Synthese (2021) 198:3339–3368
https://doi.org/10.1007/s11229-019-02283-w

KNOWING THE UNKNOWN

Four grades of ignorance-involvement and how they nourish the cognitive economy

John Woods[1]

Received: 28 February 2019 / Accepted: 3 June 2019 / Published online: 17 June 2019
© Springer Nature B.V. 2019

Abstract

In the human cognitive economy there are four grades of epistemic involvement. Knowledge partitions into distinct sorts, each in turn subject to gradations. This gives a fourwise partition on *ignorance*, which exhibits somewhat different coinstantiation possibilities. The elements of these partitions interact with one another in complex and sometimes cognitively fruitful ways. The first grade of knowledge I call "anselmian" to echo the famous declaration *credo ut intelligam,* that is, "I believe in order that I may come to know". As construed here, one knows in this anselmian way that $E=mc^2$ just in case one knows that sentence expresses a true statement, but without having to understand the proposition it expresses. Most epistemologists ignore the significance of this grade of epistemic involvement. In a second grade of epistemic involvement, knowing that $E=mc^2$ is knowing what that sentence means and understanding the proposition it express. This is knowledge in the propositional or semantic sense, and is the dominant target of epistemological investigation. Tacit and implicit (TI) knowledge occupies another tier. A typical example would be something that someone has "known all along" but, until now, hasn't had occasion to put her mind to it or formulate in words. TI-knowledge remains a minority interest in today's epistemology. Operating at a fourth grade of epistemic involvement is what I call "impact"-knowledge, which is the knowledge of a matter at its deepest and most widespread. An example, to be discussed below, is the knowledge that was generated by the Wiles proof of Fermat's last theorem. Its true importance lies not only, or even mainly, in its verification of a commonly accepted fact about numbers, but rather in its enrichment of the mathematics of elliptical curves and the promise it holds for greater advancement into the mathematical unknown. Knowledge of this fourth grade has yet to find a seat in the parliaments of epistemology. Knowledge of the anselmian sort is independent of the other three. Tacit and implicit knowledge is incompatible with anselmian and semantic knowledge but coinstantiable

✉ John Woods
john.woods@ubc.ca
https://www.johnwoods.ca/

[1] Department of Philosophy, University of British Columbia, 1866 Main Mall BUCH E370, Vancouver, BC V6T 1Z1, Canada

with impact-knowledge. Semantic knowledge is incompatible with tacit and implicit knowledge but coinstantiable with the others. Impact-knowledge is pairwise coinstantiable with the others. Below I will bring the ignorance partitions into such alignment as they have with these ones. In doing so, I'll propose a naturalized causal response epistemology designed to give these interactive distinctions the theoretical air they need to breathe.

Keywords Abduction · Anselmian knowledge · Big-box scepticism · Causal-response epistemology · Cognitive economics · Command-and-control epistemology · Consciousness · Deception · Error · Explicit knowledge · Ignorance · Ignorance-preservation · Impact-knowledge · Information · Information filters · JTB model · Logic naturalized · Misinformation · Natural knowledge · Propositional (semantic) knowledge · Tacit and implicit knowledge · Virtual truth-values

1 Natural knowledge

Ignorance is a state that knowledge-seeking beings are in, not rocks, or peanuts or prime numbers. Because ignorance is the absence of knowledge, it only stands to reason that a good theory of ignorance might be got from a good theory of what it's the absence of. Scarcely a trace of the transfinite abundance of what's true will ever be known. In *ordo essendi,* ignorance is prior. In *ordo cognoscendi,* knowledge is prior. In its application here, the distinction should be taken as saying that to achieve a knowledge of ignorance, we will first have to attain a knowledge of knowledge.

The objects of knowledge come in various stripes—that *p,* whether *p,* why *p,* what to do to bring it about that *p*, when to do it, how to do it, and so on. Ignorance inherits these variations. At times, ignorance is privation in the sense noted by Aristotle at *Physics* 192[a] 3 (Aristotle 1984). It is the absence of the knowledge that a knowledge-seeking being should have had as a matter of course. In other cases, ignorance is simply the absence of what someone might have wished to know but lacks the occasion or wherewithal for; in still others it is absence of knowledge no one would be interested in having.

It is interesting how little notice of ignorance is taken by epistemologists of the present day.[1] We might think this readily explainable. If ignorance is the absence of knowledge, would it not suffice for its specification to lay out the conditions on knowledge that not-knowing falls short of? This, certainly, is a part of the story. But

[1] A welcome exception is Arfini (2019). For the unexceptional rest, see for example, the following: In a volume as authoritative as Goldman's and Whitcomb's *Social Epistemology: Essential Readings,* ignorance makes no appearance in the index. The same is true of Bernecker's and Pritchard's *The Routledge Companion to Epistemology,* save for one citation of Descartes' argument from ignorance in the First Meditation, which hardly counts. Check Blackburn's *Oxford Dictionary of Philosophy* and you'll find nothing except "*ignoratio elenchi*". Save for that one reference, there is no mention in these three places of autoepistemic inference or arguments from ignorance. Honderich's *The Oxford Companion to Philosophy* comes to much the same end. It acknowledges the more general form of Descartes' use of it in its scant entry on arguments from ignorance, and a briefer mention of *ignoratio elenchi.* See Goldman and Whitcomb (2001), Bernecker and Pritchard (2013), Blackburn (2005), and Honderich (2005).

when we press to determine the nature, scope and limits of human knowing, we find ourselves in the eye of a squall brought on by a lack of philosophical consensus on those very matters. What I will propose here is that when we get closer to the right answers about knowing, we begin to see that the contrast between knowledge and ignorance is a good deal more complex than we might have supposed at the outset. It is commonly supposed that ignorance is an unqualified net cost, effected merely by the failure to know. This is an over-simple assessment. There are situations galore in which the cost of ignorance is compensated for by the spur it gives to overcoming it, sometimes in ways that lead to a knowledge of new and sometimes surprising things. In some of these cases, ignorance is itself a byproduct of epistemic success. Often enough, the solution to a problem opens up new horizons for the future investigation of things not yet known, disclosing a spread of previously unrecognized ignorance. But if the field prospers, the *originating* ignorance that started it all will have yielded substantial epistemic dividends.[2] What is wanted here is an epistemology with flexibility enough to take these complexities into proper account. For that we'll need an epistemology for natural knowledge.

The expression "natural knowledge", is ambiguous as between what it's knowledge of and whom it is known by. I read it in the second way. Natural knowledge is the knowledge acquired by beings like us, who spend our earthly lives under the chrononic and gravitational tug of the space-time world.[3] I don't accept the view that, for a natural object to have knowledge, the object of the knowledge must itself be an object of nature. I lack the space to litigate the matter here, and will instead introduce it as a working assumption.

Like knowledge, ignorance is subject to briskly debated metrics of scope and depth. If epistemic scepticism were true, the scope of ignorance would vary inversely with the scope of knowledge. In a naturalized epistemology, knowledge enjoys a larger prosperity and ignorance a correspondingly smaller one. Ignorance can also be conditioned by considerations of depth. The Wiles proof of Fermat's last theorem is a good indication of how deeply layered our ignorance can be.[4] Other factors also apply. Sometimes ignorance is triggered by error,[5] sometimes by unobservance,[6] and sometimes by deception or misinformation.[7] Also important are the limits of human knowability and the variabilities of agency-type, as with the lone individual Ruth Barcan, the multiagent NATO and the AI marvel AlphaGo. Theoretical progress

[2] See here Kitcher: "Typically, a flourishing science is incomplete. At any time, it raises more questions than it can currently answer. But incompleteness is no vice. On the contrary, incompleteness is the mother of fecundity …. A good theory should be productive; it should raise new questions and presume that these questions can be answered without giving up its problem-solving strategies." Kitcher (1983; pp. 46–48).

[3] Just to be clear, when I speak of the space-time world, I mean the world we ourselves inhabit as natural beings. For present purposes, considerations from relativity theory don't apply. I am not invoking four-dimensional Minkowki space.

[4] Wiles (1995; pp. 443–551).

[5] Rescher (2007, 2009).

[6] The word "ignore" derives from the fourteenth century French "ignorer" ("to be unaware of"), derived in turn from the 14th c. Latin ("not to know, be unacquainted"), but also ("take no notice of, disregard").

[7] See, for example, Arfini and Magnani (2016; pp. 612–627), and Arfini et al. (2018; pp. 37–50).

with ignorance requires a judicious way with epistemological assumptions. In coming sections, I will explain why I see a causal response theory of knowledge (CR) as the best fit for what a theory of ignorance calls for. It is rooted in the conviction that a philosophical theory of cognition that slights or outright ignores what everyone already knows is headed for a bad outcome.

2 The cognitive economy

Human organisms are several things at once. We are natural objects. We are pieces of biochemical organization. We are ecological beings in constant reciprocal interaction with our environment. We are social beings connected to complex networks of fellow beings in multiply interdependent ways. We are information-processing systems, at times coding up to fully conscious belief and desire, and an active participant in our community's economies. *Money* economies are interactive information-systems for the generation and circulation of wealth. *Cognitive* economies are interactive information-systems for the generation and circulation of knowledge. In matters both ontic and epistemic, the cognitive economy is prior. Participants in the cognitive economy are kitted up with *cognitive agendas* as standard equipment. In explicit form agendas are attack plans for knowledge. In implicit form, they are "opportunistic", that is, they are causal dispositions to extract cognitive benefit from information-flow. Attack agendas are comparatively expensive to discharge. Opportunistic agendas are considerably less so. There is a significant difference between striving to find things out and having them bestowed simply by showing up and making the scene, and having one's information-receptors in good working order.[8]

A healthy economy produces and circulates its goods in the markets in which it operates. A sick money economy under-serves demand and miscirculates its products. It will die if its production capacities buckle or its circulation procedures seize up. Knowledge economies are like this too. A knowledge economy that can't meet the demands of its markets will crash. If it crashes on a broad enough scale, it will take the species *homo sapiens* down with it.[9]

Space does not permit a full exposition, still less the complete defence, of the theory of knowledge that I take to be best-attuned to the empirically discernible realities of the human cognitive economy. The theory of choice is a variation of Goldman's causal reliabilism, in what has come to be known as the *causal-response* (CR) model of knowledge.[10] To help ease our way to a discussion of the four grades of knowledge and their implications for ignorance, I shall advance the key features of my variation of the theory in question also as working assumptions. It differs from Goldman's model by eliminating the justification-condition as a general requirement for knowledge and causalizing the belief-condition instead. The CR- model

[8] For a discussion of how abductive reasoning operates in cognitive economies, readers could consult my (2012; pp. 148–161).

[9] For a good appreciation of the ups and downs of naturalizing the investigation of reasoning, see Urbanski and Klawiter (2018; pp. 583–597; Sect. 3), "What about the cognitive turn in logic?".

[10] Goldman (1967; pp. 357–372). Woods (2013). A summary account appears as "Logic naturalized", in 2016. See also Magnani (2015; pp. 13–36, 2018; p. 44).

is a senior partner in the naturalization of inference project, set out in my *Errors of Reasoning*. Informally expressed and with some detail omitted, here are the key assumptions: Human beings are fallible, yet also abundant and versatile knowers. Although we are error-committing beings, we have the interest and wherewithal to detect and correct error. There is an adaptive advantage to this. It off-loads the impossibly heavy cost of error-avoidance no-matter-what to the more economical and amply realizable methods of *post facto* repair. Epistemology's heaviest burden of proof falls on philosophers who deny or downplay these empirically discernible realities. The heaviest burden is borne by big-box sceptics, who deny our knowledge of the external world or other minds, or causality or the past, and on and on.[11] While the proof burden is heavy, the proof standard is even higher. I take it as given that it is standard that big-box sceptics have yet to meet. Until the sceptics up their game, it is open to a theorist of knowledge to respect "what everyone knows."

The causal-response model is a rival of the command-and-control (CC) model of knowledge, of which the justified-true-belief (JTB) model is the dominant exemplar. On the CC-model, beyond sense-perception itself, knowledge is the product of the knower's intellect which operates as the executive authority of knowledge-acquisition. Overstating it slightly, according to the CC-theorist, the successful knower has to be smart. On the CR-model, smartness is valued when it is needed for knowledge, but it inclines to the view that, in the main, what's needed for knowledge is good information-processing devices working in the way nature designed them to work. Certainly, they work well enough to enable the detection of another empirically discernible fact about how the human cognitive economy functions:

> *Enough, enough, enough, enough:* Beings like us are right enough enough of the time about enough of the right enough things to survive, flourish and occasionally build great civilizations (Hereafter "Enough^4")

One of the benefits of the suppressing of big-box scepticism is the way it facilitates the normative assessment of human cognitive practice without having to resort to the fiction of ideal rationality in models of perfection that no human being would have any chance of instantiating. Such models are the safe harbours of scepticism. They purport to define what good reason really is, to explain why beings like us are so bad at it, and to rub it in by declaring its ideal standards to be normatively authoritative for us. By putting sceptism aside, and selecting Enough^4 as our guide, we'd have a decent shot at showing that when reasoning is in compliance with that condition is reasoning that operates in the way that reasoning should. We would have a grounding for the thesis that, except when there is good cause to say otherwise, reasoning is good when it is reasoning that normally occurs. This is the *normative-normal convergence thesis*.

Some philosophers grant the survival, flourishment and civilization building, but deny that it takes knowledge or true belief to bring them about.[12] My own view is that this is another case of big-box scepticism that doesn't meet its proof-burden.

[11] See, for example, Unger (1975).

[12] See, for example, Stich (1990).

The core finding of the CR-approach can be informally set out as follows:

CR-knowledge: Agent X knows that *p* at *t* on information *I*, if *p* is true and in processing *I* at *t*, X's belief-forming devices are causally stimulated to produce the belief that *p*, X's cognitive devices are in good working order – and working here in the way that they should; *I* is good and properly filtered information, and there is no disabling interference caused by negative externalities.[13]

Information is good when it is accurate and up-to-date. It is properly filtered when it satisfies the Harman Clutter Avoidance Principle: *Do not clutter up the mind with trivialities*.[14]

It is clear upon inspection that for large ranges of cases and by far the majority of them, the processes that satisfy the conditions set out here operate out of sight of the eye's mind, beyond the reach of the heart's command and unnnegotiable by tongue or pen (or keystroke). To prosper in the cognitive economy we have to have good equipment. But having it is different from knowing what it is and how it works. Were it otherwise, we could have a conscious command of the details. It isn't in fact otherwise; and even if it were, having that knowledge would not be a condition on good cognitive performance. This seriously motivates an epistemology that takes official notice of the fact that much of our knowing is brought to pass unconsciously in what I'll call the cognitive down-below. Lots of what we know but, comparatively speaking, *not much of it overall*, is acquired in the cognitive up-above. These are matters that require further attention if the objectives of this paper are to be met. Before quitting this section, it would pay us to heed the cost of justificationist epistemologies and their link to big-box scepticism. On the facts of the matter as we have them here, beings like us are much more favoured in their knowings than they are in their belief-justifyings. If that were so, the imposition of a J-condition on knowing would eliminate most of what we currently experience ourselves as knowing. Knowledge, so construed, is very expensive. Expenses that high are a direct link to big-box epistemic scepticism.

3 The cognitive down-below

Not all information-processing occurs in the cognitive up-above, that is, in the full light of conscious awareness. Most of the information that a human organism will process will be processed in the cognitive down-below. Information-processing down-below has in varying degrees all or most of the following properties: It is mechanism-centred, unconscious, automatic, inattentive, involuntary, non-linguistically structured, semantically inert, deep-down, parallel and computationally luxuriant. Information-processing up-above has also in varying degrees: all or most of the

[13] As formulated here, the conditions are most naturally and accurately read as characterizing propositional or semantic knowledge. It is also adaptable to the other grades of knowing, but the present formulation will do for illustrative purposes.

[14] Harman (1986; p. 2). In Gabbay and Woods (2005b; pp. 398–417).

following properties. It is agent-centred,[15] conscious,[16] controlled, attentive, voluntary, linguistically structured, semantically loaded, surfacely contextualized, linear and computationally puny. (Shiffrin 1997)[17] Reasoning down-below, and knowledge too, is interesting. When it happens it happens, as we said, out of sight of the mind's eye, beyond the reach of the heart's command, and unnegotionable by tongue or pen (or keystroke).[18] It is important to emphasize that when these properties are instantiated, they need not be instantiated in equal degree. For example, as we might expect, it is true that we are conscious of anything that's caught our attention, but we don't pay attention to all the things we are conscious of.[19]

Much of the causal traffic in the human down-below is entirely subcognitive, a matter of energy-to-energy transductions. At some level, the causal flow converts energy-to information, at which point primitive cognitive contact is made. Once information enters the picture, the causal flow starts to take on a productive potential for inference and, in time, produces tacit and implicit knowledge. Under the right conditions some of it breaks the surface into the light of conscious day. It is an understatement that there is much yet to learn of the transitions from energy-to-energy transductions and energy-to-information conversions to the knowledge produced by Wiles' elliptical curves proof. One of the unsolved puzzles is that knowledge is an information-thirsty state to be in, but for most of everything we know at *t* there isn't room for in the conscious mind at *t*. Of course, no one thinks that the holding capacity of the conscious mind kills the knowledge that's not presently in it. So somehow the knowledge is preserved and stored in the down-below. In an influential book, Fred Dretske popularized the idea that belief is information in completely digitized form, where all else is in analogue form. He is certainly correct in seeing that information down-below is differently structured from information up-above and yet, even so, that information down-below is capable of conversion to information up-above. There is reason to think that analogue-digital conversion doesn't accurately capture the essence of belief, but there is no doubting the value of the differently-structured thesis and its attendant down-up conversion claim.[20] This brings us now to a central claim of CR-epistemology. It adopts a key notion from thermodynamics.

> *The phase-transition thesis*: Information down below is subject to phase-transitions from one state to a qualitatively different state up above, and also capable of reverse phase-transition back down. In the passage upwards, information

[15] Later I will expand this notion (with suitable adjustments) to *multiagents*.

[16] The philosophical literature on consciousness also produces a deluge of rivalrous opinion. Two recent volumes of note are Paglieri (2012), and Jacquette (2018). It is a matter of interest that in neither of these volumes does ignorance appear in the index. Also of interest are Bachmann (2000), Breitmayer (2014), and Aru and Bachman (2017; pp. 128–135).

[17] Further details, can be found in chapter four of *Errors of Reasoning*. See also Doya et al. (2011), and Hohway (2013), Zimmerman (1989; pp. 166–175), Shiffrin (1997; pp. 49–64, 1976; pp. 177–236), and Mole (to appear).

[18] See further, Wilson (2002), Wegner (2001) and Carruthers (2011).

[19] This common sense view is nicely supported in Mole's (2008; pp. 86–104).

[20] Dretske (1981). For reservations about the treatment of belief, see Gabbay and Woods (2003; chapter 7, Sect. 7).

> loses properties and gains opposite ones. On the way down properties acquired on the way up are lost and their opposites regained. In a more antique formulation, when information is in phase-transition, it retains its haecceity and loses its prior quiddity in acquiring a different one.[21]

One might ask whether there can be attention down below without consciousness. The answer is yes. Biological organisms are responsive to calls or signals from their environments, including other organisms. Doing so facilitates their management of environmental complexity. On the other hand, it hardly seems possible that all biological organisms are subject to the phase-transitions that human knowers are subject to.[22]

Since knowledge is a storable commodity, it must have features, beyond unconsciousness, that enable it to *be* stored. It is known that in thermodynamically closed systems, consciousness is an information-suppressor. In the sensorium, the intersection of the five senses, ≈11 million bits of information are processed each second. If the information is processed consciously, the count plummets to ≈40 bits per second. If linguistically formulated, the count drops to ≈16. (Zimmerman 1989; pp. 166–175). No one thinks that the human information system is thermodynamically closed. But the fact remains that there is massively more that we know than we can ever get our conscious minds around at any given time. What is it about consciously held knowledge that makes it so hard to keep at the front of our minds? If we don't store knowledge in the conscious sector of the head, what is it about its unconscious sector that makes for its more capacious storage capacity? It is not the comparative sizes of the head's two sectors. The human unconscious is not an amazon.com warehouse. Still, this is not something that need deter us for long.

After a bit of reflection it becomes clear that the difference as stark as that between all that we know at *t* and all we can be aware of it at *t* cannot be explained by limitations of size. The better explanation is that properties borne by knowledge when it is consciously had, are erased when stored down below. We know some of the properties of consciously held knowledge. We know that the information that underlies it is semantically structured, truth-evaluable and linguistically expressible. It could be argued that properties such as these are storage-busters. It is not that the conscious mind is so small, but rather that what's in it is too costly to store in quantities that match its own size. By "costly", I mean that full-presence would crash the system. On that hypothesis, these are the properties that don't make the cut

[21] A phase transition in physics is a reversable change in a substance from a given state (e.g. liquid) to a qualitatively different state (e.g. gas) at a specific combination of temperature and pressure. See for example, Callen (1985; chapter 9). There is a link between Aristotle's concept of potentiality (*dunamis*) and the phase-transitions of modern physics. For Aristotle, *dunamis* is a thing's capacity to take on a new form without losing its identity. See *Metaphysics,* Book 8, 1 1046^{a} 12, 1048^{a} 25, 27. Frege thought the same about numbers. He thought that they could take on the cross-type form of sets without losing their identities. An item's haecceity is that in virtue of which it is its self-same thing and not another thing. A thing's quiddity is that in virtue of which it is the very kind of thing it is. Then the phase-transition thesis has it that there are ranges of cases in which a thing's haecceity is unmolested by change and restoration of quiddity. An excellent modern treatment of kinds is Pelletier (2010).

[22] See here Orzumi et al. (2018; pp. 1–25).

in the passage to the vaults below, with obvious implications for tacit and implicit (TI)-knowledge.

Everything known in the down-below—in its tacity and implicity—broadens the scope of not knowing in those other ways. It is true that the more TI-knowledge we have the less TI-ignorance we have. However, most of our TI-knowledge remains below and in any case, only parts of it are occasioned to surface upwards. Accordingly,

> *Limits of semantic knowledge*: All of what is TI-known at *t* is semantically unknown at *t*. Semantic ignorance varies inversely with TI-knowledge.

We should also remind ourselves that most of the plenitude of what is actually the case never achieves even TI-landfall and remains at two removes from semantically structured overt awareness. It bears repeating that there is massively more truth than will ever be known in any sense of that term. How much of the corresponding plenitude of ignorance bears in a material way on the health of our cognitive economy? In systematic terms, hardly at all it would seem. A cognitive economy would soon sicken and expire if it over-traded in cognitive products for which there was neither need nor demand. It would be counterproductive to press to know things simply because they aren't yet known. Knowledge quests are subject to value-in-having assumptions. By and large, the things that beings like us aspire to know are things it would do us some good beyond the mere fact of knowing them. Think here of the exact number of well-individuated strands of hair on Caesar's scalp in the run-up to his assassination in 44 BC, as opposed to the information which, had it arrived on time, would have spared his life. This is why cognitive agendas are standard equipment.

4 Information

Knowledge is belief extracted from information under conditions in which the belief is true. An important point of difference with justificationist frameworks is the large role accorded by the CR-approach to the cognitive productivity of unconscious information-processing. CR-epistemology is heavily invested in the notion of information and pays circumspect heed to the well-supported findings of the informational sciences, such as they may be. All the same, proceeding informationally is not an investment without risk. It has rightly been observed that "[o]urs is supposed to be the information age, and we all share, in different degrees, the problem of coping with a deluge of information flooding over us."[23] The concept of ignorance is at least fourwise ambiguous. In its *epistemic sense*, information is a representation of what is the case. It tells us how things actually are. In its *probability* sense, information is what it is channelled from a source to receiver. The source emits signals with a certain frequency, and the information picked up by the receiver is conceived of

[23] Jaakko Hintikka, "Who has captured the notion of information?", in his *Socratic Epistemology,* pages 189–210; p. 189.

as the expected reduction of probabilistic uncertainty. In intuitive terms, the more information a signal carries, the fewer its interpretations and, correspondingly, the more it shrinks its attendant ignorance-space. Agency is needed for the reception of semantic information and is often is transmitted by agents. Agency isn't needed in the probabilistic case.

In its *complexity* sense, information has to do with codes. The information-value of a code-string is the algorithmic of Kolmogoroff complexity of the string, which can be defined as "the shortest program that computes it on some fixed universal Turing machine."[24] Finally, information in the *military* sense derives from its use in intelligence and counterintelligence work. According to the CIA World Fact Book:

> "Information is raw-data from any source, that might be fragmentary, contradictory, unreliable, ambiguous, deceptive, or wrong. Intelligence is information that has been collected, integrated, evaluated, analyzed and interpreted."[25]

Informational content is big in the probabilistic sense to the extent the signal that carries is subject to few readings. Its consciously epistemic sense, although *scant*, is likewise subject to few readings. Epistemic and probabilistic information are one another's contraries, but they hover near equivalency at receptor-termini.[26]

It is easy to see the tangled state in the ways theorists conceive of information, prompting

Hintikka to observe that it is

> "far from clear … what (if anything) is meant by these different 'informations' – or whether they are related to each other at all. These questions seem to mark a most urgent challenge to philosophical analysis." (*idem.*)

Nudged by the idioms of the cognitive up and down, in *Errors of Reasoning* I advanced the idea of the human being as an "information stack".[27] If we accept the thesis that humans are causally structured beings, it is easy to see a place for energy-to-energy transductions and energy-to-information conversions operating at varying distances from the up-above where, at times,—indeed at every moment of wakefulness—information surfaces. In a caricature, it takes a lot of causality to provide for a cognitively healthy human being. Oxygen must flow, digestion must work, neurons must fire, information sensors must activate, and knowledge must be sought and in

[24] Quoting Adriaans and van Benthem (2008; p. 12). A problem pressed by the first two kinds of information is that epistemic information can't not be true, whereas probabilistic information needn't be true. There is, however a route from the complexity conception to the probability conception by way of the set of all prefix-free programs under provisions of Kraft's Inequality.

[25] www.cia.gov/library/publications/the-world-fact/docs/history.

[26] Here is Godfrey-Smith's summary of Dretske's view of the matter: "Information, for Dretske, is found where there is contingency and correlation. Any variable in the world which has a range of possible states is a source of information. When a state of information is correlated with the state of another variable, as a consequence of physical laws, the second variable carries information about the source. For Dretske, information is a resource that organisms use to make their way through the world; cognitive systems are information consuming, or information exploiting, systems." Godfrey-Smith (1996; pp. 236–237).

[27] *Errors of Reasoning,* chapter 6, Sect. 9.

some good measure found. It could be that all four of our present information-kinds are productively implicated in an up-and-down causal stack. The two that play the most conspicuously cognitive role are the first and the fourth—accurate semantic information, and the masses of raw data which make up much of "the deluge of information flowing over us." When we bring Harman's clutter avoidance maxim into play and apply the Gabbay-Woods filtration-systems to give it operational effect, we'd have a fair chance of connecting the fourth to the first:

> *Good information as well-filtered raw data*: Good information is raw data in the "military" sense when screened by G-W processing-filters to bring information into compliance with the requirements of clutter avoidance. When this happens, data convey *intelligence*, which is the security community's word for knowledge.

For all its semantic unruliness, it is Hintikka who insists that "the notion of information should be everybody's concern." (*idem.*) He continues: "Our life is increasingly dominated by computers, which are nothing but machines for processing information. Information can even serve as a commodity (utility) in economics and decision theory. Hence it is important for each of us to master this concept intellectually and to have ways of gaining an overview [of] the different kinds of information we receive". Indeed, "information rather than knowledge or belief should be the most basic concept of epistemology" (p. 190).[28] In chapter one of *Socratic Epistemology,* Hintikka proposes to do epistemology without knowledge and without belief (11–37). Speaking for myself, I concede primacy to information, which in partnership with causality, secures a safe place for belief and knowledge.

5 The first grade of epistemic involvement: anselmian knowledge

Speaking of what everyone knows, the mass-energy equation of relativity theory comes to mind. Of course, it is not literally the case that everyone knows that $E = mc^2$; many millions more know that water is wet. Still, while some billions of people know the equation, massively fewer of them know what it means. Some know that "$E = mc^2$" formulates a principle of physics.[29] Fewer know the gist of it, namely, that anything having mass has an equivalent amount of energy. Fewer than they know that the increased relativistic mass of a body times the speed of light squared is equal to the body's kinetic energy. As we see, knowledge of the first grade tends to come in degrees, ranging from knowing that S is a true sentence without any understanding of what it means, to knowing that it is true with some understanding of what one or more but not all of its terms mean and some general understanding of its subject-matter.

[28] Hintikka wrote this in ≈ 2007. A scant 11 years later, it strikes us how understated his admonitions were.

[29] Einstein (1905; pp. 639–643).

St. Anselm of Canterbury (1033/4-1109), famous for his ontological argument for God's existence, laid down in another writing an important principle about knowledge: *Credo ut intelligam,* which I translate as "I believe [it] so that I might come to understand [it]." On this occasion, Anselm was speaking of matters of religious doctrine, not the theory of special relativity. Even so, his remark easily adapts to the points of a paragraph ago. Most people by far who know that the sentence "$E = mc^2$" expresses a true proposition don't understand the proposition it expresses. They lack a semantic grasp of it. They lack an awareness of what it actually says. Consider how that proposition could be gathered up in a fuller understanding. Anyone seeking one could achieve it if he signed up for a first-year college physics course or senior year physics in a good high school. But unless the student believes that "$E = mc^2$" expresses a true proposition, he greatly lessens the chances of his ever understanding it. This is the epistemological message of *Credo ut intelligam*, and with it comes the first grade of epistemic involvement. In honour of the philosopher who first called it expressly to mind, I will dub it "anselmian knowledge"—standing mute on any expectation of theological elucidation of, such as may be, matters divine.[30] We also learn from anselmian knowledge something *pedagogically* important.

In his *Set Theory,* Kenneth Kunen opens with the following sentence: "Conventional mathematics is based on *ZFC* (the Zermelo-Fraenkel axioms, including the Axiom of Choice)."[31] Some of Kunen's readers will have had some prior knowledge of this fact. Others will have had none. Some will know the ZFC axioms but won't be able to call them to mind now. Others will have heard of the axiom of choice but not gotten around to learning what it is. Some will be dab hands at elementary set theory, but actually might not know what's meant by "conventional mathematics".[32] There are two lessons to learn from this. One nicely illustrates Nicholas Rescher's notion of *deficient* knowledge (cited in *Ignorance*'s subtitle). The other is that a student who lacks an anselmian knowledge of the opening sentence doesn't know that it expresses a true proposition and has no occasion to believe that it does. So long as his anselmian ignorance persists, the student will have the damnest of hard times in learning set theory. The point at hand generalizes widely, encompassing most of what eventually comes to be known semantically. No one will achieve a solid knowledge of reality, unless he comes to a timely knowledge of what everyone of age two knows from mum and dad and by bumping into things. No one will know what everyone knows at age ten without having been told things by the several others of his acquaintance and, in many cases, having gone to school. And so it goes, from kindergarten all the way up to Sciences Po. Here, too, we have a direct proportionality claim, but one that throws a larger shadow.

[30] I omit the upper case on the first letter of this name to discourage the misunderstanding that the present section is an exercise in Anselmian scholarship.

[31] Kunen (2013; p. 1).

[32] For example, we may take it that dialethically inconsistent mathematics is not of the conventional kind; but what about intuitionistic mathematics?

A drag on semantic knowledge: If one's knowledge of *p* at *t* is exclusively anselmian, one's semantic, tacit and implicit, and comprehensive *ignorance* of *p* at *t* is guaranteed.

The pedagogical necessity of anselmian knowledge: In practical terms, anselmian knowledge of *p* is a necessary condition of acquiring both semantic and impact- knowledge of *p*.

Still, there are off-setting considerations

Ignorance as causal stimulation: The ignorance arising from what an anselmian knower doesn't know semantically often induces the pursuit and facilitates the acquisition of a fuller knowledge, thus making it an affordable cost.

Come back now to the point that anselmian ignorance can be a serious impediment to a knowledge of richer grade. *Persisted in*, the cognitive economy would totter. This presses us to ask how anselmian knowledge is acquired. It is, in the most substantial of ways, acquired by being *told it.* I'll come back to this in Sect. 8.

6 The second grade of epistemic involvement: Tacit and implicit knowledge

A friendly notice: In this section, I make no effort to accommodate all the meanings of "tacit" and "implicit" in everyday speech. My purpose is to confine attention to those features that matter most for the theoretical treatment of information-flow in the cognitive down-below.

Since propositional knowledge is what the anselmian knower doesn't know, it may strike us that, *in ordo essendi,* it is propositional knowledge that marks the second grade of epistemic involvement. Ranking here is a tricky matter and, without prejudice, I'm going to offer up a different candidate for second-grade involvement. In the end, rank won't matter all that much, but we would be remiss in not considering the foundational importance of knowledge that's tacit and implicit. Besides, on present assumptions, tacit and implicit (TI-) knowledge have something in common with knowledge of the anselmian kind. They both fall short of propositional or semantic knowledge, albeit for different reasons. There is, however, a consideration that breaks the tie between them. One can't have the anselmian knowledge that $E = mc^2$ without having the wherewithal to recognize sentences and to form the belief that they express true propositions. TI-knowledge is different. As we saw, when someone is in a state of tacit and implicit knowledge, his knowledge *unnnegotiable* by tongue or pen (or keystroke). So, *strictly speaking*, one does not have anselmian knowledge of it either.[33]

[33] It is important to keep in mind that the four-grades thesis is itself gradualistically instantiated. It is possible for an anselmian knower of *p* to have known its semantic content earlier. Consider two cases. In case one, *p*'s semantic content reposes in stirrable memory, if only thanks to some external stirrer. In case two, there is nothing in the anselmian's memory that is semantically retrievable. In the first case we have a faded memory of something once grasped. In the second we have amnesia, that is, *savoir perdu.*

Consider the case of a reader of Kunen's *Set Theory,* who's looking help in refreshing her once-abundant knowledge of sets. When she comes to Lemma III. 3. 42 on page 183, she reads:

"Let $\mathbb{P}$, $\mathbb{Q}$ be posets with $\mathbb{P}$ acc and $\aleph_1$ a pre-caliber for $\mathbb{Q}$. Then $\mathbb{P}$ x $\mathbb{Q}$ is acc."

The seven lines that follow provide the proof. By the time she'd got to p. 183, she had no trouble understanding the wording of the lemma. She came to its proof without a hitch. Later that day, she tells her husband, "It was a piece of cake. I knew it all along. I just couldn't bring it all together and get it down on paper." Nearly everything a professional mathematician knows of mathematics falls into this category. There are masses of it that they know but, without aid, can't quite pull together and put into words. The gap between knowing and putting into words is comfortably filled by anselmian knowledge. Even so, these knowers aren't in the condition that the anselmian knower is in when all he knew of Lemma III. 3. 42 was that it expresses a true proposition of mathematics. The state that our anonymous reader is in is the state that anyone is in when something of which he's had consciousness awareness makes its way to the depositories of memory in his cognitive down-below. The state that he's now in is the state of knowing it tacitly and implicitly. When in that state, his knowledge is deficient, but is so in ways that might appear quite different from the Rescherian condition of not being fully complete. The appearance is deceiving.

According to the phase-transition thesis, the deficits here are the absence of semantic content and linguistic structure. It soon becomes clear that this is *not* an abnormal state for a knower to be in. Most of what we now know we couldn't collect together and put into words. There are conditions under which a bountiful memory summons up its secrets, and similar conditions under which we are able to call to explicit fore some of the details of the vast reaches of background information and common knowledge. But there are no conditions in which all of our knowledge surfaces, nothing remotely close to it.[34]

If we hold fast to the CR characterization of knowledge, we see that there can't be TI-knowledge unless there is TI-belief. Neither can there be knowledge unless the imbedded belief is true. But if TI-knowledge and belief lack semantic or propositional content, what would it take for the TI-beliefs of TI-knowledge to take the truth-value true? It is a good and necessary question. The answer is that they don't take it, and cannot. They are not truth-valued. They are *virtually* truth-valued.[35]

[34] An essential feature of the English common law is the inexpressibility of court-created precedents. Precedents play an essential role in determining future cases; but when judges and scholars try to write them down and make them explicit, invariably they get them wrong. Further discussion can be found in chapter five on "Unwritten law" in the second edition of my (Woods 2018a).

[35] When we grasp that 2+ 2=4 and put it into words, what we've said comports with Convention T: 2+2=4 if and only if "2+2=4" is true. When this knowledge is stored below, it loses its semantic and linguistic character and therewith its capacity to oblige Convention T. But if it has occasion to resurface, it will recover the truth-value T.

Virtual truth-values: A TI-belief *p* takes the virtual truth-value true (or false) just in case upon surfacing to the semantically enriched up-above, the truth-value it literally takes there is true (or false). In acquiring propositional form, it requires a truth-value.

Corollary: When *p* phase-transitions down below, it reacquires its virtual truth-value and retains its capacity for literal acquisition upon re-transitioning upwards.

Preserving semantic content: Since real truth-values apply to sentences with semantic content, and the same semantic content is recovered in further resurfacings, we see that semantic content is preserved under phase transitions.[36]

In the absence of semantic content and linguistic structure, knowledge of the implicit and tacit can't be propositional knowledge. So there are two ways of knowing that fall short of propositional realization. One way is anselmian and the other is the way of the tacit and implicit. A principal difference is that under the right conditions, implicit and tacit knowledge is sometimes responsive to stimuli that explicitize it and give it propositional voice. Applied to anselmian knowledge those same stimuli are unavailing. When a novice learns set theory for the first time, he has no prior memory to refresh. When a once-learned reader reads the same book the novice does, her memory might be given a jolt and send to the surface the knowledge she once had in explicit form. The tacit and implicit carry obvious consequences for ignorance.

The ignorance-making provisions of TI-knowledge. Everything of which a person has implicit and tacit knowledge at *t* he is propositionally ignorant of at *t*. The reason why is that tacity precludes expressibility and implicity suppresses propositional meaning.

There is a growing philosophical literature about TI-knowledge. In earlier days, it appeared in the writings of Wittgenstein, Ryle and Heidegger. Somewhat later it attracted Polanyi's careful attention.[37] The gist of these writings is that TI-knowing is knowing *how*, not knowing *that*. I myself am more partial to the phase-transition view that TI-knowledge is information's causal potential to surface upwards and, in so doing, to take on the semantic and linguistic character required for propositional knowing in the cognitive light of day. As far as I can see, this potential both underwrites and bisects the knowing-how/knowing that distinction. After all, when it does surface, knowing how to X is knowing that it takes such-and-such to X.

It is tempting to think that TI-knowledge is knowledge that falls short in some way. It does, of course. It falls short of formulable explicity. But what matters more is its net contribution to the cognitive economy. It is an *indispensable* contribution,

[36] Similarly, if a piece of information is semantically loaded and now transitions downwards, it loses its content. However, upon transitioning back up, the old content is recovered.

[37] Ryle (1929–1930; pp. 91–126, 1945–1946; pp. 1–16, 1949). Wittgenstein (1953). Heidegger (1962). Polanyi (1958, 2009). See also Lowney II (2017). A good survey of these and Polanyi's works can be found in Gascoigne and Thorton (2013).

in whose absence there wouldn't be room for semantic knowledge to breathe. If we accept the assumption that only a knowledge as capacious as the TI-kind is able to prop up cognitive economies as capacious in turn as our own, its indispensability to good cognitive health speaks for itself. *Res ipsa loquitur.* If we accept the further assumption that TI-knowledge is stored as *potentialities* for causally induced semantic realization up above, we'd have a decent answer to the question of how beings of such a comparatively slight architectural capacity can absorb the demand of multiples of millions of lines of code to computerize the information systems on which cognitive economies such as ours depend.

This calls to mind the injunction to theorists to respect what everyone knows, in their efforts to contrive theories that account for what not everyone knows. It raises the question of what it takes to respect this call. If, as I propose, most of what any human will know is packaged in the form of potentialities for possible semantic realization on a much smaller scale, the answer for the most part would seem to be that the knowledge implicated in what everyone knows is TI-knowledge.

A word now about common knowledge and background information. Common knowledge is what everyone in some contextually conditioned way knows. What's common knowledge in Vancouver in 2019 differs in important ways from common knowledge in Samarkand in 2019, to say nothing of 750. Some of what is commonly known is properly global; e.g. that water is wet. As it is in the English common law, here too there are gradations of commonness. Background theories are common knowledge localized to scientific objectives. They remind us of the cognitive impossibility of launching a theory entirely *ab initio*. In all such matters, as Quine says, we are *in medias res.* The point of greater moment for what concerns us here is that, in every way that matters, background information is common knowledge, and inherits its key characteristics. Accordingly,

> *Ignorance-saturation for good cognitive health*: The cognitive economies of beings like us are made both possible and fruitful by the ignorance implicated in TI-knowledge.
>
> *The economic benefits of implicity*: Explicity is a knowledge-inhibitor – much in the way that consciousness is an information-suppressor – and therefore and correspondingly an ignorance-abettor. The cost of this deficiency is borne by off-loading the semantically explicit to the non-semantic implicit, where it remains causally efficacious and subject, in good part at least, to semantically refreshed recall.

7 The third grade of epistemic involvement: Semantic knowledge

As we turn to the third grade of epistemic involvement, we should say a brief word about another growth industry. It concerns what we are to make of what's grasped when someone understands the thing an indicative sentence gives expression to. As recently pointed out by Bradley Armour-Garb and James Woodbridge, there is a hefty abundance of quite confident "proposition-*talk*", yet little consensus about what's being talked *about.* I have no time for the proposition wars, both literally

in the sense that I have to keep an eye on page-count, and also in the sense that the issue is framed by assumptions as naïve as those that made the commonplaces of how to make cognitive contact with numbers—even the inaccessibly large ones—into something as silly as Benaceraff's "dilemma". Armour-Garb and Woodbridge also give up on the proposition fight and settle for noncommital proposition-talk. Others of a more rationally reconstructive bent propose that we conceive of propositions as "probability contents"—sets of probability spaces, each of which in turn is a set of worlds and a probability function over them. So seen, they are neither objective nor subjective. They are nothing more than abstract objects of mathematics.[38]

The problem of content also migrates to belief, which plays a central role in the CR-response model, hardly less so in the JTB one. It is fertilized by rivalrous rumbles about consciousness. As with propositions, I have no room for our present dishevelments about belief, short of remarking that I am provisionally drawn to a Ramsey- and Peirce-like conception of it, in which beliefs are dispositions to act. If we accepted "the rough equivalence of mind and brain", we might say that they are also "those neural states that encode the information" to which action is causally responsive.[39] An especially interesting way in which propositional understanding intersects with ignorance is nicely capturable in Peirce's provisions for abductive inference. Let's turn to this now.

There are two ways in which ignorance makes the scene in epistemology. One is the traditional way in which ignorance is inferred from scepticism-implying premisses. The much more customary way of introducing ignorance into the general go of things in the cognitive economy is Peirce's way. Some one has an ignorance-problem when there is semantic knowledge he would like to have but currently lacks and knows he does. There are three main responses to ignorance-problems.

Subduance: New knowledge is acquired that solves the problem.

Surrender: The agent in question removes this unmet target from his current cognitive agenda.

Abduction: A semantically formulated hypothesis is abductively inferred.

The first way, the agent overcomes his ignorance. The second way, his ignorance overcomes him. The third way, the ignorance is not overcome, but a reasoned inference is drawn to guide action in a cognitively promising way. It is commonly said that the signature property of deduction is truth-preservation and, of induction, is likelihood-enhancement. Abduction stands apart. Even at its best, abductive inference is *ignorance-preserving*.[40] This is not the place to lay out the logical structure

[38] See, in the first instance, Armour-Garb and Woodbridge (2015; pp. 84–90), and, in the second, Moss (2018; chapter 1).

[39] Zimmerman (2018; p. 1).

[40] The idea of ignorance-preservation first arose in my "Remarks on the logical structure of abduction", presented at the Symposium on Abduction and Creative Inference *Cog Sci 2004,* Chicago. It was developed further in Gabbay and Woods (2005a); chapter 3 on "The structure of abduction". An updated treatment can be found in my (2012; pp. 148–161). For an update of that update, see my (2017; pp. 137–149). The Gabbay-Woods model is given an attractive pragmatic dynamicization in Chiffi and Pietarinen (2018).

of abductive inference in all the detail it demands. It will suffice to expose how not solving an ignorance-problem can nevertheless be cognitively nourishing.

Although the originator of the idea of abductive inference is Aristotle, its principal modern source is Peirce, who captures a good part of it in the schema that bears his name.

Peirce's schema

(1) The surprising fact *C* is observed.
(2) But if *A* were true, *C* would be a matter of course.
(3) Hence there is reason to suspect that *A* is true. (*CP* 5.189)[41]

The key features of the schema can be set out as follows: The correctness of the inference provides *no* reason for believing *A* or evidence that it's true.[42] Therefore, if the surprising fact *C* poses an ignorance-problem, the abductive inference that there is reason to *suspect A*'s truth leaves the problem unsolved. The inference is therefore *ignorance-preserving*.

Peirce discusses this issue in the context of experimental science. He took it that the conclusion of a properly abduced inference provided a reasoned basis of action, namely, to put the hypothesis *A* to experimental test.[43] In our adaptation of Peirce to non-experimental contexts in the G-W model, Gabbay and I proposed that the course of action to follow upon the abductively derived hypothesis *A,* is to put it to *provisional premissory use* in the context of enquiry within which the ignorance-problem arose in the first place. Should the provisionally released proposition *A* resist efforts to rebut it, and also facilitate the derivation of well-confirmed outcomes, it could in time quite reasonably be taken for true.

Although Peircean surprises can be striking, shocking, unexpected, newly-arrived, and puzzling, they are not surprises on those accounts. A Peircean surprise is something off-course, out of line, aberrant or anomalous. Surprises are not the way things normally are. When a well-abduced hypothesis removes the element of surprise, it does so in virtue of the fact that, if true, the fact in question "would be a matter of course." Peirce sometimes says that when this subjunctive conditional holds true, it makes the surprising fact explicable. But he does *not* mean by this that either the hypothesis or the fact are now well-understood.[44] Neither do we have it that as successful abduction makes the hypothesis in any way plausible or conditionally more likely or justified. What it warrants is the suspicion that in selecting the hypothesis, the abducer has made the right guess, which is innately a matter of instinct. (*RLT*, p. 128, *CP* 5.171, 7.220)

[41] Peirce (1931–1958; volume 5; p. 189). Line numbers are mine.
[42] Peirce (1992; p. 128; *CP* 5.171, 7.220).
[43] *CP* 5.59, 6.46–6.473, 7.202–219.
[44] For a recent discussion of how Peirce's model as extended by the Gabbay-Woods schema allows for successful abductions that aren't inferences to the best explanation, see *The Reach of Abduction*, chapter 5, Magnani (2009; chapter 2), Park (2017; Sect. 4; chapter 2), and Minnameier (forthcoming).

Come back now to the CR model of knowledge. Imagine that we are now at the stage at which a properly abduced *A* has become reasonably believable in the community of enquiry in which the ignorance-problem first arose. Let X be a fully qualified member of that community. Then we have it, do we not, that *A*'s status there gives cause to believe it true, X's belief-making devices are in good working order and functioning in this instance as they should, then X knows that *A* if *A* chances to be true?

> *The epistemic clout of ignorance-preservation*: If we allow that generally speaking, the more reasonably a proposition is taken for true, the greater the likelihood of its *being* true, we can begin to see the epistemically productive role played by ignorance-preserving responses to problems initially posed by surprising facts.
>
> *Abductively generated knowledge*: What is more, by the lights of the CR model, propositions meeting these conditions are with some frequency known to be true – that, is true and generated by *well-produced* belief on good information.[45]

By this same reasoning, there is reason to think that Fermat knew *without proving it* that there is no solution to $x^n+y^n=z^n$ in the positive integers for any $n>2$. If so, we might think that it bespeaks a large liberation of mathematics from the epistemic shackles of proof. Actually it doesn't. The short reason why is that knowledge acquired in that way runs into the tough headwinds of not being readily circulable. It is widely accepted that proof is the principal delivery-system of mathematical knowledge. Even so, to a quite striking extent, *knowledge* circulates on the sayso of those who possesses it.

> *More knowledge than proof:* On CR assumptions mathematical knowledge outpaces mathematical proof, and correspondingly shrinks mathematical ignorance.

The impression otherwise arises from the expectation that the KK-hypothesis holds of CR-knowledge. In fact the KK-hypothesis fails in the CR model.

It remains to say a word about the puzzling phenomenon of *hypothesis selection*. Line (2) of Peirce's schema is a subjunctive conditional. There is an indeterminacy of propositions which, if placed in antecedent position in (2), would preserve its truth when *A* was in it. Call these candidates "merely possible". How, then, does one fix on *A*? How does one cut down the merely possible to "ballpark" possibilities? As we saw, Peirce's answer is that abductive hypothesis-selection is a form of guessing, underwritten by instinct:

> "Now the surrender which we make in [abduction] is of surrender to the Insistence of an Idea. The hypothesis, as the Frenchman says, *c'est plus fort que moi*." (*CP* 5.181)

[45] See "Reorienting the logic of abduction" for details. Good information here is *A*'s excellent track-record under premissory licence. See also Urbański and Klawiter, "Abduction".

Paul Thagard rightly doubts that the conjectures of current subatomic physics—what with its ten or more spatial dimensions—is "innately constrained".[46] My answer to this is that the general *capacity* for right-guessing is hardwired, but right-guessing in theoretical science is selectively distributed under the developmental provisions of hard learning. So seen, what Planck did in arriving at the quantum hypothesis was business as usual elevated to business that wasn't at all usual. The simple fact is that some people who are smarter than the rest of us and, in varying degrees, better at doing the things we're all pretty good at. Mark well, most by far of our everyday successful reasoning is abductive, but it takes a massive cooperative effort, and sometimes a genius, to crack the most intractable of our ignorance-problems.[47]

8 Being told things

It is time to say something about how knowledge circulates in the cognitive economy and how it matters for ignorance. Most of what a human being will ever know, he will know by having been told it. By a large measure, most of where we cognitively go awry lies in being told things that aren't the case. Equally, most of what we know is brought to pass by being told things that are true. Notice, again, that the balance between well-informed and ill-informed tellings appears not to unsettle the knowledge-error balance, and does not sufficiently disoblige the Enough[4] thesis to sound a general alarm. When we know something by being told it, a filtration device is at work. Beings like us are good (not perfect) at sorting out information whose telling causes knowledge to be had and information whose telling doesn't.[48] By and large, people don't accept ethical claims just by having been told them. Filtration mechanisms appear to favour fact over value (if you'll excuse the clumsy and overworked distinction). It is, however, not the same device that filters good information from bad information. If you tell me that D is morally impermissible, and my belief-forming device doesn't fire, it doesn't remotely follow that what you've told me isn't true. As we saw just above, there are ranges of cases in which sayso is not an efficacious transmitter of mathematical truths known without proof by the would-be sayso circulator of them. Even so, it bears emphasizing that in even larger ranges of cases, when we know something by being told it, the mere telling of it *does* push the buttons of belief. If our belief-producing devices are in good working order and

[46] Thagard (1992; p. 170).

[47] It may interest some readers that none of the works cited in this section is mentioned in the *Stanford Encyclopedia of Philosophy*'s entries "Abduction" and "Peirce on Abduction" (Douven 2011, 2017). The gist of the latter is that Peirce's views on abduction have lost whatever appeal they may have had earlier on.

[48] There is another large literature which investigates the standards that qualify sayso as reliable and the manner in which recipients of sayso attempt to apply them in their own situations. There are contexts in which this is the right way in which to proceed; for example, when judges determine the admissibility of expert testimony at trial. In the general case, it is the wrong way to go. It is but a version of justificationism as applied to the reliability of testimony-induced belief. See here Lackey (2011, chapter 29) in Bernecker and Pritchard, and the works cited there. For a different slant, see *Errors of Reasoning,* chapter 9, and *Is Legal Reasoning Irrational?,* 2nd edition, chapters 8 and 15, and Woods (2018b; pp. 1205–1257).

working here in the way they're supposed to, our belief is causally produced, as is our knowledge should the belief be true. In comparison with justifyings and provings, tellings are ubiquitous. Indeed, in its more developed form, the CR-model provides as follows:

> *The most thesis*: Most, by far, of what we will ever know is got by being told it.
>
> *The causal thesis*: Most, by far, of what we know by having been told it is the causal result of its being told us.

For this to make sense, it can plausibly be said that

> *The truthfulness thesis*: Human beings are strongly disposed to be truthful in what they say. The economic advantages therefrom speak for themselves.[49]

Let's revisit Bill's anselmian knowledge that $E=mc^2$. Suppose he came to know it last Thursday in his science class at school. Bill on Thursday became a link in transmission chain of told-knowledge, arising, on or about the day in 1905 on which Einstein first announced it. The institution of telling links up in an interesting way with the phenomena of multiagency. Let's grant that Bill wouldn't on that occasion have known (in the way that he did) that $E=mc^2$ without having been told it by his teacher Mr. Fisher. Mr. Fisher couldn't have known it unless someone told it (in a textbook say) and the teller (or author) in turn without having been told it by someone who knew it beforehand, etcetera, etcetera, all the way back to the physicist who derived it from the principles of special relativity. The proximate cause of Bill's knowing it is Mr. Fisher. The distal cause of his knowing it is Einstein. Even in this much simplified example, Bill's knowing that $E=mc^2$ is the product of a sayso-manifold.

> *The sayso-manifold thesis*: With regard to an agent's knowing that p at t a sayso-manifold is the historical worldline aggregating all the tellings at $t' \leq t$ causally implicated in his knowing at t that p.

Manifolds, in turn, are the products of multiagents.

> *The multiagent thesis*: Sayso manifolds induce cognitive multiagancies made of the teller-told pairs of the tellings from which the manifold is constructed.[50]

It is easy to tell a plausible story about the sayso-manifold generation of everything anselmianly known at t. The story can be adapted to the telling-histories of things known semantically at t. The same, however, cannot be said for the causal

[49] See here Lewis (1975; pp. 3–35, 1983).

[50] The formal dynamics of telling haven't yet had much of an innings in the philosophical mainstream. Of most direct significance for epistemology is the theory of telling investigated by public announcement logics (PAL), originating with Paya's (1989; pp. 201–216). PAL extends multiagent epistemic logic to model the communicational consequences of announcements to multiagents. Think for example of the Prime Minister's televised speech announcing on behalf of the governments its intention to place tariffs on U. S. steel effective a week thence. Further references can be found in the note on p. 323 of *Errors of Reasoning,* chapter 9 on "Being told.".

antecedents of TI-knowledge at *t*. The reason is that the contents of both anselmian and semantic knowledge can be written down. The domain of TI-knowledge divides rather nicely between the aftermath of having once known things semantically that are now deposited in occurrent and deep memory, and the knowledge that hasn't yet surfaced, and may never will have occasion to.

There are three cases of multiagent knowledge-production that I'd now like briefly to touch upon. Their implications for the mix of knowledge and ignorance is three already-discussed grades of them is rather striking. The first is a famous case reported in *Physical Review Letters* under the title "Charm photoproduction cross section at 20 GeV", and ably discussed by John Hardwig.[51] The letter reports experimental results in particle physics, achieved by a team of ninety-nine co-authors of varying backgrounds and expertise—a complexly structured multiagent. Hardwig writes that

> "[o]bviously, no one person could have done this experiment – in fact … no one university or national laboratory could have done it – and many of an author like this will not even know how a given number in the article was arrived at." (p. 347, n. 10)

Hardwig adds,

> "Of course, only a few people actually wrote the article, but it does not follow that these people are masterminds for the whole procedure or that they completely understand the experiment and the analysis of the data."

The charm result was a nontrivial advance in particle physics. It was achieved by a multiagent, none of whose subcontractors understood it. Each subagent had some semantic command of certain essential components of the experiment's disclosure, but none of the whole of it. We have it from this that

> *The scientific indispensability of anselmian knowledge*: To a significant extent, anselmian knowledge is indispensable to team-science.

Further,

> *Significance for TI-knowledge*: Much of what the charm subagents had anselmian knowledge of is not a plausible candidate for their TI-knowledge of it.

We are now left to speculate about the condition that a multiagent is in when it produces knowledge like the knowledge produced by the charm experiment team. All things considered, I think it best to say that

> *What multiagents know*: Although prime candidates for the variable *situatedness* of it, the knowledge a multidisciplinary multiagent M *produces* is not a

[51] Hardwig (1985; pp. 335–349).

knowledge *possessed* by M, and not always or routinely by all of its subagents, and sometimes by none of them.[52]

Corollary: There is bound to be a significant lag between the production of knowledge and the semantic possession of it by those told of it thereafter.

This reinforces the suggestion that anselmian knowledge has a bedrock importance in multiagent science. Certainly it is not the sole preserve of kiddies and know-nothings.

The next two cases need not detain us long. The first is considerably more complex than the charm experiment case. The second is a case apart, and well on its way to becoming an epistemological phenomenon of puzzling proportions, with arresting implications for the cognitive enrichments of ignorance. The first of this pair is the Wiles proof of Fermat's Last Theorem. The Wiles proof, while nominally and importantly his, was the end-product of a 200 year chained multiagency dating from the Pythagorean triples of Babylonian times, third century Diophantine equations, the sum of squares problem dating from the tenth century, Fermat's own partial solutions of them, flowing to the work of Mordell and Kumner, and advanced in 1903 by Gerd Faltings. With the help of a SWAC computer, Faltings' results were extended and further advances were made by Wagstaff, Bubler, Crandell, Ennvald and Metsäkylä. The basic strategy for the proof was laid out in Gerhard modularity theorem in 1986, and improved by Katz. The basic strategy for proving the theorem outright was Gerhard Fey's modularity theorem of 1984, with later refinements by Serre and, in turn, Ribet's 1986 proof of the epsilon conjecture. Wiles proved a special case in 1995. It remained to prove to Taniyama-Shimura conjecture for semistable elliptic curves. With the aid of Horizontal Iwasawa theory and an Euler system advanced by Kolyvagin and Flack and Katz, Wiles announced his proof in 1993. The proof was flawed, With the help of Richard Taylor, he reworked Horizontal Iwasawa theory and got the desired result. Fermat's Last Theorem was now proved after 300 and 58 years of multiagent striving and supportive sayso.[53]

Much the same story can be told of the proving of Poincaré's Conjecture after 106 years of multiagent in-chaining.[54] The Wiles proof runs to 108 pages. Consider now anyone who kept himself abreast of things. Is there any need to expose the details of the interplay of the first three grades of epistemic involvement at work here? And of the unevenness of the distribution of its grades in the community of

[52] In his (1997), Peter Galison asks how it is possible for different "scientific cultures" with different disciplinary backgrounds to coordinate and produce a knowledge of particle accelerators, for example. More generally, how is multidisciplinary cognitive multiagency possible? The four insets above are part of the answer to Galison's question.

[53] For excellent historical and analytic coverage of the actual workings of what I've been calling mathematical multiagencies, see Ferreirós (2016).

[54] I've drawn here on Sect. 9.7 of *Errors of Reasoning,* pp. 318–320. In writing those pages I was told the Fermat part by Cornell et al. (1997), and Singh (1998). Now, for those of you who hadn't been told it before, I am happy to be the teller. It enrolls you and me in multiagent sayso-manifold ensuing from mathematical antiquity.

people who know it? Before moving to the second case announced lines ago, we'll turn briefly to the fourth grade of epistemic involvement.

9 The fourth grade of epistemic involvement: Impact-knowledge

Impact-knowledge is knowledge of a comprehensive kind. "Comprehensive" conveys three meanings that matter for knowledge. One means well-understood. Another means exhaustive. The meaning of the third can be found in the words of Kanamori concerning the theorem that Wiles and Co proved:

> "In a substantial sense, the statement is dated and unto itself has no intrinsic interest whatsoever, and it only grew in historical significance as it withstood more and more techniques, techniques that have enriched mathematics considerably … [Wiles] actually established the Simura-Taniyama Conjecture about elliptic curves in algebraic geometry through a beautiful synthetic proof, and this among mathematicians has been seen *as a great advance*." (p. 32: emphasis added)[55]

How so?

> "This truth [=Fermat's] among a myriad was definitely stalked, but in any case the proof considerably enriched the theory of elliptic curves, a theory that has its origins centuries ago in the study of planetary motion. There is still the question of [sayso] authority, but the understanding brought about by the proof is acknowledged in large part through its *potentiality, its opening up of new possibilities*." (Emphases added)

On Kanamori's telling, it is possible, indeed likely, that anyone having a semantic knowledge of Fermat's theorem and an anselmian knowledge of the fact that Wiles and Co proved it true, will utterly lack an impact-knowledge of it in our present sense of the world. This tells us something important about the state of prosperity in the mathematical economy.

> *Proofs as start-ups*: Oftentimes the *ultimate* epistemic value of a proof lies not only in knowing the truth of what it proves. It lies, rather, in the impetus it gives to mathematical start-ups for opening up new fields of study.

It is now past time to bring these reflections to a close. To that end, I'll come back to what I promised as something "well on its way to becoming an epistemological phenomenon of puzzling proportions with arresting implication for the cognitive

[55] Kanamori (2013; pp. 21–35). Angus MacIntyre has claimed that the central Modularity Thesis in Wiles' proof is provable in Peano arithmetic, prompting Kanamori to suggest that, if so, "there could be a new proof of the Fermat Last Theorem in Peano Arithmetic, possibly by passing through further analysis to Modularity Thesis". (p. 32). See also MacIntyre (2011; especially the appendix to chapter 1), "The impact of Gödel's incompleteness theorems on mathematics." But as Ferreirós notes, "… such convictions are not mathematical facts; we are facing *beliefs*", *Mathematical Knowledge,* p. 183.

enrichments of ignorance." Proofs as long as Wiles' are not the rarities we might think. In 2006, there were three published accounts detailing Grigori Perelman's proof of Poincaré's Conjecture. One ran to 200 pages, another to 326, and the third weighed in at 521 pages.

In this same boat are the four papers of Shinichi Mochizuki purporting to solve the *abc* conjecture. Consider expressions of the type $a+b=c$. Informally put, the conjecture is that if a goodly number of small primes divide a and b, only a few large ones divide c.[56] It is considered a rather elementary-seeming claim about addition and multiplication, and widely agreed that a proof of the *abc* conjecture would in one fell swoop solve an enormous number of unsolved Diophantine equations. Most number theorists agree that a proof of this conjecture could easily be the greatest mathematical advance this century. The four Mochizuki papers run to 500 pages and rest on a further 500 pages of backup material. As of 2015, "only four mathematicians say that they have been able to read the entire proof."[57] Now the number is only slightly larger.[58] To construct the proof, Mochizuki invented a new branch of number theory called "Inter-universal Teichmüller theory" (IUTT) which flummoxed the experts. "Looking at it, you feel a bit like you might be reading a paper from the future, or from outer space", blogged one expert. In 2014 Mochizuki was prompted to write that to understand his proof, it would be necessary "for researchers to deactivate the thought patterns that they have installed in their brains and taken for granted for so many years." (Castelvecchi, p. 179)[59]

Mochizuki is "attempting to reform mathematics from the ground up, starting from its foundations in the theory of sets …. (*ibid*, p. 181) As of the fall of 2014, only one mathematician is known to have claimed to have verified the proof. It is perhaps now (in early January, 2019) too early to see which way the disputed proof will go. On the side of neglect is the sheer difficulty critics and boosters have in making the proof intelligible. To which Mochizuki is quoted by Klarreich as suggesting that these difficulties arise from a "lack of sufficient time to reflect deeply on the mathematics under discussion" in combination with "a deep sense of discomfort, or unfamiliarity, with new ways of thinking about familiar mathematical objects." Two fates await it. One is that it will die of neglect. The other is that it will catch hold enough to enable its longer-time survival as an object of anselmian knowledge in the upper reaches of number theory.

Not only is stress of proofs on hardcopy-production intensifying and the burden of bringing them to print more and more widely distributed in the large multiagencies

[56] Let $q(abc)$ be the quality of the triple (abc), defined as $\log c/\log(\text{rad}(abc))$. The radical rad (n) of a positive integer n is the product of n's distinct prime factors. The *abc* conjecture puts it that, for every real number ε, there exist only finitely many triples (abc) of coprime positive integers with $a+b=c$ such that $q(abc)>1+\varepsilon$.

[57] Castelvecchi (2015).

[58] Klarreich (2018).

[59] Graham Priest gave much the same advice to metaphysicians who have difficulty bending their minds to objects of which it is the case that nothing whatever is they. Dissenters, he proposed, should "get their understanding rewired." Priest (2016).

that support them, there is now a growing reliance on computational methods of proof-making by self-learning mechanisms. Here is Kanamori on this point:

> "The weight of these various examples shows how far away mathematics now is from being comprehended by any formal notion of proof and any theory of mathematical knowledge, and how *the limits of human intelligibility are being put to the test*." (*idem*; emphasis mine)

And perhaps, I would add, in process of being *surpassed.* It is always interesting to speculate about what an alien civilization might think of us and conceivably do about it. Closer to home are the phenomena of Big Data, Deep Learning, and self-teaching theorem-provers. Computer-assisted proofs now produce theorems whose proofs, it is said, cannot be fathomed by even the best of our ilk. For two millenia and more, proof ruled the roost in mathematics and man has ruled the proofs.[60] As we have it now, proofs still call the shots but, in ranges of cases, they appear not to be ruled by us. They are ruled by the computers we've built to teach themselves how to prove things beyond human ken of how it's done. This bespeaks an interesting epistemic future for frontier mathematics:

> *anselmian knowledge regnant*: To the extent these developments hold true, the more deeply frontier mathematics lies in the repose of anselmian knowledge.
>
> *Remark*: London's Royal Society has a famous motto: *Nullias in verba*, or "take no one's word for it." It is wise and necessary advice in its place. Over-applied, it's a choke-hold on mathematics and renders it epistemically barren.

In wrapping up this section, it is necessary to distinguish unintelligibility in fact from unintelligibility in principle. For most mathematicians, the Mochizuki proof won't be intelligible to them until they achieve a working acquaintance with IUTT. That might not in fact happen but could in principle. Computer-generated proofs provide a different kind of case, when their computational complexity is too much, both in fact and in principle, for the human knower to take in. Each case presents intriguing epistemological challenges, but they are challenges of a different order.

10 Concluding remarks

The four grades of epistemic involvement, each in turn subject to gradations, scatter the seeds of ignorance across the cognitive economy in ways both productive and inhibiting, leaving overall a balanced ledger of costs and benefits—the foundation of the fact that we are right enough enough of the time about enough of the right enough things to survive and prosper and have built the great cathedrals of mediaeval France. This helps us see that the cognitive economy is itself a causally thrumming information-latticed multiagency dynamically animated by the drive to know—by cognitive yearning—that burns in the breasts of its human subagents. It

[60] For an informal overview, see Weinberger https://www.wired.com/story/neverunderstand/.

is a jujitsu kind of economy, well-versed in flipping cognitive threat to compensating advantage. Its knowledge-ignorance playbook is causally loaded and, to a high degree, elusive of semantic grasp. Accordingly, it is itself the site of fruitful ignorance, subject to intermittent and always incomplete semantic redress. The account presented here is a sketch of a theory, much of which latter is a work in progress, and some of whose central tenets have been issued here as "working hypotheses". One such is that a causal-response epistemology is best attuned to the lived realities of our individual and collective cognitive economics. That, for me, is a theoretical given whose chief value, unlike justifictionist theories, is that it makes knowledge an affordable commodity in cognitive economics. It has not been the burden of the present essay to make a detailed case for the causal-response model of knowledge. In the space that's been afforded me, I have tried to make the beginnings of a case for a further refinement of the basic model, one which exposes more of the structure of cognitive economies. In my submission, a CR-epistemology will prosper only if it recognizes in a principled way the four-tieredness of human knowledge and, correspondingly, the four-tieredness of human ignorance. Crucial to the operational economics of these tiers is some detailed grasp of the cognitive prosperity that ensues from the causally sourced inter-tier criss-crossings and intra-tier combinations of the known and not-known. The principal assertion of these pages is that the epistemology that's best attuned to the realities of cognitive life is not only a causal-response one but also one that recognizes the cognitive value of interactions between and within four-tiered knowing and four-tiered not-knowings.

Much of the paper rests upon conjectures that have attempted to lend some coherence to the facts as they currently present themselves. Like all abductively derived hypotheses, it is desirable that they be put to the test. Many of the paper's working hypotheses have been based on facts disclosed in various of the sciences, which is where confirmation or otherwise would be most likely to be found. One thing is clear: In the manner in which it is presently structured, confirmation of a four-tier causal-response epistemology (FTCR) is beyond the reach of philosophy proper. In naturalized settings, it is widely agreed that getting knowledge right requires solid but circumspect working partnerships with all the sciences of cognition. I say "circumspect" with reason. One is that in several of its branches, cognitive science has uncritically adopted philosophical assumptions which, if true, would make human knowledge exceedingly unaffordable. It would do so as a consequence of shrinking knowledge to a shell of its former self. Leading the list of price-fixers are scientists who restrict the reach of nonmonotonic knowings to agents who comport with the rules of the probability calculus. Such rules are epistemically suffocating and purveyors of big-box scepticism on a scale that makes one's teeth ache. They should be avoided at all cost.[61]

[61] Again because space is stretched, I'll briefly come back to what I take to be a fatally mistaken epistemological course. It is the course of postulating the ideal rational agent and defining it as a being which revises its belief in strict accordance with the rules of the probability calculus, which its endorsers invest with a normative authority that binds the belief-revisions of human life. One such rule is that a rational believer will close his beliefs under consequence. There being at least ω of them, there is no finite degree to which any human reasoner approximates to that standard. Another example is from Harman (1986). Suppose that you are head of a CID investigation of a serious crime. As of now, the evidence supports the judgement that Spike McGurk is a shade away from being charged with the crime. It is now early

Acknowledgements For sagacious advice and constructive encouragement I warmly thank the two editors of this special number, Lorenzo Magnani and Selene Arfini. I also thank my UBC colleague Alirio Rosales for many generous hours of fruitful instruction and productive criticism. Thanks, too, to Mike Bennett of the UBC mathematics department for welcome instruction on how the Mochizuki proof is structured. For their constructive suggestions, I also thank this journal's two referees.

References

Adriaans, P., & van Benthem, J. (Eds.). (2008). *Philosophy of information,* a volume of the *Handbook of the philosophy of science,* edited by Dov Gabbay, Paul Thagard and John Woods. Amsterdam: North-Holland.

Arfini, S. (2019). *Ignorant cognition: A philosophical investigation of the cognitive features of not-knowing*. Cham: Springer.

Arfini, S., Bertolotti, T., & Magnani, L. (2018). The diffusion of ignorance in online communities. *The International Journal of Technoethics, 9,* 37–50.

Arfini, S., & Magnani, L. (2016). Cognitive autoimmunity: Knowledge, ignorance and self-deception. *Logic Journal of IGPL, 24,* 612–627.

Aristotle. (1984). *The complete works of aristotle.* In J. Barnes (Ed.), *The Revised English Translation, two volumes*. Princeton: Princeton University Press.

Armour-Garb, B., & Woodbridge, J. A. (2015). *Pretense and pathology: Philosophical fictionalism and its applications*. Cambridge: Cambridge University Press.

Aru, J., & Bachman, T. (2017). In and out of consciousness: How does conscious processing (d)evolve over time? *Frontiers in Psychology, 8,* 128–135.

Bachmann, T. (2000). *Microgenetic approach to the conscious mind*. Amsterdam: John Benjamins.

Bernecker, S., & Pritchard, D. (2013). *The Routledge companion to epistemology*. London: Routledge.

Blackburn, S. (2005). *Oxford dictionary of philosophy* (2nd ed.). Oxford: Oxford University Press.

Breitmayer, B. G. (2014). *The visual(un)conscious and its discontents*. New York: Oxford University Press.

Callen, H. B. (1985). *Thermodynamics and an introduduction to thermostatistics* (2nd ed.). New York: Wiley.

Carruthers, P. (2011). *The opacity of mind: An integrative theory of self-knowledge*. Oxford: Oxford University Press; paperback in 2013.

Castelvecchi, D. (2015). The biggest mystery in mathematics: Sinichi Mochizuki and the impenetrable proof. *Nature News*, *526*, 178–181.

Chiffi, D., & Pietarinen, A.-V. (2018). Abductive inference within a pragmatic framework. *Synthese, 10,* 15. https://doi.org/10.1007/s11229-018-1824-6.

Cornell, G., Silverman, J. H., & Stevens, G. (Eds.). (1997). *Modular forms and Fermat's Last Theorem*. New York: Springer.

Douven, I. (2011, 2017). "Abduction" and "Peirce on abduction". In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy*. Substantive revision Fri April 28, 2017.

Doya, K., Iski, S., Pouget, A., & Rao, R. P. N. (Eds.). (2011). *Bayesian brain: Probabilistic approaches to neural coding*. Cambridge, MA: MIT Press.

Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.

Einstein, A. (1905). Ist die Träged eines Körpers von Seinem Energieinhalt ahbankängig? *Annalen der Physik, 18,* 639–643.

Footnote 61 (continued)
tomorrow morning. Twenty items of new evidence have been logged overnight. For that to have happened in accordance with the Bayesian belief-update rules, ≈ one million calculations would have to have been made. Had 30 bits of new evidence arrived overnight, ≈ a billion calculations would have been required for update. The very idea that these rules doom you and your team to irrationality is as amusing as it is bemusing. In matters of cognition, as in all things, it is always a good idea to watch out with whom we plan to do business.

Ferreirós, J. (2016). *Mathematical knowledge and the interplay of practices*. Princeton: Princeton University Press.
Gabbay, D. M., & Woods, J. (2003). *Agenda relevance: A study in formal pragmatics,* volume 1 of their *A practical knowledge of cognitive systems.* Amsterdam: North-Holland.
Gabbay, D. M., & Woods, J. (2005a). *The reach of abduction: Insight and trial,* volume 2 of their *A practical logic of cognitive systems.* Amsterdam: North-Holland.
Gabbay, D. M., & Woods, J. (2005b). Filtration structures and the cut-down problem for abduction". In K. A. Peacock & A. D. Irvine (Eds.), *Mistakes of reason* (pp. 398–417). Toronto: University of Toronto Press.
Galison, P. (1997). *Image and logic: A material culture for microphysics*. Chicago: University of Chicago Press.
Gascoigne, N., & Thorton, T. (2013). *Tacit knowledge*. Durham: Acumen.
Godfrey-Smith, P. (1996). *Complexity and function of mind in nature*. Cambridge: Cambridge University Press.
Goldman, A. I. (1967). A causal theory of knowing. *Journal of Philosophy, 64,* 357–372.
Goldman, A. I., & Whitcomb, D. (Eds.). (2001). *Social epistemology: Essential essays*. New York: Oxford University Press.
Hardwig, J. (1985). Epistemic dependence. *Journal of Philosophy, 82,* 335–349.
Harman, G. (1986). *Change in view: Principles of reasoning*. Cambridge, MA: MIT Press.
Heidegger, M. (1962). *Being and time*. Oxford: Blackwell.
Hintikka, J. (2007). *Socratic epistemology: Explorations of knowledge-seeking by questioning*. Cambridge: Cambridge University Press.
Hohway, J. (2013). *The predictive mind*. Oxford: Oxford University Press.
Honderich, T. (Ed.). (2005). *The Oxford companion to philosophy, new edition*. New York: Oxford University Press.
Jacquette, D. (2018). Cracking the hard problem of consciousness. In D. Jacquette (Ed.), *The Bloomsbury companion to the philosophy of consciousness* (pp. 258–286). London: Bloomsbury.
Kanamori, A. (2013). Mathematical knowledge: Motley and complexity of proof. *Annals of the Japan Association for Philosophy of Science, 21,* 21–35.
Kitcher, P. (1983). *Abusing science, the case against creationism*. Cambridge, MA: MIT Press.
Klarreich, E. (2018). Titans of mathematics clash over epic proof of abc conjecture. *Quanta Magazine*. https://www.quantamagazine.org/titans-of-mathematics-clash-over-epic-proof-of-abc-conjecture-20180920/#.
Kunen, K. (2013). *Set theory,* volume 34 of *Studies in logic*. London: College Publications.
Lackey, J. (2011). Chapter 29: Testimonial knowledge. In S. Bernecker & D. H. Pritchard (Eds.), *The Routledge companion to epistemology*. London: Routledge.
Lewis, D. (1975). Languages and language. In K. Gunderson (Ed.), *Minnesota studies in the philosophy of science VII* (pp. 3–35). Minneapolis: University of Minnesota Press; reprinted in Lewis.
Lewis, D. (1983). *Philosophical papers* (Vol. 1). New York: Oxford University Press.
Lowney, C. W., II (Ed.). (2017). *Charles Taylor, Michael Polanyi and the critique of modernity: Pluralist and emergentist directions*. London: Palgrave Macmillan.
MacIntyre, A. (2011). *Kurt Gödel and the foundations of mathematics: Horizons of truth*. Cambridge: Cambridge University Press.
Magnani, L. (2009). *Abductive cognition: The epistemological and eco-cognitive dimensions of hypothetical reasoning*. Berlin: Springer.
Magnani, L. (2015). Naturalizing logic and errors of reasoning vindicated: Logic reapproaches cognitive science. *Journal of Applied Logic, 13,* 13–36.
Magnani, L. (2018). The urgent need of a naturalized logic. In G. Dodig-Crmkovic & M. J. Schroeder (Eds.), *Contemporary Natural Philosophy and Philosophies,* a special guest-edited number of *Philosophies* (Vol. 3, p. 44).
Minnameier, G. (forthcoming). The economics of morality and the fabric of social sciences. *Economic Perspectives.*
Mole, C. (to appear). The creativity of generative adversarial networks. *Philosophy of Science*.
Mole, C. (2008). Attention and consciousness. *Consciousness Studies, 15,* 86–104.
Moss, S. (2018). *Probabilistic knowledge*. New York: Oxford University Press.
Orzumi, M., Albantakis, C., & Tononi, G. (2018). From phenomenology to the mechanisms of consciousness: Integrated information theory 3.0. *PLoS Computational Biology, 10,* 1–25.

Paglieri, F. (Ed.). (2012). *The role of the natural and social context in shaping consciousness,* volume 86 of *Advances in Consciousness Research.* Amsterdam: John Benjamins.
Park, W. (2017). *Abduction in context: The conjectural dynamics of scientific reasoning*. Cham: Springer.
Paya, J. (1989). Logics of public communications. In M. L. Emrich, M. Hadzikadic & Z. W. Ros (Eds.), *Proceedings of the fourth international symposium on methodologies for intelligent systems* (pp. 201–216). Amsterdam: North-Holland.
Peirce, C. S. (1931–1958). *Collected works, in eight volumes.* Cambridge, MA: Harvard University Press.
Peirce, C. S. (1992). Reasoning and the logic of things. In K. L. Kettner (Ed.), *The Cambridge conference lectures of 1898 (RLT)*. Cambridge, MA: Harvard University Press.
Pelletier, F. J. (Ed.). (2010). *Kinds, things and stuff*. New York: Oxford University Press.
Polanyi, M. (1958). *Personal knowledge*. Chicago: University of Chicago Press.
Polanyi, M. (2009). *The Tacit Dimension, with a Foreword by Armatya Sen*. Chicago: University of Chicago Press.
Priest, G. (2016). *Towards non-being* (2nd ed.). Oxford: Oxford University Press; first edition in 2005.
Rescher, N. (2007). *Error: Our predicament when things go wrong*. Pittsburgh: University of Pittsburgh Press.
Rescher, N. (2009). *Ignorance: On the wider implications of deficit knowledge*. Pittsburgh: University of Pittsburgh Press.
Ryle, G. (1929–1930). Are there propositions?. *Proceedings of the Aristotelian Society, 30*, 91–126.
Ryle, G. (1945–1946). Knowing how and knowing that. *Proceedings of the Aristotelian Society, 46*, 1–16.
Ryle, G. (1949). *The concept of mind*. London: Hutchinson.
Shiffrin, R. M. (1976). Capacity limitations in information processing, attention, and memory. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Attention and memory* (Vol. 4, pp. 177–236). Hillsdale, NJ: Erlbaum.
Shiffrin, R. M. (1997). Automatism and consciousness. In J. D. Cohen & J. W. Schooler (Eds.), *Scientific approaches to consciousness* (pp. 49–64). Mahwah, NJ: Erlbaum.
Singh, S. (1998). *Fermat's Enigma*. New York: Anchor.
Stich, S. P. (1990). *The fragmentation of reason*. Cambridge, MA: MIT Press.
Thagard, P. (1992). *Conceptual revolutions*. Princeton: Princeton University Press.
Unger, P. (1975). *Ignorance: A case for scepticism.* Oxford: Clarendon Press, paperback in 2002.
Urbański, M., & Klawiter, A. (2018). Abduction: Some conceptual issues. *Logic and Logical Philosophy, 27,* 583–597.
Wegner, D. (2001). *The illusion of conscious will*. Cambridge, MA: MIT Press.
Weinberger, D. (2017). Our machines now have knowledge we'll never understand, Bachchannel, April 18, 2017. https://www.wired.com/story/neverunderstand/.
Wiles, A. (1995). Modular elliptic curves and Fermat's Last Theorem. *Annals of Mathematics, 141,* 443–551.
Wilson, T. (2002). *Strangers to ourselves*. Cambridge, MA: Harvard University Press.
Wittgenstein, L. (1953). *Philosophical investigations*. Oxford: Blackwell.
Woods, J. (2012). Cognitive economics and the logic of abduction. *Review of Symbolic Logic, 5,* 148–161.
Woods, J. (2013). *Errors of reasoning: Naturalizing the logic of inference* volume 45 of *Studies in logic*. London: College Publications; reprinted with corrections in 2014.
Woods, J. (2017). Reorienting the logic of abduction. In L. Magnani & T. Bertolotti (Eds.), *Springer handbook of model-based science* (pp. 137–149). Heidelberg: Springer.
Woods, J. (2018a). *Is legal reasoning irrational?* volume 2 of *Law and society* (2nd ed.), revised and extended. London: College Publications; first edition in 2015.
Woods, J. (2018b). What strategicians might learn from the common law: Implicit and tacit understanding of the unwritten. *IfCoLoG Journal of Logics and Their Applications, 5,* 1205–1257.
Zimmerman, M. (1989). The nervous system and the context of information theory. In R. F. Schmidt & G. Thurs (Eds.), *Human physiology* (2nd ed., pp. 166–175). Berlin: Springer.
Zimmerman, A. Z. (2018). *Belief: A pragmatic picture*. New York: Oxford University Press.

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