

# Śrīharṣa on the Indefinability of Knowledge-Events

## Nilanjan Das

### 1. Introduction

In contemporary epistemology, it is widely recognised that states of knowing are subject to an *anti-luck condition*. Gettier (1963) showed that, in order to know, not only must an agent have a belief that is true and justified, but it must also be free from a certain kind of epistemic luck. But it turned out to be difficult—if not impossible—to articulate what this anti-luck condition on knowledge is, without appealing back to knowledge itself. In the wake of Gettier, many epistemologists proposed non-circular analyses of knowledge: analyses that seek to characterise either knowledge itself (as an epistemological kind) or the application-conditions of the word “know” in knowledge-independent terms. But these analyses are subject to counterexamples. Given this miserable track record, some epistemologists beginning with Williamson (2000) have recommended a “knowledge first” approach to epistemology. As fleshed out by Williamson, the approach involves at least two commitments.<sup>1</sup> First, it is committed to the view that knowledge should be treated as a *sui generis* mental state that need not be (metaphysically or conceptually) decomposable into a *mental* condition like belief and other *non-mental* conditions like truth and reliability. Second, it is committed to the view that knowledge should be treated as *explanatorily fundamental* in epistemology: other normatively significant notions in epistemology, such as evidence and justification, should be analysed in terms of knowledge.

Like history *simpliciter*, the history of philosophy too repeats itself.<sup>2</sup> In this paper, my aim is to explore an analogous episode in the history of Sanskrit philosophy, where a number of epistemologists who were active in South Asia between the second half of the first millennium CE and the first half of the second millennium CE grappled with the problem of epistemic luck. These epistemologists were concerned not exactly with the notion of knowledge, but rather with the notion of *pramā*. Instances of *pramā* are mental events of *learning* or *knowledge-acquisition*. Suppose I look out of my window and undergo a veridical perceptual experience as of there being a hawk on my fence. So, I perceptually learn, i.e., acquire the knowledge, that there is a hawk on my fence. Or, suppose I see a thin trail of smoke emanating from a hill at a distance, and conclude that there is fire on the hill. So, I inferentially learn, i.e., acquire the knowledge, that there is fire on the hill. Here, both my perceptual experience and my inferential judgement are awareness-events (*jñāna*)—

---

<sup>1</sup> For discussion, see McGlynn (2014). For a helpful disambiguation of different ways of construing the “knowledge first” approach, see Ichikawa and Jenkins (2017).

<sup>2</sup> Vaidya (2022) and Williamson (2022) have recently discussed the question of whether there are “knowledge first” approaches in Sanskrit epistemology. I disagree Vaidya’s interpretation of the later Nyāya view as an example of the “knowledge first” approach, because both the early and the later Nyāya epistemologists typically seek to conceptually analyse the notion of knowledge-events (*pramā*) into mental conditions such as awareness-events (*jñāna*) and non-mental conditions such as accuracy (*yāthārthya*). See Das (2021) for more details.

experiences or thoughts—in undergoing which I acquire some knowledge.<sup>3</sup> They are instances of *pramā*. For convenience, we shall call these *knowledge-events*.<sup>4</sup> At least some Sanskrit epistemologists recognised that such knowledge-events are subject to an anti-luck condition: an experience or a judgement that is accurate or true as a matter of luck cannot be a knowledge-event.

Yet, the lessons that they drew were varied. The *optimists*—for example, early Nyāya philosophers like Udayana (10<sup>th</sup>/11<sup>th</sup> century CE)—thought that we could articulate an anti-luck condition on knowledge-events without falling back on the notion of learning or knowledge-acquisition. The *pessimists*—especially the non-dualist Vedāntins like Śrīharṣa (12<sup>th</sup> century CE) and his later followers like Citsukha (13<sup>th</sup> century CE)—rejected this claim: they argued that there was no way of characterising the anti-luck condition on knowledge-events independently of our notion of learning or knowledge-acquisition. Finally, the *ameliorationists*—especially, later Nyāya philosophers like Gaṅgeśa Upādhyāya (14<sup>th</sup> century CE) and his commentators—agreed with the pessimists, but revised the concept of knowledge-events so that it would not only apply to experiences and judgements that are free from epistemic luck, but also to epistemically lucky ones.<sup>5</sup> In earlier work (Das 2021), I have explored the ameliorationist approach in detail. In this chapter, I wish to consider the view of the pessimists: a view developed by Śrīharṣa in *A Confection of Refutation* (*Khaṇḍanakhaṇḍakhādyā*, henceforth the *Refutation*).

Śrīharṣa is a defender of non-dualistic Vedānta, a view that emerges from a certain reading of the Upaniṣads, which are the last part of the Vedic corpus and therefore sometimes called “*vedānta*” (literally, “the end of the Veda”). Non-dualistic Vedāntins accept a form of monism: the view that there is a single entity that ultimately exists (*paramārthasat*), i.e., exists independently of our attitudes like beliefs, desires, judgements and so on: namely, consciousness (*viññāna*). Nyāya and Vaiśeṣika philosophers reject this view. They make two claims. The first is an *ontological* claim: there are many kinds of entities that ultimately exist. The Vaiśeṣika metaphysicians offer a list of six ontological categories—substance (*dravya*), quality (*guṇa*), motion (*karman*), universals (*sāmānya*), ultimate differentiators (*antyavišeṣa*),

---

<sup>3</sup> The term “*jñāna*” is sometimes translated as “cognition.” Typically, philosophers and cognitive scientists take cognitive states to be mental states like beliefs and judgements whose contents can be directly used for theoretical reasoning, verbal reports and controlling action. But some Sanskrit philosophers think that non-conceptual perceptual experiences—which count as “*jñāna*”—aren’t like this. So, it is better to use a term like “awareness.” Even though a construction like “*S* is aware that *p*” in English ascribes a factive mental state insofar as it entails that *p*, philosophers use constructions of the form, “*S* is aware of *o* as being *F*,” which don’t always entail that *o* is *F*. I will use the latter kind of construction stipulatively: on my view, *S* is aware of *o* as being *F* if and only if *S* perceives/judges/suspects *o* to be *F*. None of these attitudes entail that *o* is *F*, even though they entail that *o* exists (which is an assumption shared by Śrīharṣa’s Nyāya-Vaiśeṣika interlocutors).

<sup>4</sup> Not all states of knowing are knowledge-events. First, knowledge-events are awareness-events, i.e., occurrent mental states like experiences and thoughts. By contrast, states of knowing can be occurrent as well as dispositional. Second, not all states of knowing are states through which we learn or acquire knowledge; for example, remembering is a paradigmatic state of knowing, but we needn’t learn that *p* by remembering that *p*.

<sup>5</sup> See Matilal (1986); Ganeri (2016) closely follows Matilal’s exposition. A complete translation of Śrīharṣa’s *A Confection of Refutation* (*Khaṇḍanakhaṇḍakhādyā*) can be found in Jha (1913/1986) and a partial one in Granoff (1978). For discussions of Śrīharṣa’s philosophical views more generally, see Granoff (1978), Phillips (1999), and Ram-Prasad (2002) and Das (2018).

and inherence (*samavāya*)—that is supposed to exhaust everything that ultimately exists. The second is an *epistemological* claim: our ordinary methods of knowing can help us know various facts about ultimately existent objects. The Nyāya epistemologists (henceforth the Naiyāyikas) offer a list of four methods of knowing (*pramāṇa*)—perception (*pratyakṣa*), inference (*anumāna*), analogy (*upamāna*), and testimony (*śabda*)—which are supposed to give us epistemic access to the constituents of ultimate reality. In the *Refutation*, Śrīharṣa dissents from both these claims: he refutes the definitions (*lakṣaṇa*) that the Nyāya and Vaiśeṣika thinkers propose for their preferred ontological and epistemological categories. His aim is to show that neither the polycategorical ontology that the Vaiśeṣikas defend, nor the Nyāya story about how we gain epistemic access to objective features of reality, is defensible. While defending this claim, Śrīharṣa argued that knowledge-events are *indefinable* precisely because there is no good way of articulating the anti-luck condition on knowledge-events.

What is novel about Śrīharṣa’s treatment of this problem is that he takes his arguments for indefinability of knowledge-events to also undermine the view that knowledge-events form a unified, *sui generis* kind of mental events. This, as I shall show, has to do with his commitment to a view about attributions of knowledge-events.

*The Epistemic Priority Thesis.* Knowledge-events are *epistemically prior* to other non-factive mental states and events: when we are trying to determine whether an agent has undergone a knowledge-event, we don't initially ascribe to them some other non-factive mental event, and then check if that event meets some further conditions (like truth or reliability) necessary for it to count as a knowledge-event; rather, we treat certain mental events by default as knowledge-events until a defeater comes along.

Śrīharṣa argues that this thesis—when taken in conjunction with the arguments for the indefinability of knowledge-events—should give us reason to doubt whether our ordinary attributions of knowledge-events are reliably tracking any unified, *sui generis* kind of mental events. I argue that Śrīharṣa’s arguments can be extended to states of knowing more generally. Thus, his arguments not only undermine the idea that knowledge is analysable, but also one of the key theses of “knowledge first” epistemology, i.e., that knowledge is a *sui generis* mental state.

## 2. Śrīharṣa’s Argument

Śrīharṣa’s argument against Nyāya epistemology exploits a tension between two commitments about knowledge-events that many Naiyāyikas endorse. Many (if not all) Naiyāyikas accept the following definition of knowledge-events.

*The Nyāya Definition of Knowledge-Events.* An awareness-event (*jñāna*)—an experience or a thought—is a knowledge-event if and only if it satisfies two conditions:

- (i) It is a firsthand awareness (*anubhava*): it is not a recollective awareness.
- (ii) It is accurate (*yathārtha*): it represents the world the way it is.

This characterisation of knowledge-events imposes two necessary conditions on knowledge-events. The first condition—call it the *non-mnemic condition*—says that a knowledge-event cannot be a recollective awareness. This may seem surprising, but it follows from the Nyāya conception of recollection (*smaraṇa*). Recollection, according to Naiyāyikas, just is *information retrieval*, i.e., a process of retrieving bits of information one has acquired through earlier non-recollective awareness-events. Understood in this way, recollection is not a *generative* source of knowledge: when we undergo a recollective awareness, we simply retrieve what we already were aware of; we don't thereby acquire any piece of knowledge that we didn't already possess. This implies that recollective awareness-events aren't knowledge-events. The second condition—call it the *accuracy condition*—says that a knowledge-event must be accurate. Different Naiyāyikas explain this condition differently. But the basic idea is clear: if an experience or thought misrepresents the way the world is, then it cannot be an event of knowledge-acquisition. While both these conditions can plausibly be treated as necessary conditions on knowledge-events, it is not obvious that they are jointly sufficient.

The *Nyāya Definition of Knowledge-Events* conflicts with another commitment of the Naiyāyikas:

*Nyāya Infallibilism.* An awareness-event is a knowledge-event only if it is produced by a totality of causal conditions (*kāraṇasāmagrī*) that could not have given rise to an inaccurate awareness-event.<sup>6</sup>

Take three examples that Śrīharṣa considers in this connection.

*Guesswork.* You place a few shells in your fist and ask me, “How many shells are there in my hand?” I have no idea. On a hunch, I judge, “There are five shells in that fist.” So, I say out loud, “There are five.” My awareness is correct: you have exactly five shells in your hand.

*Mist and Fire.* I look at a hill and see what looks like smoke emerging from it. So, I judge that there is smoke on the hill. I am wrong: all I see is a wisp of mist. I had previously observed (in kitchens, etc.) that smoke is always accompanied by fire. On the basis of those observations, I had judged that, wherever there is smoke, there is

---

<sup>6</sup> It is worth distinguishing *evidential* infallibilism from *causal* infallibilism. Evidential infallibilism says that an agent can come to know that *p* only if the evidence on the basis of which they believe (or judge) that *p* entails that *p*. The Nyāya thinkers like Udayana are not evidential infallibilists: they think that we can come to learn inductive generalisations, or facts on the basis of testimony, even though our evidence for judging the relevant contents to be true doesn't decisively rule out the possibility that they are false. However, they are causal infallibilists: they think that, whenever an agent comes to know that *p*, their awareness is formed on the basis of causal conditions that couldn't lead to an accurate judgement. Thanks to Mark Siderits for pressing me on this point.

fire. Now, I remember that generalisation. So, I conclude that there is fire on the hill. My judgement turns out to be true: there is fire on the hill.

*Horns and Cows.* From a distance, I see an animal with horns. Earlier, I had observed many cows with horns. On the basis of these observations, I *falsely* judged that all animals with horns are cows. Now, I recall that generalisation. So, I conclude that the animal is a cow. My judgement is true: the animal is a cow.

In these cases, it seems that I could easily have made a false judgement. Thus, my judgement is true as a matter of luck. According to *Nyāya Infallibilism*, therefore, these cannot be knowledge-events: the totality of causal conditions that yield these judgements could easily have given rise to an error. Yet, according to the *Nyāya Definition of Knowledge-Events*, the judgements that I make in these cases of epistemic luck are knowledge-events: they are accurate, non-recollective awareness-events. This is the tension. Call this the *problem of epistemic luck*. Śrīharṣa argues that this tension cannot be resolved by any definition of knowledge-events. Any satisfactory (and therefore non-circular) definition of knowledge-events will have to include an anti-luck condition that doesn't appeal back to the notion of learning or knowledge-acquisition itself. But there is no such anti-luck condition.

In what follows, I will reconstruct Śrīharṣa's argument by paying closer attention to his text.

### 3. Response 1: Infallibilism

Śrīharṣa's argument against the Nyāya conception of knowledge-events primarily targets a definition given by Udayana in his *Garland of Definitions (Lakṣaṇamālā)*: "A knowledge-event is a firsthand awareness of the way something is (*tattvānubhūtiḥ pramā*)" (LM 9.1). This is a version of the *Nyāya Definition of Knowledge-Events*. In an initial barrage of arguments, Śrīharṣa claims that neither the notion of firsthand awareness (*anubhūti*) nor the notion of the way something is (*tattva*) can be adequately characterised. But then he turns to the definition as a whole. Here, his argument hinges on a case like *Guesswork* (KKh 383.20-384.11). In that case, when you ask, "How many shells do I have in my fist?" I think, "There are five." My awareness is accurate: there are exactly five shells in your hand. And it is firsthand: since its content isn't derived from any earlier firsthand awareness by recollection, it is non-recollective. Yet, it is obvious that this cannot be a knowledge-event: when I judge that you have five shells in your hand, I don't thereby come to know this. This is a problem for the *Nyāya Definition of Knowledge-Events*.

Śrīharṣa considers two possible Nyāya responses to this problem. The first response seems to be the most natural: the way I arrive at my judgement in this case could easily have led me astray. Since I was randomly making up my mind about the number of shells in your hand, I could easily have made a false judgement about how many shells you had in your fist. Thus, Śrīharṣa's Nyāya interlocutor could revise their conception of knowledge-events.

*The Revised Nyāya Definition of Knowledge-Events.* An awareness-event is a knowledge-event if and only if it satisfies three conditions:

- (i) It is a firsthand awareness-event: it is not a recollective awareness-event.
- (ii) It is accurate: it represents the world as it is.
- (iii) It arises from a totality of causal conditions that couldn't give rise to any inaccurate or erroneous awareness.

This account preserves the spirit of *Nyāya Infallibilism*: it entails that an awareness can be a knowledge-event only if its causal ancestry guarantees its accuracy. But we might wonder: What is it about the causal ancestry of knowledge-events, which explains why the knowledge-events are infallible in this way? There are two different answers to this question—defended by Naiyāyikas and Bhāṭṭa Mīmāṃsakas respectively—which, in turn, imply two distinct ways of fleshing out the *Revised Nyāya Definition of Knowledge-Events*.

#### 4.1 Two Kinds of Infallibilism

Let us begin by distinguishing two kinds of infallibilism. The Bhāṭṭa Mīmāṃsakas—Kumārila Bhaṭṭa (7<sup>th</sup> century CE) and his commentators—argued that the normal or default (*autsargika*) state of any awareness is to be a knowledge-event. An awareness will deviate from this normal or default state—and become inaccurate—just in case the causal conditions are *abnormal*, i.e., just in case the causal conditions that give rise to the awareness include certain epistemic defects (*doṣa*) that are normally absent. Thus, what is necessary for a knowledge-event to arise is the absence of epistemic defects. This yields:

*Defect Infallibilism.* For any kind *K* of knowledge-events (perceptual, inferential, testimonial and so on), the totality of causal conditions that are necessary for producing any instance of *K* must include the absence of certain positive conditions—epistemic defects—which guarantee the inaccuracy of the resulting awareness.<sup>7</sup>

Nyāya and Vaiśeṣika thinkers such as Jayanta Bhaṭṭa (9<sup>th</sup> century CE), Śrīdhara (10<sup>th</sup> century CE), and Udayana disagree with this claim (NM I 442.13-444.2; NK 516.1-2; NKA 211.1-220.2).<sup>8</sup>

---

<sup>7</sup> This is sometimes called the theory of intrinsic knowledgehood (*svataḥprāmāṇya*) with respect to production (*utpatti*): on this view, a knowledge-event arises simply from the normal causes that give rise awareness-events of a certain kind (as long as those causes are non-defective); no positive factors such as epistemic virtues are necessary. Kumārila Bhaṭṭa's commentators—Umveka Bhaṭṭa, Sucarita Miśra and Pārthasārathi Miśra—defend different varieties of the theory of intrinsic knowledgehood in their commentaries on Verse 47 in *The Detailed Commentary in Verse (Śloka-vārttika)* ad *Mīmāṃsāsūtra* 1.1.2 (ŚVT 54.1-17, ŚVK 90.13-25, NRK 45.7-20).

<sup>8</sup> The disagreement revolves around the epistemic status of testimony. Both the Bhāṭṭas and the Naiyāyikas accept the status of testimony as a *sui generis* source of knowledge, and accept the Veda—a text regarded as authoritative by all Brahminical thinkers—to be a source of knowledge with respect to ritually correct action (*dharma*). But Bhāṭṭas think that testimony can be a source of knowledge insofar as certain epistemic defects belonging to the speaker—which make testimony inaccurate—are absent. By contrast, the Nyāya thinkers think that testimony can be a source of knowledge only insofar as the speaker not only lacks these defects but also

Udayana notes that the question of whether *Defect Infallibilism* is true can only be settled by looking at specific sources of knowledge, such as inference (NKA 215.1-216.2). Take cases of defective inference like *Mist and Fire* and *Cows and Horns*. In *Mist and Fire*, the fire that I infer is the target property (*sādhya*), and the hill to which I ascribe that fire is the site (*pakṣa*) of the inference. The inferential mark (*liṅga*)—the perceived feature of the hill on the basis of which I infer the fire—is smoke. But, since this smoke is not in fact present in the site, my judgement that there is smoke on the hill is erroneous. It is this epistemic defect that prevents me from acquiring knowledge in this case. Next, consider *Cows and Horns*. Here, the target property that I infer is the cowhood, while the site of the inference is the animal before me. The horns of the animal I see are the inferential mark on the basis of which I infer its cowhood. In this case, even though the inferential mark is present in the site, the target property doesn't invariably accompany or *pervade* the inferential mark; for not all animals that have horns are cows. This, in turn, makes my judgement that whatever has horns is a cow false. It is this epistemic defect which, in this case, prevents me from acquiring knowledge.

Udayana claims that the mere absence of such epistemic defects isn't sufficient to generate inferential knowledge-events. In good cases of inference, other positive conditions must be present: the agent's inferential judgement must be based on prior judgements that accurately indicate (i) that the relevant inferential mark is present in the site and (ii) that it is invariably accompanied or pervaded by the target property. This suggests that the inferential knowledge-events causally depend on certain positive conditions over and above the absence of the epistemic defects which generate awareness-events that aren't knowledge-events. This diagnosis seems to be true of other kinds of knowledge-events, such as those derived from perception, analogy and testimony. The general view, then, is that knowledge-events require not only the absence of epistemic defects but also the presence of certain epistemic virtues (*guṇa*) that guarantee their accuracy. This suggests:

*Virtue Infallibilism.* For any kind *K* of knowledge-events (perceptual, inferential, testimonial and so on), the totality of causal conditions that are necessary for producing any instance of *K* must include certain positive conditions—epistemic virtues—that guarantee the accuracy of the resulting awareness.<sup>9</sup>

---

possesses certain epistemic virtues (*guṇa*) that makes them trustworthy (*āpta*). This disagreement bears indirectly on the epistemic status of Veda. Since Bhāṭṭas regard the Veda as an authorless text, they take the Veda to be a source of knowledge precisely because it is not vitiated by the defects of an author. By contrast, the Naiyāyikas like Jayanta and Udayana regard the Veda as the creation of an omniscient, omnipresent God-like being called Īśvara: not only does Īśvara lack the epistemic defects that make testimony inaccurate, but also possesses certain epistemic virtues that make Him trustworthy.

<sup>9</sup> Contemporary virtue epistemologists treat epistemic or intellectual virtues either as faculties or as traits that promote some intellectual good. Virtue reliabilists, like Sosa (1991), think of intellectual virtues as faculties or qualities that helps the agent maximize their surplus of true beliefs over false ones. In contrast, virtue responsibilists, like Zagzebski (1996), treat intellectual virtues as traits of character that promote intellectual flourishing. Both camps, however, treat virtues as dispositions. However, the Nyāya epistemologists typically treat epistemic virtues as causal factors which are necessary for the production of knowledge-events, but may or may not be dispositional properties.

It is worth understanding the general picture that emerges from this discussion.

For both virtue and defect infallibilists, what distinguishes knowledge-events from other kinds of awareness-events—such as error and uncertainty—is that the causal history of knowledge-events is different from the causal history of other kinds of awareness-events. According to the defect infallibilists, the causal history of knowledge-events is characterised by the absence of certain positive conditions—the epistemic defects—whose presence would make the resulting awareness-events inaccurate. By contrast, according to the virtue infallibilists, the causal history of knowledge-events involves certain positive conditions—the epistemic virtues—over and above the absence of epistemic defects. These positive conditions—insofar as they guarantee the accuracy of the resulting knowledge-events—explain why these are accurate.

#### 4.2 Śrīharṣa on Virtue Infallibilism

A virtue infallibilist could say that, in cases of epistemic luck like *Guesswork*, these epistemic virtues go missing. In *Guesswork*, my guess that you have five shells in your fist cannot be a knowledge-event even if it is accurate: what generates my awareness in that case is a cognitive process that could easily lead to inaccurate judgements in other cases, so the causal ancestry of that awareness doesn't include any epistemic virtue that could guarantee its accuracy.

Śrīharṣa is unhappy with *Virtue Infallibilism*. His unhappiness comes across clearly in his treatment of an infallibilist response to *Guesswork*.

[The opponent:] One should insert the qualifying clause ‘produced by causal conditions that don't err (*avyabhicārin*)’ [in the definition that a knowledge-event is a firsthand awareness of how something is].

[Reply:] No. For, then, the expression “how something is” (*tattva*) will be useless. And you cannot accept even an accidentally fact-conforming (*kākatālīyasaṃvāda*) awareness to be produced by a totality of causal conditions that are shared with erroneous (*vyabhicārin*) awareness-events. This is because, then, since an erroneous awareness would not differ with respect to its causal conditions [from accurate ones], there would be the undesirable consequence that they would be accurate. For it is not the case that its accuracy is causeless, since, then, a problem of overgeneration would follow in virtue of the absence of a determining factor. Given that this [awareness] is not erroneous, it must necessarily be said that its causal conditions are indeed invariably connected to non-erroneous awareness-events. (KKh 387.1-7)

Śrīharṣa is making two claims in this passage. The first is relatively trivial: adding clause (iii) to *Nyāya Definition of Knowledge-Events* makes the accuracy condition laid down by (i) redundant. We would no longer need to define a knowledge-event as a firsthand awareness of how something is. The second point is more important: the reasoning that leads Udayana to



*Virtue Infallibilism* should also compel him to say that accidentally accurate awareness-events must arise from certain accuracy-guaranteeing epistemic virtues. Let's unpack this claim.

Udayana's argument for *Virtue Infallibilism* depends on a principle about causation (NKA 211.1-213.2): namely, that if effects of a certain kind *E* are a species of effects of a more general kind *E\**, then the totality of causal conditions which are necessary for producing any instance of *E* must include causal conditions that are not included amongst the causal conditions which are necessary for producing any arbitrary instance of *E\**. Now, consider knowledge-events of a particular kind (perceptual, inferential, etc.). Since these are a subspecies of awareness-events of the relevant kind, the causal conditions necessary for producing knowledge-events of that kind must include some further conditions that are not included amongst the causal conditions which are necessary for producing awareness-events of that kind. These special causal conditions—which distinguish the causal history of knowledge-events from mere awareness-events of the relevant kind—are the epistemic virtues.

The problem is this. Just like knowledge-events, *accurate* awareness-events of a particular kind (perceptual, inferential, etc.) too are a subspecies of awareness-events of that kind. So, given Udayana's principle about causation, they too must be produced by certain specific causal conditions that are not included amongst the causal conditions which are necessary for producing awareness-events of that kind. For, if these accurate firsthand awareness-events were produced by the exactly same causal conditions that also produce other awareness-events of the relevant kind (including inaccurate ones), then even those other inaccurate awareness-events would be accurate. This reasoning, if sound, should yield the result that the causal history of accurate awareness-events must be different from that of inaccurate awareness-events. This is significant: it implies that, even awareness-events that are *accidentally accurate* must be produced by at least some causal conditions that play no role in the production of inaccurate awareness-events. So, accidentally accurate awareness-events—such as my lucky judgement that you are holding five shells in your fist—must also arise from causal conditions that could not give rise to any inaccurate or erroneous awareness. So, the infallibilist response fails.

A general concern about this argument is that it relies heavily on the assumption that *Virtue Infallibilism* can only be motivated by relying on Udayana's principle about causation. This assumption can be rejected. But Śrīharṣa does have a second response against *Virtue Infallibilism* that doesn't make this assumption (KKh 389.11-16).

The response depends on two closely related examples. The first is *Mist and Fire*, the case where I misperceive mist on a hill as smoke and infer the presence of fire on the hill on the basis of that smoke. The second is:

*Extended Mist and Fire.* I look at a hill and see what looks like smoke emerging from it. So, I judge that there is smoke on the hill. I am wrong: all I see is a wisp of mist. I

had previously observed (in kitchens, etc.) that smoke is always accompanied by fire. On the basis of those observations, I had judged that, wherever there is smoke, there is fire. Now, I remember that generalisation. So, I conclude that there is fire on the hill. My judgement is true: there is *both* smoke *and* fire on the hill.

The difference between this case and *Mist and Fire* is that, in the latter, my judgement that there is smoke on the hill could be false; in this case, it is true. But, in both these cases, the inferential mark on the basis of which I arrive at my judgement is defective: what I perceive as smoke is in fact not smoke; it is mist. Yet, the inferential judgement that I make is accurate: there is in fact fire on the hill. Since this judgement is an accurate firsthand awareness, the defender of the *Nyaya Definition of Knowledge-Events* is committed to the claim that the judgement is a knowledge-event. Even though the problem here is structurally the same as in the case of *Guesswork*, this pair of examples is a bit more robust. First, in both these cases, my judgement is uncontroversially a state of certainty. So, the two cases cannot be ruled out by appealing to a response that is available in the case of *Guesswork*: since in that case I merely guess that you have five shells in your fist and there is nothing to make me certain about this, the relevant awareness is not a state of certainty (KKh 384.7-10). Second, in *Extended Mist and Fire*, the kinds of epistemic virtues that Naiyayikas typically regard as necessary for knowledge-events are indeed present: since there is smoke on the hill and smoke invariably accompanies fire, my final inferential judgement is based on my correct judgement that there is smoke on the hill, and my correct judgement that wherever there is smoke, there is fire. So, at least *Extended Mist and Fire* cannot be excluded from the scope of knowledge-events by appealing to the idea that certain epistemic virtues that are necessary for the production of inferential knowledge-events have gone missing.

Of course, the Naiyāyika could indeed insist that the epistemic virtue that is required for the production of inferential knowledge-events is not merely a correct judgement about the inferential mark's being present in the site or its being invariably accompanied by the target property. Something stronger is needed: namely, that the agent must have *learnt* or come to *know* that the inferential mark is present in the hill, and that it is pervaded by the target property. But the Naiyāyika cannot, strictly speaking, appeal to such an epistemic virtue within their definition of knowledge-events without making that definition blatantly circular. So, *Virtue Infallibilism* fails.

### 4.3 Śrīharṣa on Defect Infallibilism

Given the failure of *Virtue Infallibilism*, the Nyāya infallibilist could always switch to *Defect Infallibilism*: they could argue that an accurate firsthand awareness is a knowledge-event just in case it arises from a totality of causal conditions that doesn't involve any epistemic defects. In *Mist and Fire* and the extended version thereof, my final inferential judgement arises from a defective awareness of an inferential mark: what I take to be smoke on the hill is in fact mist. However, as Śrīharṣa himself emphasizes elsewhere, it is difficult to articulate what epistemic defects are without appealing back to the notion of knowledge-events.

This argument occurs in a different context: while discussing the proposal that knowledge-events are awareness-events that are free from disconfirmation (*avisamvādirjñāna*) (KKh 430.13-16). Under one interpretation of this proposal, it is equivalent to an indefeasibility analysis of knowledge-events: the view that knowledge-events are awareness-events that cannot be defeated (*bādhita*)—indicated to be inaccurate—by other awareness-events. As Śrīharṣa rightly notes (KKh 432.2-7), this account will lead to undesirable consequences unless it is properly restricted: erroneous experiences and thoughts can count as *defeaters* (*bādhaka*) against a knowledge-event insofar as they can indicate, from the subject’s perspective, that the relevant knowledge-event is inaccurate. So, the only kind of defeaters that knowledge-events should be invulnerable to have to be knowledge-events themselves. This leads the opponent to say that a knowledge-event is an awareness that cannot be defeated by an awareness that is produced by a non-defective method or instrument.

Śrīharṣa’s response falls into two parts. First, he notes that this property of *being an awareness that is produced by a non-defective instrument* (*aduṣṭakaraṇajanyatva*) can itself be treated as a defining characteristic of knowledge-events (KKh 432.7). But, unless the opponent can independently specify what defectiveness consists in, the proposal in question will remain uninformative (KKh 432.7-8). Recall how we understood *Defect Infallibilism*: we took epistemic defects to be causal conditions of awareness-events that guarantee, or at least are conducive to, their being inaccurate. Thus, the opponent initially claims that “defectiveness is a specific characteristic which is conducive to an opposite awareness, and which resides in the causal conditions of that awareness” (KKh 432.8-10). Śrīharṣa’s response is simply that we don’t know what “opposite awareness” (*viparītajñāna*) means here (KKh 432.10-11). If an opposite awareness were just an inaccurate awareness, then the proposal under consideration would be equivalent to the view that a knowledge-event is simply an awareness which is produced by a totality of causal condition that could not have led to an inaccurate awareness. But this view would inherit the two problems for *Virtue Infallibilism*. First, given Udayana’s principle about causation, even accidentally accurate awareness-events could be taken to have been produced by causal conditions that couldn’t have led to any inaccurate awareness. Second, in *Extended Mist and Smoke*, both the judgements that my inferential judgement are based on are accurate. Even though I misperceive the mist as smoke, my judgment that there is smoke on the hill is accurate. Moreover, I also know that fire invariably accompanies smoke. Thus, despite the defectiveness of the inferential mark, the presence of these epistemic virtues guarantees that the resulting inferential judgement will be accurate. So, it is unclear that *Defect Infallibilism* can be successful in ruling out accidentally accurate awareness-events like my inferential judgement in *Extended Mist and Fire*.

In their desperation, the opponent could make one final gambit. They could argue that any epistemic defect that prevents an awareness from being a knowledge-event is *just that*: it is a causal factor that is conducive to the production of awareness-events that aren’t knowledge-events. Thus, in this context, the kind of opposite awareness that an epistemic defect yields is, quite simply, an awareness that isn’t a knowledge-event. But this makes the definition of knowledge-events circular and, therefore, uninformative.

[The opponent:] Knowledge-events are excluded by the expression “opposite.”

[Reply:] No. For that is what is being defined. Given that the nature of those knowledge-events, insofar as it is excluded from what isn't a knowledge-event, remains unapprehended so far, what is it from which the exclusion is to be apprehended [by means of the expression “opposite”]? Without the awareness of that [nature of knowledge-events] as excluded from what are not knowledge-events, it is impossible for there to be awareness of that [nature] as excluded from what are not knowledge-events. So, there will be either faults of reflexive dependence and symmetric dependence, or a regress. (KKh 432.13-18)

According to the proposed definition, knowledge-events are just awareness-events which arise from causal conditions that are free from epistemic defects. According to the last characterisation of epistemic defects, these are just causal factors that are favourable to the production of awareness-events that aren't knowledge-events. So, in order to grasp what the distinguishing features of knowledge-events are on the basis of the proposed definition, we need to antecedently know what distinguishes knowledge-events from awareness-events that aren't knowledge-events. Under one interpretation, this solution makes the proposed definition straightforwardly circular: as Śrīharṣa says, there is either a fault of reflexive dependence (where knowledge-events are characterised in terms of themselves) or a fault of symmetric dependence (where knowledge-events are characterised in terms of defects, while defects are characterised in terms of knowledge-events). Under another interpretation, this solution leads to a regress: since this definition—in order to be informative—requires us to antecedently know what distinguishes knowledge-events from awareness-events that aren't knowledge-events, we would need a further definition that informs us of this distinguishing characteristic; if that definition appeals to the notion of an epistemic defect, we will need another definition, and then another, and so on *ad infinitum*.

The upshot is this. The infallibilist strategy—favoured by virtue infallibilists like Udayana as well as defect infallibilists like Kumāriḷa and his commentators—cannot solve the problem of epistemic luck: it will either fail to rule out cases like *Extended Mist and Fire* or will make our characterisation of knowledge-events circular. This dilemma has some similarity to the generality problem for process reliabilists, i.e., the problem of individuating the causal process that give rise to a belief with the right fineness of grain.<sup>10</sup> If the causal process is individuated too coarsely, then it will be less reliable, making the resulting belief unjustified; if it is individuated too finely, then it will be more reliable, making the belief justified. Here, too, the problem is that of specifying the set of causal conditions that gives rise to a knowledge-event. But, unlike in the case of the generality problem, here the fineness of grain is not at issue. The problem, rather, is this. If the accuracy-guaranteeing causal conditions that give rise to a knowledge-event were specified without reference to knowledge-events, then those causal conditions might sometimes generate accidentally accurate awareness-events.

---

<sup>10</sup> See Conee and Feldman (1998) for a seminal treatment of this problem.

So, the infallibilist would have failed to solve the problem of epistemic luck. By contrast, if those causal conditions were specified with reference to knowledge-events, then the infallibilist's characterisation of knowledge-events would be circular.

#### 4. Response 2: Inaccuracy

A less conciliatory response to cases like *Mist and Fire* will be to say that, in these cases, my judgement that there is fire on the hill is not accurate at all. This response is conservative: it tries to reconcile the original *Nyāya Definition of Knowledge-Events* with *Nyāya Infallibilism* by claiming that the putative “accidentally accurate” awareness-events aren't accurate at all. Śrīharṣa states this response as follows:

Since that awareness arises from a defective instrument, its intentional object [e.g., the inferred fire] is in fact distinct from the actual target property, and so on. (KKh 389.16-18)

The argument is this. In *Mist and Fire*, when I infer that there is fire on the hill, the fire that I infer is not the actual fire that is present on the hill; rather, it is a fire that accompanies the smoke that I erroneously take to be present on the hill when I see mist. So, the content of my inferential judgement is not that:

(1) The hill contains fire.

Rather, it is expressed by:

(2) The hill contains a particular fire that pervades the smoke I saw.

Since no such fire exists on the hill, my inferential judgement is false.

Śrīharṣa's initial response is concessive (KKh 389.18-19). Even if we grant that the particular fire that I infer to be present on the hill is distinct from the actual fire that is present on it, my inferential judgement could nevertheless be partially accurate insofar as it ascribes the presence of something of that kind—*some* instance of firehood—to the hill. So, with respect to that part of its content, the inferential judgement would still be a knowledge-event according to the *Nyāya Definition of Knowledge-Events*. This response won't be effective against the opponent (KKh 289.19-21). They could argue that this is not really a part of the content of the inferential judgement: when we infer the presence of a particular fire on the hill, we do not thereby also separately draw the more general conclusion that *some* instance of firehood is present on the hill.

In response, Śrīharṣa appeals to a type of case that he hasn't discussed so far: *Horns and Cows* (KKh 390.3-4; KKhPV 393.7-12). In that scenario, I infer that the animal before me is a cow on the basis of the fact that it has horns. On a natural construal of this case, what I inferentially ascribe to the animal is a universal, i.e., cowhood. While there can be multiple

instances of fire, universals like cowhood are unitary: there is only a single property of cowhood that is shared by all and only cows. Thus, even though the opponent may claim that the particular fire that I infer in *Mist and Fire* isn't really present on the hill, they cannot claim that the cowhood that I infer is in fact absent from the animal before me. If the animal is present before me is a cow, it will indeed possess cowhood. This, in turn, will make my inferential judgement accurate. But, given that the inference is based on a defective inferential mark, i.e., the possession of horns, which isn't invariably accompanied by the target property, i.e., cowhood, the resulting inferential judgement can only be true as a matter of luck. So, it cannot be a knowledge-event. Thus, the opponent's strategy of explaining away the accuracy of epistemically lucky inferential judgements won't succeed.

Śrīharṣa pre-empts two possible responses from his Nyāya opponent (KKh 390.4-6). First, they could claim that the inferred property of cowhood in *Horns and Cows* is an imaginary or conceptually constructed (*kalpita*) property which is distinct from the real cowhood that is present in the animal before me. Or, they could claim that my inferential judgement in *Horns and Cows* ascribes cowhood to the animal before me by means of an imaginary or conceptually constructed relation of inherence (*samavāya*). But this response—as Śrīharṣa correctly notes—will be costly for his Nyāya opponent. The Naiyāyika is committed to a *realist theory of content*: the view that only ultimately existent particulars and properties can serve as intentional objects of awareness. This commits them to a *misplacement theory of error* (*anyathākhyātivāda*): the view that, even in cases where an object *o* is presented or represented to an agent as *F* when it is in fact not *F*, both *o* and *F*-hood, as well as the relation between the two, should ultimately exist.<sup>11</sup> If the Naiyāyika were to claim that, in such cases of error, either *F*-hood or the relation by which *F*-hood is ascribed to *o* in fact doesn't ultimately exist, then they will be jettisoning their own realist theory of content. In effect, they would be embracing an alternative theory sometimes attributed to the Mādhyamika Buddhists: the view that, in cases of error, non-existent particulars and properties can appear. This would undermine the realist ambitions of Nyāya epistemology. Given that it will often be difficult to subjectively distinguish erroneous awareness-events directed at conceptually constructed objects from ordinary knowledge-events that track how things are independently of our attitudes, we may not be able to show that our awareness-events do in fact constitute knowledge-events about ultimately existent particulars and properties. By contrast, the *misplacement theory of error* implies that, even in cases of error, our awareness-event do in fact latch on to ultimately existent particulars and properties.

The second response that Śrīharṣa entertains is somewhat different: the claim is no longer that the inferential judgement in *Horns and Cows* ascribes an imaginary or conceptually constructed property or relation, but rather that it ascribes identity with cowhood to some other property present in the animal before me (KKh 390.6-7). So, the content of my inferential judgement is not that:

(3) This animal is a cow.

---

<sup>11</sup> For discussions of this view, see Dasti (2012) and Vaidya (2013).

Rather, it is that:

- (4) The property which pervades the possession of horns and which is present in this animal is identical to cowhood.

Since the property in question isn't cowhood (given that cowhood doesn't invariably accompany the possession of horns), my inferential judgement will be false. This will allow the Naiyāyika to maintain that so-called "accidentally accurate" awareness-events aren't accurate after all.

Śrīharṣa's response is two-fold (KKh 390.7-11). First, even if (4) expresses the content of the inferential judgement, that content still implies—and has as its part—the content that the animal in question possesses cowhood. So, with respect to that part of its content, the inferential judgement will remain accurate and therefore would qualify as a knowledge-event. But this is not a result that the Naiyāyika wants. Second, we should be sceptical of the proposal that (4) actually represents the content of the inferential judgement that I make in *Horns and Cows*. As Śrīharṣa notes, the causal conditions that give to my inferential judgement are in fact part of a cognitive process that necessarily yields ascriptions of an association (*samsarga*), i.e., a relation other than identity, rather than those of identity. If I judge that an object *o* possesses a feature *X*, and that whatever possesses *X* possesses *Y*, then those two judgements should (under favourable circumstances) give rise to the inferential judgement that the object *o* possesses *Y*. Here, the relation between *o* and *Y* is an association, not an identity. The same pattern of inference takes place in *Horns and Cows*.

To bolster this point, Śrīharṣa considers a variant of *Horns and Cows*. In this version of the example, I not only make an inference, but also express my reasoning out loud to an audience: "The animal before me is a cow, because it has horns; whatever has horns is a cow, e.g., the animal in my byre." A linguistically competent hearer, who listens to my utterance and takes it at face-value, should judge that the animal before me is a cow and therefore ascribe an association with cowhood to the animal before me. For, in this scenario, the content of the sentence I utter *just is* an association between the referents of the words that are part of that sentence, i.e., an association between the animal picked out by "the animal before me" and the property of cowhood picked out by "is a cow." So, it would be perverse for a trusting hearer who understands my uttered sentence to make a judgement with a different content. Yet, given that my argument is based on my earlier defective inference, it cannot give rise to any (testimonial) knowledge-event in the hearer. Thus, this would be a case where the hearer will make an accurate judgement that fails to be a knowledge-event.

Of course, the opponent could stubbornly insist that, even in this case, the content of the hearer's judgement is false because it involves the misascription of identity with cowhood to some other property. But, then, this could lead to a problem of overgeneration elsewhere. For, now, the opponent could treat pretty much any case of error as involving a misascription of identity. If they were to do so, they would be blurring a distinction between two kinds of

error: awareness-events that are erroneous in virtue of misascribing identity and awareness-events that are erroneous in virtue of misascribing an association. For example, when I mistake mother-of-pearl to be silver, what I misascribe to the mother-of-pearl is an association with silverhood. By contrast, when I misidentify a man on the street to be Devadatta, what I misascribe to the man is identity with Devadatta. The distinction between these two kinds of error is borne out by introspective evidence. In the first case, I introspectively judge myself to have ascribed the association with silverhood to the object before me. In the other, I introspectively judge myself to have ascribed the identity with Devadatta to the man before me. To claim that all error involves the misascription of identity is to blur the distinction in content between these two kinds of error. For example, the defender of the view that my inferential judgement in *Horns and Cows* would also be committed to saying that my introspective judgement in *Horns and Cows*—which takes the form, “With respect to this [animal], an association with cowhood has been inferred by me; surely, this is just a cow”—is erroneous. Thus, the defender of this view will not only proliferate error amongst inferential judgements, but also amongst introspective ones.

Perhaps, the opponent can bite the bullet here: they could accept the consequence that we introspectively misconstrue the content of our inferential judgements in cases like *Horns and Cows*. Śrīharṣa explores one last way of convincing such an opponent (KKh 395.8-13). Consider the following variant of *Horns and Cows*:

*Extended Horns and Cows.* From a distance, I see an animal with horns. Earlier, I had observed many cows with horns. On the basis of these observations, I falsely judged that all animals with horns are cows. Now, I recall that generalisation. So, I conclude that the animal is a cow. Then, I come closer to the animal, and notice that it has dewlaps. Earlier, I had observed many cows with dewlaps. On the basis of these observations, I correctly judged that all animals with dewlaps are cows. So, I conclude that the animal is a cow.

In this case, my second inference suffers from the fault of establishing what has already been established (*siddhasādhana*): its conclusion had already been established by the first inference. Yet, Śrīharṣa’s opponent is committed to saying that the first episode of reasoning—insofar as it is based on a deviating reason—cannot yield a true conclusion. By contrast, *ex hypothesi*, the second inference (unlike the first episode of reasoning) is not based on any defective inferential mark (since only cows have dewlaps) and therefore can yield an inferential knowledge-event. But this can only happen if the conclusion of the second inference is true. So, Śrīharṣa’s opponent faces a dilemma in cases like this (KKh §264): either they must implausibly say that the second inference in this case doesn’t suffer from the fault of establishing what has already been established, or they must abandon the view that the first defective inference leads to a false conclusion. The first option is risky: unless properly qualified, it could lead to the result that the fault of establishing what has been established never arises at all. The second option will undermine the *Nyāya Definition of Knowledge-Events*: the opponent will be admitting that even a defective reason can help us to draw a true conclusion in cases like *Horns and Cows*.



The lesson: the Naiyāyika cannot easily explain away the accuracy of epistemically lucky inferential judgements in cases like *Mist and Fire* and *Horns and Cows*.

## 5. Knowledgehood and the Epistemic Priority Thesis

So far, we have seen how Śrīharsa exploits a tension within the Nyāya conception of knowledge-events. On the one hand, Naiyāyikas are committed to the simple view that a firsthand awareness is a knowledge-event just in case it is accurate. On the other hand, they embrace a kind of infallibilism: they think that a knowledge-event is produced by a totality of causal conditions that couldn't have led to any error. Cases of epistemic luck like *Guesswork*, *Mist and Fire* and *Horns and Cows* create trouble for this combination of this views. In each of those cases, the final awareness seems accurate, but not a knowledge-event because its causal conditions could easily have led the agent astray. As Śrīharsa has argued, the Naiyāyika cannot exclude such cases from the class of knowledge-events either by adding a blanket infallibility condition on knowledge-events, or by explaining away the accuracy of such awareness-events.

What moral should we draw from this discussion? For a contemporary defender of a “knowledge first” approach, this might just reveal a flaw in the Nyāya approach to knowledge-events. The Naiyāyikas presuppose that knowledge-events can be defined in terms of a *mental condition* like a firsthand awareness plus some *non-mental conditions* such as accuracy, reliability, and so on. However, a defender of a “knowledge first” approach might argue that this assumption is wrong: if knowledge-events are a *sui generis* kind of mental events, they may not after all be decomposable (either metaphysically or conceptually) into mental and non-mental conditions.

Śrīharsa's Nyāya interlocutor Udayana himself seems to anticipate a proposal of this kind (NVTP 51.16-52.4 ad *Nyāyasūtra* 1.1.1).

*Knowledge-Events as a Natural Kind.* An awareness-event is a knowledge-event if and only if it is an instance of the natural kind property (*jāti*) of *being a knowledge-event* (*pramātvā*), what we shall call *knowledgehood*.

Natural kind properties, on the Nyāya-Vaiśeṣika view, cannot be decomposed—either conceptually or metaphysically—into other properties. So, on this view, even though knowledge-events are a sub-species of awareness-events, they are unified by a property—knowledgehood—that cannot be decomposed (conceptually or metaphysically) into (i) a mental property like the property of being an awareness (*jñānatva*) or being a firsthand awareness (*anubhūtitva*) and (ii) a non-mental property like accuracy. In that sense, they are a *sui generis* kind of awareness. It is this latter feature of this definition that makes it analogous to a contemporary “knowledge-first” approach to epistemology. Even though defenders of this latter approach needn't treat states of knowing as a *natural* kind, they nevertheless hold

the view that those states form a *sui generis* mental kind.<sup>12</sup> *Knowledge-Events as a Natural Kind* makes a similar claim.

## 6.1 Not a Natural Kind

Like Udayana, Śrīharṣa too is sceptical of the idea that knowledgehood is a natural kind property. For Udayana, two natural kind properties cannot *cross-cut*: if  $K_1$  and  $K_2$  are natural kinds, it cannot be the case that (i) there exists an entity  $x$  that belongs to both kinds  $K_1$  and  $K_2$ , (ii) there exists an entity  $y$  that is of kind  $K_1$  but not of kind  $K_2$ , and (iii) there exists an entity  $z$  that is of kind  $K_2$  but not of kind  $K_1$ .<sup>13</sup> This constraint on natural kinds—sometimes called the *cross-cutting constraint* (*sāṅkaryā*)—amounts to the constraint that natural kinds should be *categorically distinct*, i.e., that there shouldn't be any smooth transition from one to another.<sup>14</sup> But if two kinds  $K_1$  and  $K_2$  were to cross-cut each other, there would be plenty of things that are of both kinds, but there would also be things that are of kind  $K_1$  but not of kind  $K_2$ , and also things that are of kind  $K_2$  but not of kind  $K_1$ . In that case, it would seem that the distinction between the two kinds is merely conventional; for nature itself doesn't really draw a clear boundary between them. In that case, they shouldn't be treated as natural kinds. Following Udayana, the problem that Śrīharṣa points out for knowledgehood is this (NVTP 51.13-52.11; KKh 444.2-3). For Nyāya-Vaiśeṣika philosophers, perceptual awareness-events form a natural kind in virtue of instantiating the property of epistemic directness (*sākṣāttva*). But note that some perceptual awareness-events are knowledge-events, while others, e.g., inaccurate ones, are not. So, if knowledge-events were to form a natural kind, there would be two natural kinds that cross-cut each other.

Śrīharṣa realizes that this argument isn't decisive. Someone who wishes to defend the proposal that knowledge-events are a natural kind might just deny that perceptual awareness-events form a natural kind, or that the cross-cutting constraint is a genuine constraint on natural kinds. And nothing crucial about the relevant conception of knowledge-events really hangs on whether we recognise knowledge-events as a *natural* kind. As long as there is a sufficiently non-disjunctive property of awareness-events—knowledgehood—on the basis of

---

<sup>12</sup> Williamson's (2000, ch. 1) claim that knowledge is a *sui generis* mental state is best understood as the claim that particular states of knowing, e.g., my seeing here right now that there is a coffee-mug on my desk, are instances of a *mental kind* that cannot be reductively analysed in terms of other kinds or types of mental states like belief and non-mental conditions like truth, reliability, etc. Even though knowledge cannot be reductively analysed in terms of other mental states and non-mental conditions on this view, particular states of knowing are nevertheless instances of a unified kind of mental state. This interpretation fits Williamson's claim that knowledge is the *most general factive stative attitude*, in the sense that other such factive attitudes, such as *perceiving* and *remembering*, are sub-species of knowledge. Nagel (forthcoming), more recently, has cast her version of "knowledge-first" epistemology in terms of knowledge being a type of mental state. Other authors—notably Kornblith (2002) and Kumar (2014) who are sympathetic to naturalism in epistemology—argue that knowledge is a natural kind; "knowledge first" epistemologists don't have to accept this stronger view.

<sup>13</sup> In his commentary *Kiraṇāvalī* on Praśastapāda's *Padārthadharmasaṅgraha*, Udayana explains the distinction by appealing to a set of six kind-blockers (*jātibādhaka*) (Kir 23.3-4). Cross-cutting (*sāṅkaryā*) is one of them. For discussion, see Pellegrini (2016).

<sup>14</sup> For a defence of this constraint, see Ellis (2001). For the claim that there are cross-cutting natural kinds, see Dupre (1993), Khalidi (1998), and Hacking (2007).

which our attributions of knowledge-events are made, it would be permissible to define knowledge-events in terms of knowledgehood.

Therefore, instead of focusing on the question of whether knowledgehood is a natural kind property, Śrīharṣa raises a different question. If knowledgehood is supposed to capture the application-condition (*pravṛttinimitta*) of the expression “knowledge-event,” then it must be the condition that triggers our application of that term. Thus, knowledgehood should (causally) explain our use of the term “knowledge-event” in ordinary discourse. How does it do this? Śrīharṣa explores possible answers to this question while focusing on our *self-attributions* of knowledge-events.

Śrīharṣa begins with the proposal that the presence of knowledgehood in various awareness-events produces our attributions of knowledge-events *by itself*, i.e., in the absence of any awareness on our part that an experience or a thought has the status of being a knowledge-event (KKh 444.3-6). Clearly, this is a non-starter. On this picture, we are just knowledge-ascribing automata, whose ascriptions of knowledge-events are directly manipulated by the presence of knowledgehood outside in the world. No judgements about knowledgehood are required on our part. But, as Śrīharṣa points out, if this were the case, then we couldn’t possibly be in doubt or be mistaken about whether an awareness is a knowledge-event. For, then, whenever an awareness had the status of being a knowledge-event, the knowledgehood present in that awareness-episode would on its own make us correctly ascribe that epistemic status to that awareness. This seems bad.

The only other option is to say that knowledgehood explains our use of the term “knowledge-event” only insofar as our use of this term is triggered by our awareness of various experiences and thoughts as instances of knowledgehood. This, in turn, raises a different question. How do we determine whether an awareness is an instance of knowledgehood? There are two options: either we have *epistemically direct* access to knowledgehood, or we have *epistemically indirect* access to knowledgehood. Śrīharṣa argues that none of these options really work.

## 6.2 No Direct Access

What does it mean to have *epistemically direct* access to knowledgehood? As Śrīharṣa understands it, the position is roughly this. We are equipped with an internal monitoring mechanism, by which we can attend to our own conscious mental occurrences, and thereby can become aware of them and their features. For instance, when I am undergoing a perceptual experience as of there being a wall before me, I may not only attend to the wall that is outside in the external world, but I may also simultaneously attend to the perceptual experience itself and become aware of myself as undergoing that perceptual experience. This internal monitoring mechanism—which enables me in this case to become aware of my perceptual experience—is what the Nyāya-Vaiśeṣika philosopher calls the *inner sense* or the *manas*.

The inner sense's access to various conscious mental occurrences is in most cases *epistemically direct* in the sense that, in order to become aware of their own conscious mental occurrences, an agent doesn't have to base their awareness on something else. This is supposed to be analogous to the case of sensory perception, where in order to become perceptually aware of her external environment, the agent doesn't have to first become aware of some piece of evidence and then infer various claims about the external world on that basis. This kind of epistemically direct access to one's own awareness-events and other conscious mental occurrences is therefore construed as a form of "inner perception" or "introspection", and the *manas*, accordingly, is thought of as an "inner sense." The Nyāya-adjacent position that Śrīharṣa considers is this. Just as we become aware of our own conscious experiences and thoughts by means of inner perception, so also we can detect the property of knowledgehood that belongs to those awareness-events by means of inner perception.

Here is Śrīharṣa's response (KKh 445.12-15). We cannot ascertain whether an awareness is a knowledge-event by inner perception alone. For, even when we are aware of all the introspectable features of an awareness by means of inner perception, we could still be in doubt or be mistaken about whether it is a knowledge-event. The examples that Śrīharṣa considers are cases where *positive introspection* fails for knowledge-events: cases where an agent comes to know something, but doesn't (or can't) come to know that they have come to know it, because they either doubt whether they know it or falsely think that they don't know it. A commentator on *the Refutation*, Ānandapūrṇa, describes an example similar to Radford's (1966) case of the unconfident examinee (KKhPV 446.6-8).

*The Circumspect Jeweller.* An apprentice to a jeweller, unbeknownst to herself, is an expert at discerning precious metals like silver. So, when they see a piece of jewellery made of silver, they can immediately tell that it is made of silver. But, since they are not sure of their own ability to tell precious metals apart from ordinary ones, they doubt whether they have learnt that the jewellery is made of silver. So, they know that the piece of jewellery before them is made of silver, but still doubt whether they know this.

In this case, the agent might indeed know everything about their judgement that there is to know by introspection. But that still wouldn't dispel their doubts about its epistemic status. But if knowledgehood were indeed an introspectable feature of awareness-events, this wouldn't be the case. If positive introspection for knowledge-events can fail in this way, knowledgehood cannot be a feature of awareness-events which we can detect by inner perception. Since inner perception is the only way we can have direct access to properties of awareness-events, it follows that we don't have direct access to knowledgehood.

### 5.3 The Problems of Indirect Access

This leaves us with the possibility that we may have epistemically indirect access to knowledgehood. Śrīharṣa entertains two different ways of fleshing out this possibility (KKh

445.15-17). The first option is to say that an agent gains access to the property of knowledgehood inherent in her own awareness-events by means of inner perception, but this kind of inner perception is *indirect*, i.e., dependent on the agent's antecedent awareness of some *symptom* (*cihna*) that the relevant awareness-events possess. In other words, when an agent becomes aware of their own experience or thought, they first pick up on this symptom, and that triggers an introspective recognition of that awareness as a knowledge-event. The second way of fleshing out the relevant proposal doesn't appeal to inner perception. On this approach, we become aware of knowledgehood by *inference* on the basis of some symptom possessed by the relevant awareness-events. In either case, crucially, the symptom that forms the basis of our awareness of knowledgehood is something that can be introspectively grasped independently of knowledgehood itself. It is this feature of the two proposals that Śrīharṣa will exploit in his arguments against them.

Śrīharṣa's opening move is to say: "Let's grant that we become aware of knowledgehood on the basis of certain symptoms. But are there many such symptoms or just one?" Suppose Śrīharṣa's interlocutor says that there is just one symptom which underlies our self-attributions of knowledge-events (KKh 445.17-18). But, if that is true, there is no reason to posit an unanalysable property such as knowledgehood. The thought is that, when we defined knowledge-events in terms of knowledgehood, our definition was supposed to capture the extension (or intension) of the term 'knowledge-event' as it is commonly used by laying down its application-conditions. But now it turns out that there is a single symptom which constitutes the reason for which we apply the term 'knowledge-events' to our experiences and thoughts. Then, there is no reason to treat knowledgehood as the application-condition of the term; the symptom itself should do the job.

What if there are many such symptoms? Śrīharṣa presents two distinct challenges here. First, Śrīharṣa challenges his interlocutor to say what these many symptoms are (KKh 445.18-19). Suppose there is a series of different symptoms  $S_1, S_2, S_3, \dots$ , such that each  $S_i$  *decisively indicates* that the relevant awareness is a knowledge-event. Śrīharṣa's argument is that his opponent won't be able to list any such symptom: for any symptom that the opponent may mention, it will suffer from one of the problems that Śrīharṣa has already described for the proposed definitions of knowledge-events.

However, one needn't be convinced by this argument. One might argue that in order to infer the presence of knowledgehood on the basis of some symptom, the relevant symptom doesn't have to be a *decisive* indicator of knowledgehood and therefore needn't rule out all awareness-events that are not knowledge-events. For example, red spots on one's skin aren't a decisive indicator of measles, but they surely are a *reliable* indicator of measles. And on that basis, one can know that a person is suffering from measles. Similarly, certain features of an awareness—its accuracy, its possession of some kind of epistemic pedigree, and so on—may serve as a *reliable* indicator of knowledgehood under a range of circumstances, even though none of them may decisively indicate the presence of a knowledge-event by ruling out all instances of epistemic luck.

Even though Śrīharṣa doesn't explicitly address this objection, the second challenge that he raises seems to speak to it. This challenge arises from a positive proposal about attributions of knowledge-events: the *Epistemic Priority Thesis*. According to this proposal, knowledge-events are *epistemically prior* to other non-factive mental states and events. When we are trying to determine whether an agent has undergone a knowledge-event, we don't initially ascribe to them some other non-factive mental event, and then check if that event meets some further conditions (like truth or reliability) necessary for it to count as a knowledge-event. Rather, we (if we're rational) treat certain mental events by default as knowledge-events until a defeater comes along.

In Śrīharṣa's own work, the *Epistemic Priority Thesis* is articulated in two distinct ways. In some contexts, it is expressed as a thesis about *self*-attributions of knowledge-events (KKh 445.19-20). On this version of the thesis, in order to judge that their own experience or thought constitutes a knowledge-event, the agent doesn't need to check whether the epistemic credentials of that experience or thought are any good, e.g., whether it is produced by certain epistemic virtues, or the absence of epistemic defects.<sup>15</sup> As long as there is no positive reason to suspect that the awareness is inaccurate or is produced by a defective set of causal conditions, the agent is required not to be uncertain about the status of the awareness as a knowledge-event. Call this:

*The Default Knowledgehood Thesis.* If an agent undergoes an experience or a thought with the content that *p*, then, in the absence of any rebutting or undercutting evidence against their awareness, they are required by rationality not to doubt that they have come to know (or learnt) that *p*.

This is sometimes called the theory of *intrinsic knowledgehood* (*svataḥprāmāṇyavāda*). On this story, if a rational agent doesn't have any doubt about the epistemic credentials of an awareness, then—provided that they are aware of that awareness—they will typically never hesitate to ascribe the status of being knowledge-event to that awareness. The agent's judgement that the awareness is a knowledge-event won't be dependent on their antecedent grasp of any symptom—like accuracy or reliability—that the relevant awareness possesses over and above the mere fact that it is an awareness that portrays the world in a certain way.

Why should we think that this story about self-ascriptions of knowledge-events is true? The best *a priori* argument for this theory is a regress argument offered by Bhāṭṭa Mīmāṃsakas like Kumāriḷa Bhatta and his commentators. Suppose we reject this theory and say that, to judge that they have learnt something, an agent must always antecedently judge that their awareness has the epistemic credentials required for being a knowledge-event, e.g., that it is accurate or that it is reliably formed. Now, take an agent who has formed such judgements. If the agent is rational, why should they rely on such judgements unless they also think that

---

<sup>15</sup> This theory was defended by Kumāriḷa Bhatta in vv. 52-53 in his *Detailed Commentary in Verse* (*Ślokavārttika*) on *Mīmāṃsāsūtra* 1.1.2. In the fourth chapter of the *Refutation*, Śrīharṣa tells us that he has defended this view in another text called *The Intent of the Lord* (*Īśvarābhisandhi*), a work that is lost to us.

these judgements are knowledge-events? To assure themselves that these judgements are knowledge-events, they might appeal to another set of judgements about the epistemic credentials of these judgements. This will lead to a regress. To avoid the regress, we must grant that there is at least one set of judgments whose status as knowledge-events the agent rationally takes for granted without relying on a further set of judgements.<sup>16</sup> But, then, why shouldn't we say that about the very first awareness?

In the fourth chapter of the *Refutation*, Śrīharṣa seems to generalise this story from the case of self-attributions of knowledge-events to the case of attributions of knowledge-events to others (KKh 1322.14-15). He argues that we can criticise a sceptic, who refuses to recognise our ordinary waking experiences and thoughts as knowledge-events, by appealing to a kind of suppositional reasoning (*tarka*) which is driven by a default assumption (*utsarga*). Unlike other standard instances of suppositional reasoning accepted by Śrīharṣa's Nyāya-Vaiśeṣika interlocutors, the point of such reasoning is not to show that some undesirable consequence follows from the sceptic's view, but rather to show that the sceptic's position doesn't fit our default picture of how things normally are in the world. The underlying principle is this: for any two mutually incompatible and exhaustive features  $F_1$  and  $F_2$ , if  $F_1$  is *predominantly* or *normally* present amongst a class  $C$  of objects rather than  $F_2$ , then (absent defeaters) we should accept that any arbitrary object chosen from  $C$  has  $F_1$  rather than  $F_2$ . As Kumārila and his commentators claim, (given the threat of regress) we are entitled—if not required—to assume without further evidence that our awareness-events are knowledge-events under normal conditions, e.g., when we are awake or when our awareness-generating mechanisms aren't malfunctioning. So, in the absence of any defeater that suggests that the conditions are abnormal, we should treat our awareness-events (irrespective of whether they belong to ourselves or others) as knowledge-events. Thus, Śrīharṣa writes by appealing to Kumārila:

For example, this kind of suppositional reasoning would apply against someone who, in the absence of defeat, ascribes the absence of knowledgehood to an awareness that arises for a well-functioning and wakeful agent, even though there is no difference with respect to the presence of any source of knowledge (*pramāṇa*) that could determine [the awareness'] status of being a knowledge-event or of not being a knowledge-event. However, it would not apply against someone ascribes knowledgehood [to such an awareness]. It is this suppositional reasoning by adopting which [Kumārila] has said:

Therefore, the status of an awareness as a knowledge-event is obtained in virtue of its having the nature of awareness. That is cancelled due to the awareness of the object's being otherwise, or a defect that arises from the causes [of the awareness]. (KKh 1322.15-20)

---

<sup>16</sup> This argument occurs in v. 56 in Kumārila Bhaṭṭa's *Detailed Commentary in Verse (Śloka-vārttika)* on *Mīmāṃsāsūtra* 1.1.2.

As Śrīharṣa is careful to emphasise, a default assumption—like the assumption that our experiences and thoughts are knowledge-events under normal conditions—is not a hypothesis that we are uncertain about or even treat as highly likely. A default assumption is not fundamentally different from a state of certainty, i.e., a judgement or a determination (*nirṇaya*) (KKh 1322.22-1323.11). When we conclude that an awareness is a knowledge-event on the basis of the default assumption that our experiences and thoughts under normal conditions are knowledge-events, there remains no residual uncertainty about the epistemic status of that awareness. We simply judge that the relevant awareness is a knowledge-event. Thus, in this way, a background default assumption not only leads to the destruction of uncertainty, but also functions—at least from our perspective—as a source of knowledge with respect to facts about whether we know. Once again, on this view, knowledge-events are epistemically prior: to rationally ascribe knowledge-events to ourselves and others, we don't need to determine whether the relevant awareness-events satisfy any further non-mental conditions such as accuracy or reliability; given our background default assumption, in the absence of defeaters, we can rationally ascribe such knowledge-events to ourselves.

What is important in this picture—but is left implicit by Śrīharṣa—is that we cannot rationally accept the default assumption that our experiences and thoughts are normally knowledge-events on the basis of any judgements that—independently of that assumption—confirm or provide evidence that those experiences and thoughts are accurate or reliable. If such judgements were required to justify this default assumption, a Kumāri-la-style regress argument would be difficult to avoid.

This brings out a key feature of the position that Śrīharṣa is driving us towards. On this view, we simply don't have any independent evidence that our ordinary attributions of knowledge-events are reliably tracking instances of a unified kind property—knowledgehood—in ourselves and others. Rather, if Śrīharṣa is to be believed, such attributions are typically based on two factors: one is independent of the context of attribution, and the other is not. The context-invariant factor is our default assumption that our experiences and thoughts are knowledge-events under normal conditions. The context-dependent factor is the lack of any available evidence that suggests that the conditions are abnormal. When both these factors are present, we can—and often do—rationally ascribe knowledge-events to ourselves and others. But neither of these two factors guarantees that our attributions of knowledge-events are in fact tracking a unified or *sui generis* kind of mental events. First, the assumption that our awareness-events are knowledge-events under normal circumstances isn't justified by any independently acquired evidence. So, it may indeed be false. Second, even if that assumption is true, the context-dependent factor—namely, the lack of evidence that suggests that conditions are abnormal—doesn't guarantee that conditions are in fact abnormal. So, it is compatible with this contextualist story that, in different contexts of ascribing knowledge-events, we are treating very different kinds of awareness-events as knowledge-events. Thus, there is no way of showing that our ordinary attributions of knowledge-events are reliably tracking a unified or *sui generis* mental kind.

## 6. Scepticism



The difference between Śrīharaṣa and contemporary defenders of “knowledge first” epistemology lies in the fact that the latter adopt an anti-sceptical stance about our ordinary attributions of knowledge to ourselves. They assume that there is a unified mental kind—the *most general factive stative mental attitude*—that is reliably tracked by our ordinary ascriptions of knowledge.<sup>17</sup>

By contrast, Śrīharaṣa thinks that his arguments against the definability of knowledge-events pave the way for thoroughgoing scepticism about our ordinary attributions of knowledge-events. He states his conclusion as follows.

Moreover, whatever may be stated as the defining characteristic of knowledge-events, if it gives rise to linguistic usage pertaining to knowledge-events insofar it is not an object of awareness or insofar as it is merely an object of an awareness [which isn’t a knowledge-event], then there will be a problem of overgeneration. If it gives rise such ascriptions in virtue of being known, then it will be difficult to determine if it is not determined what knowledge-events are. (KKh 448.22-449.1)

Suppose our ordinary attributions of knowledge-events reliably track a unified mental kind. Assume that what instances of that kind share is a property X, a defining characteristic of knowledge-events. But, if we reliably track that mental kind, then we must have the capacity to reliably make attributions of knowledge-events only in cases where the property X is instantiated. But, if we don’t know what this property X is—either because we are unaware of this property, or are aware of this property but don’t know that it distinguishes knowledge-events from awareness-events that are not knowledge-events—then it is hard to explain how we might have such a capacity to reliably make attributions of knowledge-events. By contrast, if this property X generates our attributions in virtue of being known to us but we cannot determine this property, even then it will remain hard to show that we know what the defining property of knowledge-events is. In either case, we end up with a sceptical conclusion: namely, that we cannot show that our ordinary attributions of knowledge-events are reliably tracking a unified mental kind.

A tempting response to this problem is to say that, even if we cannot determine any extensionally (or intensionally) adequate defining characteristic of knowledge-events, we nevertheless can reliably track instances of that characteristic. Śrīharaṣa is unmoved by this response.

[The opponent:] Let the defining characteristic be undetermined, but it is in fact such [i.e., does give rise to linguistic usage pertaining to knowledge-events].

[Reply:] No. This is because, then, there will be the undesirable consequence that there will be no response to someone who says that [the defining characteristic] is in

---

<sup>17</sup> See footnote 13. Nagel (2013, forthcoming) discusses this anti-sceptical assumption explicitly.

fact not such, and because it will follow that it is futile to determine what knowledgehood is. And let linguistic usage pertaining the reality of objects such as a pot also proceed due to a knowledge-event itself [without any intervening ascriptions of knowledge-events]. No further elaboration is needed. (KKh 449.1-5)

The position, according to Śrīharṣa, faces three objections. The first two objections revolve around two flaws of the Nyāya position. First, if the Naiyāyika has no way of articulating what the defining characteristic of knowledge-events is, then they have—by their own admission—no way of showing that our attributions of knowledge events are made in response to the presence of a defining property of knowledge-events. But, then, they have no way of refuting a sceptic who makes the opposite claim: namely, that our attributions do not reliably track the presence of such a property. The second objection is *ad hominem*: if the Naiyāyika says that we reliably track the presence of the defining characteristic of knowledge-events without in fact being able to determine what that defining characteristic is, then the whole Nyāya project of discovering such a defining characteristic will be pointless.

The third objection is better. The Naiyāyika claims that there is a defining characteristic of knowledge-events, which is reliably tracked by our ordinary attributions of knowledge-events. Yet, they cannot tell us any believable story about how we know that there is such a characteristic. If this were permitted, then we should also be allowed to claim that an ordinary object such as a pot exists, without putting forward any story about how it is that we know such an object to be real. But, surely, if an agent asserts that a pot exists, they can reasonably be asked, “How do you know that the pot exists?”<sup>18</sup> But, if the agent then fails to come up with an answer to that question, then it seems unreasonable for them to persist in asserting that the pot exists. But the Naiyāyika’s stance seems to suggest that this charge of unreasonableness doesn’t apply when it comes to the defining characteristic of knowledge-events. This, then, raises a challenge for the Naiyāyikas: they must explain why the defining characteristic of knowledge-events should be treated differently from ordinary objects such as pots.

A defender of “knowledge first” epistemology will resist this sceptical argument. Śrīharṣa claims that it is unreasonable for his opponent to assert that there is a defining characteristic *X* of knowledge-events, on the basis of which we make our attributions of knowledge-events, without being able to say what this characteristic *X* is and how we recognise it. There are two—mutually compatible—ways of resisting this claim. The first is to say that our recognition of this characteristic is *tacit* or *implicit*, and therefore not accessible to conscious reflection or verbal expression. So, even though we are able to know that this characteristic is instantiated by our awareness-events, we cannot consciously determine or articulate what it is. Alternatively, one could appeal to failures of the KK principle, the principle that whenever one knows that *p*, one knows (or is in a position to know) that *p*. Thus, even though we may

---

<sup>18</sup> This is one of the observations on the basis of which Williamson (2000, ch. 11) motivates the view that knowledge is the norm of assertion. Śrīharṣa himself states the principle in a stronger form (KKh 90.4-5): “A person who is constrained by the sources of knowledge and speaks of the existence of a pot should state a source of knowledge with respect to that.”

know that the defining characteristic  $X$  of knowledge-events is instantiated by our awareness-events, we don't know that we know this.<sup>19</sup> Due to the lack of such higher-order knowledge, we fail to articulate how we recognise it.

These responses miss the point of the argument that Śrīharaṣa is offering here. Śrīharaṣa's argument is *dialectical*, indexed to the context of a debate about the thesis that there is a defining characteristic of knowledge-events in virtue of recognising which we ordinarily make attributions of knowledge-events. The point, more abstractly, is this. Imagine a context of a conversation where the content that  $p$  isn't part of the stock of information shared by the participants of the conversation. Suppose that, in that conversation, an agent  $S$  asserts that  $p$ , but then cannot answer the follow-up question from their interlocutor, "How do you know that  $p$ ?" Since it cannot be settled in that context whether  $S$  in fact knows  $p$ , the content of  $S$ 's assertion—even though the assertion may originally have been appropriate because it was based on knowledge—won't be accepted by the other participants. At that point,  $S$  cannot just persist in asserting that  $p$  without offering any further explanation of how it is that they know that  $p$ . Similarly, for the defender of the view that there is a defining characteristic of knowledge-events that guides our ordinary attributions of knowledge-events, it is inappropriate to *persist in asserting their view* when they cannot explain how they know this.

If Śrīharaṣa's positive account of how we ascribe knowledge-events is correct, this failure to answer the "How do you know?" challenge isn't induced by a rectifiable epistemic defect, e.g., the tacitness of some piece of knowledge that could in principle be made explicit, or the lack of some piece of higher-order knowledge that in principle could be acquired. It is in fact *impossible* to show that our ascriptions of knowledge are made in response to *any* defining characteristic of knowledge-events. For Śrīharaṣa, we typically ascribe knowledge-events to ourselves and others, *not* because we recognise some unifying feature of knowledge-events, but rather because, first, we assume that our awareness-events are normally knowledge-events and, second, we have no reason to suspect that conditions are abnormal. But the problem is that the assumption in question is not justified on the basis of any independently acquired evidence, and the lack of evidence that conditions are abnormal doesn't guarantee that conditions are normal. This leaves open the possibility that, in different contexts of ascription, our ascriptions of knowledge-events pick out heterogenous mental events, which have nothing in common with each other except perhaps that they called "knowledge-event" in those contexts.

This throws the sceptical problem that Śrīharaṣa develops here into sharper focus. Since the defender of the claim that knowledge-events—or states of knowing more generally—form a

---

<sup>19</sup> See Williamson (2000, chs. 4 and 5). As I have noted in §6.2, Śrīharaṣa isn't committed to any strong form of introspection principle about knowledge-events: he thinks that an agent can be subject to error and uncertainty about the presence and the absence of their own knowledge-events. However, he would reject Williamson's strong anti-luminosity thesis that there is no non-trivial condition  $C$  such that, whenever  $C$  obtains, one is in a position to know that it obtains. While defending the idea that consciousness is self-manifesting (*svaprakāśa*), he defends the claim that, whenever we are conscious of something, we are in a position to know that we are conscious of something (KKh 81.1-8).

unified mental kind cannot show that our ascriptions of such events or states reliably track any such kind, they cannot continue to assert that there is such a unified mental kind.<sup>20</sup>

## References

### Primary Texts and Abbreviations

- KKh Śrīharṣa's *Khaṇḍanakhaṇḍakhādya* in *Śrīharṣhapraṇītam Khaṇḍanakhaṇḍakhādya* *Ānandapūrṇaviracitayā Khaṇḍanaphakkikāvibhajanākhyayā vyākhyayā 'vidyāsāgarī' ti prasiddhayā sametam Citsukha-Śaṅkaramiśra-Raghunāthakṛtaṭīkāvalambinyā tippaṇyā sanātham*. Edited by Gaṅganātha Jhā and Lakṣmaṇa Drāviḍa Śāstrī. Benares: Chowkhamba Sanskrit Book Depot, 1904-1914.
- KKhPV Ānandapūrṇa Vidyāsāgara's *Khaṇḍanakhaṇḍakhādya* *phakkikāvibhujana* in *Śrīharṣhapraṇītam Khaṇḍanakhaṇḍakhādya* *Ānandapūrṇaviracitayā Khaṇḍanaphakkikāvibhajanākhyayā vyākhyayā 'vidyāsāgarī' ti prasiddhayā sametam Citsukha-Śaṅkaramiśra-Raghunāthakṛtaṭīkāvalambinyā tippaṇyā sanātham*. Edited by Gaṅganātha Jhā and Lakṣmaṇa Drāviḍa Śāstrī. Benares: Chowkhamba Sanskrit Book Depot, 1904-1914.
- Kir Udayana's *Kiraṇāvalī* in *Praśastapādabhāṣyam: With the Commentary Kiraṇāvalī of Udayanācārya*. Edited by Jitendra S. Jetly. Baroda: Oriental Institute, 1971.
- LM Udayana's *Lakṣaṇamālā* in *Lakṣaṇamāla by Udayanācārya*. Edited by Śaśinātha Jhā. Darbhanga: Mithila Institute, 1963
- NK Śrīdhara's *Nyāyakandalī* in *Nyāyakandalī Being a Commentary On Praśastapādabhāṣya, With Three Sub-commentaries*. Edited by Jitendra S. Jetly and Vasant G. Parikh. Baroda: Oriental Institute, 1991.
- NKA Udayana's *Nyāyakusumāñjali* in *The Nyāyakusumāñjali of Śrī Udayanācārya with Four Commentaries: The Bodhinī, Prakāśa, Prakāśikā (Jalada) and Makaranda by Varadarāja, Vardhamānopādhyāya, Mecha Thakkura and Rucidattopādhyāya and with Notes by Śrī Dharmadatta (Bachchā Jhā)*. Edited by Padmaprasāda Updhyāya and Dhunḍirāja Śāstrī. Varanasi: Chowkhamba Sanskrit Series, 1957.
- NM I Jayanta Bhaṭṭa's *Nyāyamañjarī* in *Nyāyamañjarī: Sampādakagrathitanyāyasaurabhākhyatippaṇīsamanvitā*. Vol. I. Edited by K. S. Varadacharya. Mysore: Oriental Research Institute, 1969.
- NRK Pārthasārathimiśra's *Nyāyaratnākara* in *Ślokavārttikam Kumārilabhaṭṭapādaviracitaṃ Pārthasārathimiśraviracitayā Nyāyaratnākaravyākhyayā Sanātham*. Edited by Dwarikadas Sastri. Varanasi: Tara Publications, 1978.
- NVTP Udayana's *Nyāyavārttikatātparyapariśuddhi* in *Nyāyavārttikatātparyapariśuddhi of Udayanācārya*. Edited by Anantalal Thakur. New Delhi: Indian Council of Philosophical Research, 1996.

---

<sup>20</sup> Thanks to Arturs Logins, Rosanna Piscasia, Mark Siderits, Davey Tomlinson, an anonymous referee for this volume, and audiences at University of Oxford, University of Manchester, and Columbia University for their helpful comments.

- ŚV Kumāriḷa Bhaṭṭa's *Ślokavārttika* in *Ślokavārttikam Kumāriḷabhaṭṭapādaviracitaṃ Pārthasārathimiśraviracitayā Nyāyaratnākaravyākhyayā Sanātham*. Edited by Dwarikadas Sastri. Varanasi: Tara Publications, 1978.
- ŚVK *The Mīmāṃsāślokavārttika of Kumāriḷabhaṭṭa: With the Commentary Kāśika of Sucaritaṃśra*. Parts I and II. Edited by K. Sāmbaśivaśāstri. Trivandrum: CBH Publications, 1990.
- ŚVTT *Ślokavārtikavyākhyā (Tātparyaṭikā) of Bhaṭṭombeka*. Edited by S. K. Ramanatha Sastri. Madras: University of Madras, 1940.

### Other Sources

- Conee, Earl and Richard Feldman, 1998, "The Generality Problem for Reliabilism", *Philosophical Studies*, 89(1): 1–29.
- Das, Nilanjan (2018). Śrīharṣa. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2021 Edition). URL = <https://plato.stanford.edu/archives/win2021/entries/sriharsa/>.
- Das, Nilanjan (2021). Gaṅgeśa on Epistemic Luck. *Journal of Indian Philosophy* 49 (2):153-202.
- Dasti, M. R. 2012. Parasitism and disjunctivism in Nyāya epistemology. *Philosophy East and West* 62 (1):1-15.
- Dupré, J. (1993). *The Disorder of Things : Metaphysical Foundations of the Disunity of Science*. Cambridge MA: Harvard University Press.
- Ganeri, Jonardon, (2007). *The Concealed Art of the Soul: Theories of Self and Practices of Truth in Indian Ethics and Epistemology*. Oxford: Oxford University Press.
- Ganeri, Jonardon (2017). Śrīharṣa's Dissident Epistemology. In Ganeri, Jonardon (ed.), *The Oxford Handbook of Indian Philosophy*. Oxford: Oxford University Press.
- Gettier, Edmund L. (1963). Is Justified True Belief Knowledge? *Analysis* 23 (6):121-123.
- Granoff, Phyllis E. (1978). *Philosophy and Argument in Late Vedānta: Śrī Harṣa's Khaṇḍanakhaṇḍakhāḍya*, Dodrecht: D. Reidel.
- Hacking, Ian (2007). Natural Kinds: Rosy Dawn, Scholastic Twilight. *Royal Institute of Philosophy Supplement*, 61: 203–239.
- Ichikawa, Jonathan & Jenkins, C. S. I. (2017). On Putting Knowledge 'First'. In Joseph Adam Carter, Emma C. Gordon & Benjamin Jarvis (eds.), *Knowledge First: Approaches in Epistemology and Mind*. Oxford: Oxford University Press.
- Jha, Ganganatha (1913/1986). *The Sweets of Refutation. An English Translation of the Khaṇḍanakhaṇḍakhāḍya of Shri-Harṣa*, Allahabad: "Indian Thought Series". [1986 Reprint Edition: Delhi: Satguru Publications].
- Kornblith, Hilary (2002). *Knowledge and Its Place in Nature*. Oxford: Oxford University Press.
- Kumar, Victor (2014). "Knowledge" as a natural kind term. *Synthese*, 191(3), 439–457.

Matilal, Bimal Krishna (1986). *Perception: An Essay on Classical Indian Theories of Knowledge*, Oxford: Oxford University Press.

Matilal, Bimal Krishna (1990). *Logic, language, and reality: Indian philosophy and contemporary issues*. Delhi: Motilal Banarsidass.

McGlynn, Aidan (2014). *Knowledge First?*. New York: Palgrave Macmillian.

Nagel, Jennifer (forthcoming). Seeking Safety in Knowledge. In *Proceedings and Addresses of the American Philosophical Association*.

Pellegrini, G. 2016. Differentiating *jāti* and *upādhi*: Towards a further exegesis of the six *jātibādhakas* according to Navya Nyāya. *Proceedings of the Meeting of the Italian Association of Sanskrit Studies. Rivista Degli Studi Orientali* 28: 73-91.

Phillips, Stephen H. (1999) *Classical Indian Metaphysics: Refutations of Realism and the Emergence of "New Logic."* Chicago and La Salle, Illinois: Open Court.

Ram-Prasad, Chakravarthi (2002): *Advaita Epistemology and Metaphysics*. London: Routledge.

Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, Timothy (2023). Vaidya on Nyāya and Knowledge-First Epistemology. *Oxford Studies in Epistemology* 7: 365-375.

Vaidya, A. J. 2013. Nyāya perceptual theory: disjunctivism or anti-individualism? *Philosophy East and West* 63 (4):562-585.

Vaidya, A. (2023). Elements of Knowledge-first Epistemology in Gaṅgeśa. *Oxford Studies in Epistemology* 7: 336-364.