief and reflective knowledge, volume I*. Oxford University Press.
- Sosa, E. (2010). How competence matters in epistemology. *Philosophical Perspectives*, 24(1), 465–475.
- Sosa, E. (2021). *Epistemic explanations: A theory of telic normativity, and what it explains*. Oxford University Press.
- Soteriou, M. (2020). Dreams, agency, and judgment. *Synthese*, 197(12), 5319–5334.
- Stumbrys, T., et al. (2012). Induction of lucid dreams: A systematic review of evidence. *Consciousness and Cognition*, 21, 1456–1475.
- Thompson, E. (2014). Waking, dreaming, being: Self and consciousness in neuroscience, meditation, and philosophy. *Waking, dreaming, being*. Columbia University. <https://doi.org/10.7312/thom13709>.
- Turri, J. (2016). Knowledge judgments in Gettier cases. *A Companion to Experimental Philosophy*, 337–348.
- Turri, J., Buckwalter, W., & Blouw, P. (2015). Knowledge and luck. *Psychonomic Bulletin & Review*, 22(2), 378–390.
- Whitely, C. (2020). Aphantasia, imagination, and dreaming. *Philosophical Studies*, 178, 2111–2132.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.

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