philosophies

MDPI

Article

Could a Computer Learn to Be an Appeals Court Judge? The Place of the Unspeakable and Unwriteable in All-Purpose Intelligent Systems

John Woods

The Abductive Systems Group, Department of Philosophy, University of British Columbia, Vancouver, BC V6T 1Z1, Canada; john.woods@ubc.ca

Abstract: I will take it that general intelligence is intelligence of the kind that a typical human being—Fred, say—manifests in his role as a cognitive agent, that is, as an acquirer, receiver and circulator of knowledge in his cognitive economy. Framed in these terms, the word "general" underserves our ends. Hereafter our questions will bear upon the *all-purpose* intelligence of beings like Fred. Frederika appears as Fred's AI-counterpart, not as a fully programmed and engineered being, but as a presently unrealized theoretical construct. Our basic question is whether it is in principle possible to equip Frederika to do what Fred does as an all-purpose participant in his own cognitive economy. Can she achieve a sufficiency of relevant similarity to him to allow us to say that she herself can do what Fred can do, perhaps even better? One of the things that Fred can do—or at least could learn from experience to do—is discharge the duties of an Appeals Court judge. As set down in the ancient doctrine of *lex non scripta,* Fred must be able to detect, understand and correctly apply certain tacit and implicit rules of law which defy express propositional formulation and linguistic articulation. Fred has an even more widespread capacity for the epistemically tacit and implicit, clearly one of his most cost-saving kinds of intelligence. Indeed, most by far of what Fred will ever know he will know tacitly and implicitly. So we must ask: how tightly bound to the *peculiarities* of Fred's cognitive enablement conditions is the *character* of the intelligence that he manifests? And how far down Fred's causal make-up does intelligence actually go?

Keywords: cognitive economics; energy-to-energy transduction; energy-to-information conversion; filters; implicity; jiu-jitsu advantage; ignorance; inconsistency; phase transitions; relevant similarity; storage and retrieval; tacity; the cognitive down-below; told knowledge; unwritten law; young truth

Citation: Woods, J. Could a Computer Learn to Be an Appeals Court Judge? The Place of the Unspeakable and Unwriteable in All-Purpose Intelligent Systems. *Philosophies* **2022**, *7*, 95. https://doi.org/10.3390/philosophies7050095

Academic Editor: Woosuk Park

Received: 5 May 2022
Accepted: 18 July 2022
Published: 26 August 2022

> "Formal logic must not be too formal. It must represent a fact of experience, or else it is in danger of degenerating into a mathematical recreation."
>
> C. S. Peirce [1] (2. 710)

> "You should never express yourself more clearly than you are able to think."
>
> Niels Bohr to Abraham Pais, viva voce.

1. Fred

1.1. Beings like Fred

This paper [1,2] is divided into three unequal parts. Section 1, the longest, is about Fred. Much shorter is Section 2 about Madam Justice Flanagan, made so by the fact that most of what is to be said about her cognitive make-up will already have been said about Fred's. Frederika is the focus of Section 3, the shortest of the three sections, made so by unsettled uncertainties arising from peculiarities unearthed in the preceding parts, which still await further empirical elucidation.

As the next wave of the artificial is about to make itself known to us, an obvious question will be in what its newness consists. Is it new enough and different enough to justify whatever ordinal number might lie in its name? A serious candidate for this title

is what is sometimes called general intelligence. The idea here is that computers can do some of the intelligent things that we ourselves do, and do them more quickly and more accurately. So, then, it lies in computation's future that it can do every kind of intelligent thing that we humans do, and if so, at what cost for what benefit? "General" here carries the sense of "all-purpose" or "comprehensive". In this essay, Fred plays the role of a typical human being [2–5].[3] When we ask the present question with him in mind, we are asking about all the ways in which his behaviour is a manifestation of intelligence. If we found that we could grace the computational question with an affirmative answer, the chances would enlarge that all kinds of intelligent human behaviour, even at their humanly most perfect, could be discharged by a machine with good cost-benefit outcomes. But it is Fred, I think, rather than the most perfectly idealized human, who is the real nut for AI's Deep Learning to crack. Powerful computational engineering has been an established fact of life for just over sixty-years or so, and has transformed human tasking more than any other technological advance. Anyone with a grown-up memory of 1960 has first-hand grasp of these developments—sixty straight years of creative destruction. Although setbacks and vexations were par for the course there, the larger truth is one of one impressive success after another. I have little doubt that *n*th wave AI, whether we like it or not, will be marked with like successes and, however things go with her, Frederika will be a machine to reckon with. The pun was inadvertent. What I mean is that Frederika will be no pushover. Still, the question of what it would take to build her to do at least as well as Fred presses for some sharp notice of what it has taken to make it the case that Fred himself is as good at intelligence-manifestation as we find him to be. And it also invites an answer to "How good is that?" Whether the big push can be accomplished is but one of the questions to raise about the present striving for third-wave advancement of AI. Of at least equal importance is whether it should, if possible, *be* accomplished, and by what means and under whose authority? How would the new regime be regulated, and how far would the redundancy of intelligent *human* agency be planned for or allowed to go? And to what extent should we place the true knowledge of yet unknown things beyond all human understanding once that knowledge has been computationally attained? Each of these questions is complex, not in the way that Aristotle rightly complained of, but rather in the sense of motivating further questions in search of well filled-out answers. The whether-possible question is riddled with questions and unclarities of its own, all circling the central notion of sufficient relevant similarity. If we desire Frederika to manifest intelligence in all the ways in which Fred manifests his own, how like his cognitive make-up must her own be engineered to be? It is not a question for which mathematical logicians have much mind or time, and, surprising as it may seem, it is also a question that evades the insights of establishment epistemology. So if one is a logician by profession who also appreciates the undetachable importance of epistemology for logical inference and proof, where is one to turn for direction? The present essay proposes a way of proceeding. To know what Fred is like in his cognitive doings, we should begin by following the

> *What Actually Happens Rule*: In the first instance, attend with care to the observable regularities in Fred's cognitive behaviour. In the second, make your best fist of their abductive accountancy, that is, of bringing them into some overall coherence. In the third instance, do not let theory override reality without just cause.
>
> *Corollary*: And keep in mind that Fred's role here is that of Everyman.

It is a matter of some importance that what the rule bids us to do, digital computers do much the same thing. If they hold course, Fred and Frederika alike will process large swarms of empirical data for the patterns that can be found within. If we thought of computer pattern-searches as digital empiricism, that would not be far off the epistemic character of Fred's pattern-searches.

As we proceed, I will take particular note of human cognitive capacities which strike me as especially noteworthy—indeed as rather amazing—and often overlooked or scanted

by establishment philosophy. These will serve as a checklist for what we might also ask of Frederika.

The present essay is itself a work in abductive progress. In it, more questions are posed than answered, partly because of the lack of space and partly because of the lack of well worked-out answers. The first thing to know about abduction is that it is a *limitation-mitigation* device. In other words, if Fred is to know these things, he must have the very limitations which make mitigation possible. I see abduction as a form of reasoning spurred by ignorance-problems. There is something the enquirer does not at present know and for which he has set himself out to find the answer. It lies in the nature of problems of this kind that they cannot be solved with the enquirer's present knowledge, even when supplemented by such epistemic updates as lie within his timely reach, as for example, by checking in at the library or clicking on a search-engine or by asking someone in the know. What is more, even though the problem cannot be solved with present updated resources, the enquiring mind is not prepared to let the matter go. So there is no recourse but to do some serious thinking.[4] C. S. Peirce calls this thinking "originary". Its object in the case presently before us is to implement the What Actually Happens Rule. Again, the rule bids us to examine with care the empirically discernible behavioural regularities in Fred's cognitive practices, and then to formulate hypotheses which, if true, would hold those practices to good theoretical account.

Peirce places great emphasis on originary thinking, one of whose most common modalities is what he calls good-guessing, that is, the openness to be moved by hypotheses that come to one *by la force majeure* of an insistent idea [1] (2. 443) [6] (873 13–15, 638 14–15) [1] (5. 181). In some cases, candidate hypotheses are already to hand,[5] but of even greater importance are the hypotheses that embody wholly new ideas or, as Peirce says, the "positively creative" "*entia ratione*" which "are quite as real [as antecedently existing things]" [1] (773 2–3). This is an extraordinary insight if true, and which one would think not obviously true. But before dismissing it from further consideration, consider the smashing breakthrough of Riemann's invention of the n-dimensional manifold, without which there might not have been differential geometry and, without that, the loss of the tensor calculus on Riemannian manifolds which opened the way to general relativity theory. Lest we make too much of these extraordinary powers, we should note that the human cognitive economy is not chockablock with geniuses. Even so, while the man in the street [6] is not so gifted, everyone who is so is in all essentials just like Fred. To refine what we are asking of Frederika, we ask whether she can match or better *all* manifestation of human intelligence, never mind the net spread of their manifestation in the human population at large. And to answer that question, we must give due notice of abduction's prominence in the inferential practices of humanity.

1.2. Abduction

Peirce is the modern founder of abductive logic and the author of some of its richest insights. To help keep us on track in this essay, I shall from time to time pivot to him as a point of contrast. There are, however, two points on which there is scholarly confusion that can be cleared up now. Peirce says repeatedly that when an abductive hypothesis is successfully arrived at, it provides an explanation of the phenomenon that triggered the search for it. [7] Some scholars rightly point out that the success-conditions for abductive reasoning do not always give a result that squares with the actual meaning of the word "explanation" [7]. [8] There is point to this, since for Peirce, what it means for a hypothesis to explain the phenomenon in question is to render it necessarily *deducible* from it. [9] Since this does not catch other meanings of "explain", it is best to stick with a meaning equally acceptable to Peirce and more at ease with natural speech: Hypotheses must *account for* the data in question. Note, however, the strictness of what Peirce demands of data accountability.

It is important to see what Peirce's deductivist abduction is not. It is not his view that when a successful abduction is attained, of a fact which, if true, would explain the fact that awaits on it (the target fact), the truth of the *explanans* would follow of deductive necessity.

His claim rather is that the *abduction* performs its explanatory function only if the target-fact follows *deductive necessity* from its *explanans*. Judging from Fred's own abductive doings, the bar of deductivist abduction requires such heavy lifting that it can only be said that it transgresses What Actually Happens and leaves us with two questions to ask of Frederika. Given that abduction is a limitation-mitigation device, could she be engineered to be so bound without imperilling the relevant similarity requirement?

A second matter which is often overlooked is the sharp line Peirce imposes between being led to think-that and one's drawing-that as an inference. The distinguishing mark is conscious control. [10] In the Abstract, we asked, "How far down Fred's causal makeup does intelligence actually go?" If we replaced "intelligence" with "inference", Peirce's answer would be that it stops at the point at which consciousness subsides and the question of conscious control no longer arises. On a fair reading, the thinking that underlies hypothesis-selection, certainly when the selector yields to the insistence of an idea, does not qualify as reasoning, never mind the large cognitive advantage of our receptivity to *force majeure.* [11] But also on the chopping block is any notion of tacit and implicit knowledge and inference, and also in question is Justice Flanagan's and her colleagues' wherewithal for judge-made law in courts of appeal. As we expand our tour of what Fred is like as a cognitive being, we should be on the lookout for reasons to uphold Peirce's aversion to uncontrolled reasoning. I will say for now that there will be scant reason to yield to it [8]. [12] Peirce himself is unable to give it consistent obligence. Indeed, he sees guessing as what happens when the abducer's enabling instinct is triggered. There is little doubt that Peirce saw the guessing instinct as naturally selected for, but not, of course, for guesses one actually makes [1] (5. 171, 7. 220). There is no doubt at all that by guessing, Peirce does not meet launching shots in the dark.

> "Proposals for hypotheses inundate us in an overwhelming flood while the process of verification to which each one must be subject before it can count as at all an item even of likely knowledge, is so very costly in time, energy, and money, the Economy here would override every other consideration even if there were any other serious considerations. In fact there are *no* others." [1] (5. 602; emphasis mine.)

When an originary thinker creates a new idea and embodies it in a hypothesis, something real is created. For, if successfully abduced, the hypothesis "is gravid with young truth" [6] (683) 1. In saying so, Peirce is onto something important. Just as concepts need time and circumstance to mature (e.g., the concept of set), so "young truth needs" time and circumstance to grow up into the real thing. So we can think of young truth as truth *in potential.* And it is to Peirce that we owe the idea that young truth realizes its potential under the spur of market-conditions in free markets of competing ideas. I will come back to this below [9]. [13]

Peirce sometimes likens hypothesis-selection to a lottery. With a little artful dodging, we could make the hypothesis-space for any given abduction problem as big as we like. Then on a statistical analysis of chance probabilities, the chances of tenable selection very quickly outrun what is possible for Fred. Although Peirce mentions these difficulties, it is clear that that he is not defeated by them. For the insistence of the idea one has chosen overrides the statistical chances of selection.

It is also important to keep in mind that Peirce sees mathematical progress as paradigmatically abductive. He writes that "[o]n the other hand, it is an error to make mathematics consist exclusively in the tracing out of necessary consequences. For the framing of the hypotheses of the two-way spread of imaginary quantity, and the hypotheses of Riemann surfaces were certainly mathematical achievements" [10] (Art. 3). Concerning the manner of such achievements, Peirce goes on to say, "It cannot be said that all framing of hypotheses is mathematics. For that would not distinguish between the mathematician and the poet . . . . Detective stories and the like [also] have an unmistakable mathematical element" [10] (Art. 2). It should be noted that by mathematics, Peirce means a certain kind of practice. It is the framing of hypotheses and the tracing of their deductive consequence. In short, it is abduction. It is a remarkable concession and a deeply important insight into how truth is

acquired in mathematics, where by truth I mean real truth, not artful fictions masquerading as truths, and truths that comport with realist assumptions overall. [14]

To modern ears, the intellectualist view that reasoning is always consciously controlled seems rather old hat. It would be of some possible interest here that Peirce's thinking about control stems from his pioneering work on logic machines. Peirce accepts that numerical computation is reasoning-like in the sense that its results are entailed by its data. But if we called this reasoning, "[a]ny apparatus whatever used for experimentation would be, on the same principle, a logical machine". [15] Peirce grants that, starting from true premisses, a computer can work its way to true conclusions. But here, too, he declines to call this reasoning. It lacks the capacity for conscious approval, and it lacks the capacity for "originary thinking" [10] (pp. 168–169). A further reservation arises from Peirce's conviction, true at the time but not when Arthur Burks wrote of it in the 1940s, that whereas a human's memory-capacity is unlimited, the same could not be said for machines [11–14]. [16] If we heed the What Actually Happens Rule, we see that it is, in fact, rather more the other way around. Before leaving Peirce, there would be value in knowing that the independent co-founder of the logic of second-order quantification, and the independent modern co-originator of ionic and diagrammatic proof-methods, was also a pioneer of logic gates, lattice theory, NAND, Shannon-Weaver information theory and modern pragmatic theories of information, whose later importance is reflected in several ways, as for example, in the Peirce-Birkhoff conjecture and the Peirce-Birkhoff ring [15]. [17] It is rather awkward that the inventor of modern abductive logic should have found himself so athwart the empirically discernible behavioural regularities of Fred's cognitive practices, both in his deductivist reading of abductive success and his intellectualist approach to human reasoning. But if we want to do right by Fred, we cannot hold him to such views.

If we are to arrive at a well-made understanding of the limit of intelligence, it makes plain methodological sense to examine the hard cases that arise for people like Fred. One of the hardest—I mean hardest to get to the bottom of—is how justice is meted out by Courts of Appeal. We should pause awhile with this.

1.3. Some Legal Peculiarities

Humans are enormously varied in their cognitive wherewithal. Again, the likes of Archimedes, Aristotle, Dante, Michelangelo, Newton, Shakespeare, Bach, Riemann, Einstein and Turing are a comparatively thin and scattered minority. But there are things that lie in Fred's cognitive purview that are common to us all. One of these potentials is realized in the course of learning how to be an Appeals Court judge at the common law bar of justice. Unlike lawyers, who must achieve formal qualifications as a condition of legally permitted practice, there is no School of Judges from which aspirants obtain their formal qualifications for membership of the bench. In many common law jurisdictions—England, Canada and Australia, to name just three—judgeships at all levels are by appointment, not election. In any properly run jurisdictions, all judges are drawn to the bench after many years of practice at the bar. This is the example I will stay with here, without prejudice to the question of which is the better judicial selection system. For the fact remains that in either case, no judge will deserve his title or his office if he lacks judgement, that is, the capacity and habit of arriving at wise findings. And no high court judge [18] could be similarly positioned if he were unable to judge wisely even in matters that *preclude* written and spoken articulation, notwithstanding his duty to lay out reasons for his decision (*ratio decidendi*) in painstaking detail. Taken at face value, this is a very odd claim to stake. For if there is no formal training for judges, how are they to learn how to be one? And if their findings are to be predicated on the unspeakable and unwriteable, how is a formally unprepared person to assess to such predicates? And, putting that briefly aside, can it be said of Frederika that she, too, has the capacity for sound non-algorithmic judgement [16,17]? [19]

It falls to Courts of Appeal to judge the cases before them in strict accordance with settled law. Sometimes, however, the Court is unable to find in contemporary jurisprudence

a clear resolution of the matter at hand. Lacking the option to leave the matter undecided, the Court must in accordance with its best judgement reach a finding that best fits the present facts. When this happens, a *precedent* is created and immediately takes on the gravamen of settled law. In the old way of speaking, the finding is subject to *stare decisis*; that is to say, it is binding on decisions in future cases. The decisions of a higher court are binding on all courts below, and courts at all levels must not disoblige their own prior decisions. [20] It would only be natural to expect the wise enactor of a new law to take the pains to write it down with all the care that he can muster. After all, we expect no less of the legislative branch, where much effort is given to the orderly writing of bills and, upon passage, to wide public promulgation. It is not that way in Courts of Appeal. Here the doctrine of *lex non scripta* prevails. While the Court has a duty to give detailed reasons for the *application* of a precedent in the case it must decide, it must not accord the same service to the precedent itself. It must not write it down. What is more, in common law jurisprudence, this is a service that it *cannot* provide; for any attempt to write it down would misdescribe it and therefore nullify it as law. Of particular importance is the likelihood that when an appellate court is making a *new* law in a case unsettled by the known precedents and the present facts, it is doing something resembling what Riemann made—new mathematics—in the absence of settled provenance that would have made his method of advancement unnecessary. High court justices are not on the whole geniuses and have no need to be to perform their functions. But perform they cannot without some capacity for originary thinking. Could Frederika be engineered to do the same?

The peculiarities of legal reasoning are not the sole preserve of appellate courts. The instrumentabilities of legal practice and juridical oversight in criminal trials raise epistemological questions of the first importance. [21] For one thing, all evidence to which the jury is privy is conveyed by say-so under oath, leaving the jurors at two removes from what actually happened. Even when the evidence is conveyed in eyewitness testimony, it reaches the jury in doubly filtered form. The witness is not free to give his or her own account of what was directly observed, and is restricted by the requirement to say nothing except in answer to counsel's questions. Since counsel are *parti pris* in these matters, these have questions have the effect of cherry-picking the evidence. No case will go to trial if the defence thinks there is no reasonable prospect of acquittal, or the prosecutor thinks that the likelihood of conviction is not strong. Bearing in mind that a trial is an adversarial proceeding, it is virtually guaranteed that once the case is given to the jury to decide, the evidence given in testimony will be internally inconsistent. There is no better site than a jury room in a criminal trial for a logician to probe the interstices of the management of inconsistency. Only one of two outcomes results in a verdict: Either the accused is guilty as charged, or it is not the case that he is guilty as charged. A further requirement is imposed in no other area of jurisprudential practice. The verdict in a criminal trial must be *unanimous.* It is the only body of common law judgement held to this requirement. In a nine-person Supreme Court, sizeable junks of parliamentary provision can be struck down on a five-to-four vote of the justices.

It is easily seen that the jury is faced with a tricky abduction problem. Its targets are pre-selected—they must either convict the accused or acquit him. They must not supplement what they have heard and seen at trial by means of a search-engine, for example, and they must not talk things over with a lawyer brother-in-law. Even so, although the *evidence* given at trial is all that counts as legal evidence, it cannot be overlooked that a juror might be moved by what he saw in court, perhaps an evasive face or a possibly rehearsed answer. No trial judge known to me has denied the juror the use of all that they have heard and seen at trial. In addition to that, there will be occasions on which the judge will instruct jurors to use their common sense and draw upon their experience of the world. So, in addition to testimonial evidence, jurors are expected to draw upon background information. A final piece of verdict-enabling wherewithal comes from the trial judge's instruction on matters of law which jurors have a duty in law to obey. It is open to question whether a juror's task calls for abductive action answers to Peirce's deductivist understanding. It is hardly likely

that a juror at his humanly possible best will hit the target set by him by reason of the fact that if the juror hypothesized the truth of the evidence he himself is acting on, his verdict would follow as a matter of deductive certainty. But we can certainly rule out that since the evidence *in toto* is inconsistent, his verdict does indeed follow of deductive necessity, for whether guilty or not, the opposite verdict also follows. There is neither space nor need to decide these matters here. More of the abductive character of criminal jurisprudence can be found in Woods [18], and solving the inconsistent evidence problem is discussed in Appendix G on "Inconsistently based verdicts". [22]

What we find in the present case is something rather common in high relief. Solving abduction problems is business as usual and everyday practice in Fred's cognitive life. And like the juror, Fred's background information is known to be inconsistent much in the way that deep memory is. Some are of the view that any high-functioning big information system is pervasively inconsistent. All such systems are *inconsistency-robust* [19,20]. [23] From this we can glean two valuable insights. One is that Fred's epistemic balance is not at all upset by the truth of *ex falso quolibet*. [24] The other is that incoming information has good filtration mechanisms to screen out *spotted* inconsistencies upon arrival. I will come back to *ex falso* a section ahead.

The core of these goings-on lies in the fact that often—rather often in fact—there are things we want to know that we do not know and will not get to know by the standard methods of epistemic update. Suppose that some Omni-Knower made an appearance and offered us the gift of a knowledge-base large enough to obviate all need for abductive striving. Should we take the offer? I have two reservations about this. One is that we ourselves would likely lack storage space for so large a knowledge base as that. The other is that if a new site and rationale for originary thinking were not somehow to be found, we might suffer a grievous loss of the wherewithal for progressive and revolutionary science. Again, in very large measure, our cognitive prosperity is a product not of what we can do but rather of what we cannot. In large ranges of cases, it is in the conversion of disadvantage to advantage.

Frederika stands at least at one great remove from Fred. Her database dwarfs Fred's. If it were big enough to avoid the effort, time and other costs of abductive striving and uncertain outcomes, and if we asked whether, in those circumstances, she would be able to create out-of-the-box science, in some quarters we would find an automatic theorem-proving answer and the prospect of a troubled future. For all the earned welcome of automatic theorem-making, how ready should we be when the next great theorems of the upper reaches of mathematics are too long for Fred to prove, and once proved by Frederika, impossible for Fred to understand? [25]

1.4. The Intentional and the Impalpable

We come now to one of the most mismanaged phenomena in logic and mathematics. It arises from the plain fact that Fred cannot scratch his itchy nose with a low-hanging cardinal number. We are beings who stand in the profitable relations of aboutness, in which we have knowledge of what the objects of thought are about. This brings us to a place of fundamental and longstanding confusion. When Fred is thinking about the number π, there is some object of his intentional acquaintance which bears fundamental relations to other imaginary numbers and, as often averred, even to the positive integers. Pi is available to Fred in respects other than referring. It is something he can make true predications of, and from which in turn further knowledge can be got by inference. There are, however, relations which Fred and π cannot bear to one another. Fred cannot display π on his mantlepiece or shoot it with his iPhone, and π in turn cannot refer to Fred or in any way be palpable to him; it has no mantlepiece on which to display a portrait of Fred, and could not have Fred as lunch guest if lunch were even a possibility for π. From antiquity onwards, there has been stout resistance to these ways of speaking. Numbers, it is said, are abstract. This is true when they arise by abstraction from something already given. In its more widespread attribution, it confuses the abstract with the impalpable (to us), and this leaves plenty of

room for objects to be created by originary thinking or by what Frege called "constructive definition". Whether brought about by abstraction or by creative stipulation, the common feature of such objects is not their abstractness but their impalpability.

In the present day, there is a body of thought according to which cognitive contact by beings like us with beings like π is impossible, and made so by the fact that knowledge of π is not possible between spatiotemporal beings like us and non-spatiotemporal beings like him. The reason given is that reference requires an initial point of causal contact between one of us and one other on which the name that names it is bestowed (or somehow licensed). And since reference is a condition on attribution, and attribution a condition on belief-that, and belief-that a condition on knowing-that, there can be no human knowledge at all of anything non-spatiotemporal. So, then, it will be necessary to find the cleverest among us and mandate them to dance as fast as they can if mathematics is to have any future at all [21–24]. [26]

It is an arresting pivot and a huge mistake, and an utter molestation of What Actually Happens. All we need to show it is the poet of Peirce's mention, but a better choice would be the world's best-known protagonist in detective fiction, Sherlock Holmes. Story-making arose concurrently with the emergence of conversational speech and is one of humanity's oldest and most treasured creations. In relation to those of us who read him, Sherlock is as impalpable as the number π. But to say on that account that it cannot be known who he is or what he did is purblind refusal to heed what the whole of conversant humanity has known from the far distant then until the newly arrived now. Of course, there is a well-known further problem to take note of. Of the legions who read him, Holmes is a real object of reference, thought and knowledge who does not actually exist, and the vexations of causal inalienation are compounded by the discouragements of nonexistence. If one is not careful, one will find oneself denying cognitive contact with beings one cannot kick and of which one cannot say actually exist. Sorry, it will not wash. When a nominalist makes his weary way home after a hard day at the office, and reads his darling wee daughter to sleep with tales of Winnie the Pooh, are we to accuse him of negligence? Are we to charge him with semantic child abuse? Yes, she will in time grow out of the belief that Winnie really exists, but unless she takes on a degree in some distinguished department of analytic philosophy, there will be no thought of her never having known the delights of Winnie and his doings. So we might ask whether Frederika herself is structured for the profitable uses of intentionality. We are now half-way to how things are made true of reality in mathematics. It is a world fact that in inscribing the sentences of the first Holmes story, Sir Arthur Conan Doyle created Holmes in a way that created some of the statements that are true of him. All these creative makings are concurrent, and that made a crucial difference with creative truth-making in mathematics. When Riemann first framed the n-dimensional manifold, he created a putative mathematical object and enshrined it in a hypothesis. Truth, as we saw, came later when the marketplace of ideas signed off on it.

If we heeded What Actually Happens, we would see that everyone who is not a philosopher, and a great many of those who are, take it as given that when someone says something, there is something that he said. It is also understood that what Fred said just now could also be said by François even though François neither speaks nor understands English. A common word for what on those occasions Fred and François said is "proposition". It is open to legitimate question whether whatever that can be said by François in French can be said in English by Fred, where in each case the something said is a numerically identical proposition. But there is no doubt that when Fred or François assert anything at all, they do so by a device that expresses the proposition. The dominant forms of propositional expression are sentential utterance and inscription. Some philosophers are troubled by this line of reasoning. It smacks, they think, of undeserved commitment to abstracta. While I can tattoo on my forearm the sentence "Water is wet", I cannot put the proposition it expresses in a bottom drawer of my desk. The sad story continues: since abstract things are impalpable to us, they cannot be referred to (and so on), and cannot

even be expressed. These reservations repeat the ones we have examined and disposed of. We need not trouble with them here.

There remains, even so, a question to be answered. How does it get to be the case that Fred is able to use the English words "water", "wet" and "is" to express the very proposition he does. The (very) short answer is that he learned his language by learning to speak it, and he could not have done that on his own. Calling his native language his "mother-tongue" reflects a basic semantic fact. What enables sentences to express propositions is the meanings they acquire, and Fred learns these in the course of learning to speak. Meanings, we could say are solutions of coordination problems in a cognitive economy. So, although sourced in linguistic communities *in terra firma*, propositions are impalpable to those who express them. Philosophers who find this troublesome should reconsider their pretheoretical assumptions.

This would be the juncture to being in play in Charles Morris' famous trichotomy of the dimensions of natural language—the syntactic, semantic, and pragmatic. Loosely expressed, the syntactic is the dimension of a language's grammar, the semantic the dimension of its meaning, and the pragmatic the dimension of its usage [25]. [27] In formalized contexts, logical syntax gives a formal language's grammar and deducibility regulae, the semantic is the language's dimension of model-theoretic truth and consequence over wffs of the system, and the pragmatic is the empty dimension. This marks a radical disjuncture between the natural and the artificial. In natural language, the pragmatic dimension dominates. It is first *in ordo essendi* and it is first *in ordo cognescendi.* Again, every working element of a language' syntax and semantics originates in the usages of speech in a community of intelligent interaction bound by conventions which serve as solutions of coordination problems [26–28]. [28] We could speak of these as the semantic conventions of propositional expressive speech. Since the propositions they give rise to are impalpable to spatiotemporal language users, we have it that proposition inhabits logical space, and the rest of us reside in cognitive space. Meanings, too, are like propositions. They are sourced *in terra firma,* and whether they have bearing on what happens here, they also affect the structural arrangements in logical space. For example, they will affect a proposition's modal standing in logical space and relations in which it stands to other propositions. Think here of some given proposition following of necessity from some given others.

1.5. Inconsistency-Management

This would be the right place to expose a related difficulty with deductive logic. The problem is the failure of its practitioners to give adequate notice of deduction's three faces. It is one thing for some given propositions to *have* some given proposition as a consequence, and another thing entirely for a cognitive agent to *spot* the consequence, and another thing again for him to *draw* the consequence as an inference. Consequence-having (and its converse, entailment) obtains, or not, in logical space. Consequence-spotting and consequence-drawing occur, or not, in cognitive space. There are no people in logical space, but they are amply on hand in cognitive space.

When an occupant of cognitive space spots or draws a consequence had by a proposition, a relation is born between an earthbound being and an inhabitant of logical space that is impalpable to him. Leaving the puzzling details aside for now, it is universally accepted that when the cognitive agent—say Fred—correctly draws a proposition's consequent as a conclusion of his inference, the inference will be valid if it were not for the fact that consequence-having is a truth-preserving relation in all contexts of spotting and drawing. A special case of error-management is the enormously neglected question of how beings like Fred are able to stay on course to survive and prosper by adroit by largely subconscious management of the inconsistencies that pervade large information systems and admit of expungement in principle only at the loss of their practical value. It is a striking feature of information systems as big as our own that, notwithstanding the inconsistencies which abide there, neither blows up nor falls down, the first as paraconsistent logicians mistakenly believe [29–32], [29] and the second as Frege mistakenly believed [33,34]. [30] As mentioned

earlier, computer scientist Carl Hewitt, calls these systems *inconsistency robust* [20,35]. [31] They keep on functioning in spite of these alethic embarrassments.

We have here an example of the good that comes from paying attention to What Actually Happens. For many decades now, the received opinion has been that deep memory is the permanent home of inconsistency, and we have known since antiquity the hidden contradictions that make footfall in some of our most successful theories. We also know of cases in which the theories' sponsors soldier on to great advantage, never mind the inconsistencies within. Then, too, we are awash in concurrent information-flows from different and rival sources that hit our receptors without alarm or undue consternation. Where applicable, our information-filters separate what we keep from what we expel, but large amounts achieve unfiltered arrival and repose in large reservoirs of raw information, functioning in large part as background information. Nowhere amidst all these data is there the slightest hint of explosion or any worry over the possibility of it. We also learn something essential about the difference between consequence-having, consequence-spotting and consequence-drawing. Since, save for truth-preservation, nothing that holds for consequence-having is dispositive for spotting or drawing, *ex falso* can be a free ride for having without laying a finger on spotting or drawing. Paraconsistent logicians have wonderful insights, but they have had the misfortune of applying them to the one dimension of deductive involvement in which they do not obtain. We should bear this in mind when we talk to Frederika's engineers.

Every logician knows that *ex falso quodlibet* is entailed by the intuitive concept of consequence-having, and nearly every logical theorist takes a position on it. Approached in the "classical" manner, an inconsistent theory is not disabled by inconsistencies within. According to the nonclassical camp, the inconsistency imperils the whole rationale of theories, essentially by overwhelming a theory's capacities for cogent truth-preserving proof. For close to a century, a set piece for logical theory has been to reinforce or destroy any proof of *ex falso*. What is usually overlooked is the sheer abundance of inconsistency, other than what *ex falso* entails, that flows from truth conditions that define the natural-language consequence relation. It can be formulated as follows:

> *Consequence-having*: S′ follows of necessity from $S_1, \ldots, S_n$, just in case there is no respect in which it is in any way possible for the S_i to be true and S′ not. [32] [36]

From these conditions alone, we also have it that consequence is truth-preserving, monotonic, reflexive, transitive and adjunctive. Taking these properties alone, it is child's play that a true entailment has a truth-preserving maximally inconsistent antecedent-expansion, L: the set of all and only the propositions expressible in the language in which the entailment was expressed. L, at a minimum, is uncountably large and inconsistent at every turn. So whenever we have it that $S \vDash S'$, we have it as well that $L \vDash S'$. In a few more steps, the inconsistencies ramify, and before we know it, in a nice old phrase there is more inconsistency than you can shake a stick at. Again, if we are wise enough to give What Actually Happens its head, it is clear at once that all of this is true of consequence-having, and none of it is worth a tinker's dam for proof theory.

With great consternation, some logicians have flocked to the dogma that, in its present state, all hope of valid and sound reasoning is lost and gone forever lest some restructuring of consequence-having restore the proper order of things in cognitive space. Yet with it comes the danger of self-hoisting. For should the reasoning brought to bear in restructuring consequence-having to some prospect of cognitively positive usage embody assumptions flowing from entailment's natural state, where is the advantage to be got by hoisting oneself on one's own *pétard?* The safe and only way of proceeding lies both elsewhere and at hand. It requires the proper exploitation of arrangements presently in place. In its own domain of logical space, consequence-having has but one "normal function", which is the entailing of things and being entailed by things, and never mind its transfinite and inconsistent issue. How, we might ask, can a relation whose sole semantic function is to stew in the juices of its over-production be of any use to the human reasoner? We have the answer and the relation that delivers it already in train. It lies in our capacity for productive contact with

the impalpable. Once freed of heresy that having is "normative" for reasoning, the ways of blockage lie within our command. In the first instance is the filter by which *premisses* for reasoning are separated from propositions for which no premissory use is presently in view. Thinking here of the cases in which inferences are drawn from stated propositions, we see at once that inconsistent propositions have no premissory future in the cognitive space of the matter at hand, *never mind* how they go far in logical space. So we have a fundamental question to ask of the facts that entailment entails of itself:

- Which of these facts, beyond truth preservation, are dispositive for truth-preserving reasoning?

  The answer can be formulated as

  The first law of proof-theory: Aside from truth-preservation, *no* fact about the entailment relation which follows from entailment itself is dispositive for truth-preserving and cognitively tenable proof.
  Corollary: Not, therefore, *modus ponens*, not *modus tollens*, and not either belief-closure.

There is a reason for this. Aside from truth preservation, it matters greatly what a proof is *for* and what it is a proof *of*. It also depends on the premisses that are on hand and eligible for employment. If we endeavoured to prepare an exhaustive manual of the telically appropriate with respect to premiss-eligibility, case-to-case, we would be sure to go awry. The distinction between entailment and inference is as old as your hat, but logicians appear incapable of breaking themselves of the habit of shopping for good inference regulae in the truth conditions for entailment. A full-service logic of deduction would still the impulse and turn its attention to the greener fields of real-life reasoning behaviour. [33] The consequences for management of robust inconsistency should speak for themselves.

1.6. Talking

One of the things its beneficiaries want from AI are machines worth talking to. For that to happen, Frederika must be capable of engaging Fred in conversation in a language they speak in common. If What Actually Happens were given its due, Frederika and Fred must know how to exchange differences of opinion by conversational means. They must know how to argue. And when context requires deductive engagement, both parties must be able to master, to both respective and collective advantage, the three dimensions of deductive involvement. With this comes a spot of bother. The last time that a logician negotiates to advantage the deduction's having-spotting-drawing dimensionality was when the founder of systematic did it in *Posterior Analytics* [37], without express invocation of the distinctions as I have labelled them here. Aristotle produced the first sound and complete demonstrative metalogic for the deductive sciences. The large spot of bother—a big splash of it, really—is what it takes to engineer Frederika for these same ends, not least her capacity for spontaneous conversational responsiveness to the refutationary efforts pressed by Fred. At the heart of these questions is what it takes for Frederika to speak English in the absence of some antecedently acquired natural language of her own. This is not, of course, a fair or intelligent question. For however Frederika learned her way around any human language, it was not in the way that Fred did. And where, pray, would the acquisition parallels lie?

Since its inception in the 1960s, AI has been a device for doing some of the work that Fred does. It was a special purpose device. At that stage, Good Old-Fashioned AI, as John Haugeland whimsically put it, was a problem-solving machine that matched explicit representations of general principles to particular data-sets, enabling early forms of proof-regulation and relation algorithmic tasks. [34] From the beginning, most of the underlying platform logics for computers have been of the classical first-order kind [16,38–43]. [35] Computation is a beneficiary of its strengths and an inheritor of its limitations. Suffice it for the present that Good Old-Fashioned AI has two characteristics which I myself take Fred to *lack*. One is that it operates on *representations* of its data. The other is that its operating procedures are capable of *express articulation* in all cases. While the special-

capacity machine of Good Old-Fashioned AI was built for instantiation of general principles already known or assumed, in the second wave of the 1980s onward, the converse relation was emphasized. Machine learning emerged in which computers were programmed to derive generalizations from very large sets of simple datasets, as in pattern recognition and updating belief- or knowledge-representations in response to new information which, millisecond by millisecond, is as much of a cascading constant of the system's set-up as it is for Fred's, too. The task-completing capacities of the modern computer are now vast and greatly varied, and there is talk in some quarters of using computational measures such as those of AlphaZero to help in solving how to artificialize general intelligence; and there, as is said, is the rub. It prompts us to ask, "What problem is that, and why would it take a computer to solve it?" If it does not fall to beings like Fred to solve it, what would a software engineer, himself a being like Fred, tell the AlphaZero fraternity to do?

My reservations about characterizing Fred's linguistic and cognitive capacities in representationalist terms can be lightly sketched as follows. [36] It is a mistake to characterize a declarative sentence as representing a propositional belief. The actual connection is more basic. Such sentences *state* those beliefs when uttered with the assertive intonation contour. Similarly, a declarative sentence's normal form is *statemental.* It required some further device, such as quotation marks, to exhibit the proposition it expresses in unstated form. [37] It is the same way with adjectives. When I say that the robin's breast is red, I am not representing it to be; I am attributing the colour to the robin's front. Still less am I representing the thing before me as a robin; I am referring to it as it is by the name that robins have. It is the reflection of the fact that in the pragmatic dimension in which syntactic and semantic objects are put to use, they take on *a telic* character. The first logician who gave express notice of this was Aristotle when [37] he provides that when a proposition is a premiss (*protasis*) for a conclusion in an argument, its role is to secure the judgement expressed by the conclusion. The conclusion in turn is that which gives what the premisses are *for.* This same telic character is taken on when a truth-preserving finite sequence of propositions is converted to argumental form. Premisses, statements, arguments and their like are objects of the language in their full dimensionality. So it bears repeating that for *homo sapiens,* language was born to be used; it was telically oriented at inception. In its further stages of development, it achieved the capacity to use itself to talk about itself, at which point the syntactic and semantic dimensions become objects of study [44]. [38]

If we leap ahead to a concept-writing notation of Frege's [45] and the more standard notation derived from Peano and Schröder of Russell's 1903 [45], we find artificial languages (so-called) in two dimensions, not three. What is missing is the pragmatic and all that it means for purposive language-use. Frege and Russell twigged to this in both places there (and elsewhere) introduced a notation to mark the difference between a displayed proposition (or sentence expressing one) and the assertion of it—the assertion-stroke, '⊢'. There was a good deal of back and forth in logical theory between what a proposition asserts and what it merely implies (or presupposes or suggests). But the point to issue here is that '⊢' is of gestural usage only and carries with it none of the meaning, the significance and rich detail of what would have been revealed had the languages in question been allowed a pragmatics. Such assertive posturing was to imperil the health of theories of proof and inference from that time to this.

When we examine the meanings actually possessed by the English words "represent" and "representation", they reflect a certain kind of relation between something and some other different thing, often very different in kind. When a nation's ambassadorial representative to the U.K. casts a vote in the General Assembly, the vote is not his or hers, but has been cast in right of the country the ambassador represents. In this case, the ambassador is representing his or her principal and, in punching the Yay-button, she caused it to be the case that her country has voted affirmatively. The same relationship exists between a person and the executor of his estate. The role of the executor is to act in his principal's name and do his principal's own bidding upon the principal's decease. No nation can vote without the assistance of a human representative, [39] and no deceased person can deed the farm to his beneficiary in the absence of his legally appointed representative with pen

at the ready for signature. A quite different and less often used meaning of "represent" is often inscribed with a hyphen. So when the billboard at the theatre announces the re-presentation of the performance at 13:00 h, the company is giving notice of a *repetition* of the performance.

There is a large cluster of binary relations obtained between the likes of us and others, sometimes quite other than we. We build a maquette as a model of the new concert hall. We photograph the young bride and groom. They give their daughter the name "Kelly", while others paint the Queen's portrait, and others establish isomorphisms between set-theoretic structures, and so on. But to call any of these makings representations or the objects of representation is to avail oneself of the liberty or originary thinking, in this case the luxury of making things mean what they do not in fact mean. In short, the luxury of Humpty Dumpty semantics. [40]

Now that the third-wave enthusiasms have started to percolate, the hope is that, whereas computers are now good (very, very good) at solving a range of problems human beings also solve, often less quickly and efficiently, the time is ripe for a considerable enlargement of repertoire, in eventual fulfillment of Herbert Simon's 1965 article of faith that by twenty years later, machines will be able to do any work that the common man or the man in the street can do. The idea of artificial all-purpose intelligence captures a significant part of that hope, adjusted for prediction-time. The aim is to make machines that manifest all the kinds of intelligence implicated in a human being's participation in what cognitive-systems theorists call the *cognitive* economy [46–48]. [41] I come back to this two sections hence.

It bears emphasis that, up to now, the development of Good Old-Fashioned AI and machine learning has had little involvement with the empirical sciences of cognition and virtually none with the philosophy of knowledge [16,17,49,50]. [42] Phase-two developments showed a marked indifference to cognitive studies, and even those who employed artificial neural networks were more preoccupied with advances in engineering. "To the extent that modelers withdrew from pursuing cognitive investigations, the design of neural models was allowed much more freedom in adopting mathematical solutions *alien* to mental processes." [51–53] [43] This is hardly surprising. Computer science was invented by mathematicians and took some time before being its own thing. [44] The logical involvements of that time were of the first-order kind, and it cannot be left unsaid that the record of first-order modelling of what humans do when they draw conclusions from incoming information is rather dismal [46,48,54–56]. [45] When computer scientists turned their attention to AI, these limitations could only press more tightly. Born of a discipline which had had no real grasp of the cognitive realities of human beings, AI would still be calling the shots both conceptually and operationally, but without adequate awareness of the frailties of this inheritance. When someone asks a software engineer to computerize some conceptual material he had given him, his Rubicon has been crossed, and it is the modeller who's now in charge. It has its comical moments, to be sure, but it can be in truth a serious problem.

As the theory of reasoning has passed from the old logic to mathematics and from mathematics in turn to the AI engineers, we find ourselves at two distant removes from human cognitive reality. "The old logic" is a generic name for all logical systems dating from Aristotle until, let us say, 1879, the year that Frege published his concept-writing notation. The focus of the old logic was on the ways and means of the deductive sciences, and its discourse domain was the cluster of properties intrinsic to the successes and failures of deductive reasoning and argument—concepts such as deductive consequence, truth-preservation, proof, contradiction and inconsistency. The medium of exchange, both at the inferential and metatheoretical level, was the theorist's own home language or some other more widely used one. Also from early on, non-truth-preserving relations were given theoretical attention, beginning with Book A of Aristotle's *Posterior Analytics* [37] and rising to the heights (or otherwise) of subjective probability theory, modern-day Bayesianism, rational decision theory, error-statistical experimentation theory, ranking theory and the like. These developments are not, however, in the reference class of "the old logic". But they

are said to be enormously important for computer science. As long as we hold fast to our self-assigned duties to get to the cognitive bottom of Fred, we are in a tricky position. As we have it today, there is nothing in deductive logic that respects that duty, even as regards those comparatively few occasions when Fred is reasoning deductively; and nothing to better effect is to be found in the ample findings of nondeductive logic. Deductive logic does not give a hoot about Fred, and inductive logic squanders its own sophisticated resources on abstract idealizations of Fred, with respect to which no known approximation relation has yet been defined for him. Whatever their failures of Fred, no like charge can be brought against them in regard to Frederika. For her, such measures are peaches and cream. Why, then, would we think that measures which work well for Frederika could succeed in making her in her intelligent moments just like or even better than Fred is in his own? Why would AI theorists even frame the question, and why would philosophers so much as countenance it?

1.7. Brains

Today's Deep-Learning hopefuls are keen to repair a longstanding omission. There is now serious talk about building a computer model of the human brain. It is granted that brains play a huge and indispensable role in the operation of the human cognitive economy—some say that they are its whole operational show. But given the present and foreseeable state of neuroscience, all talk of computer modelling is idle, and everyone knows it. For an excellent but baleful overview of the failure of neuroscience to achieve any standing as an experimental science, see Matthew Cobb [57]. [46] Cobb is not one to pull his punches. Phrenology was "guff", Freud "had nothing novel or insightful to say about how the brain worked", and gene-editing involvements only reveal the "thoughtlessness" of "one smart-aleck researcher" [57,58]. [47] To this I would add that, whatever its particulars, the decision to "inferentialize" the causal pathways of the neural wiring that brought Fred to a knowledge of the bird's breast colour, steps must be taken not to give "inference" a meaning it cannot have in such attributions. Attribution of inferential activity to brains has all the theoretical allure of a sophomore's anthropocentrism. All the same, the fallback position is that it is still worth beavering away at, since a good outcome is a possibility in principle—a move, as some would say, from the idle to the empty. Still, hardly anyone in 1980 could seriously have imagined the likes of AlphaZero [59] [48] or self-driving motor vehicles [60]. [49] So who knows—really knows—that the present day's unimaginably designable cannot happen again? [61–63] [50] Very well, if we are to keep an open mind about this, we must up our game considerably about the very idea of all-purpose intelligence. For that, we'll need to take some measure of the cognitive economies of Fred and his kind.

All beings in the natural order as a condition of survival and biological prosperity are met with ecologically sourced coordination problems, each in his own way a bringer of things about and the recipient of what others have brought to pass. Beings like us are dwellers in a particular kind of ecological habitat, some of whose more distinctive features are functions of our neural makeup; Lorenzo Magnani characterizes these systems as "eco-cognitively open" and "unlocked".

> "This special kind of 'openness' is physically rooted in the fundamental character of the brain as an open system constantly coupled with the environment (that is, [it is] an "open" or "dissipative" system): its activity in the uninterrupted attempt to achieve equilibrium with the environment in which it is embedded, and this interplay can never be switched off without producing severe damage to the brain." [64–66] [51]

The permanent attachment of the brain to the ceaselessly changing environment subjects it to its own unending causal wash which, upon arrival, is dispersed to the relevant causal-processing units [67]. [52] The constancy of causal refreshment enables the system to stabilize and maintain ecological equilibrium. When a causal input carries information, it is often dispersed to the system's information processors in linguistically expressed semantic form—"Watch out for the bus!" But by far the greatest share of it is processed as

the potential not the actuality of syntactico-semantic realization. And much of that is stored in the unconscious for possible transition to propositional form under the right cues. There is a common saying which captures the drift of this idea: "Good gracious, how interesting! Just think! I knew it all along and did not realize I did until just now." This is well worth noting. The widely distributed experience of not knowing that one knows until something nudges a later awareness of it is a fact so large, so representative and so common as to merit some principled attention from epistemologists. Let us call this the "unknowing knowing datum". This is ripe occasion to recall Peirce's intellectualist repudiation of subconscious knowledge. But it will also be a good place to record some relevant observations. Writing of the facts which the logician cannot ignore,

> "come within the range of every man's normal experience, and for the most part in every wakening hour of his life." [1] (1. 241)

Moreover, they constitute

> "the universal data of experience that we cannot suppose a man not to know and yet be making enquiries." [1] (4. 116)

So it is not clear whether Peirce can hold his intellectualist serve.

It is also helpful to bear in mind how much a human being will have come to know even by the onset of speech, and yet how little of it he carries around in the front of his mind. We have no reason to suppose that we retain all the information that washes over us, but there is good reason to think that we retain and store very large quantities of it in memory and everywhere else information resides in the rest of the human corpus. Studies of the dissipative brain suggest an analogue in quantum field theory.

> "In the dissipative quantum model of the brain the vacuum code is taken to be the memory code. Again memory is represented by a given degree of ordering. A huge number of memory records can thus be stored, each one in a vacuum of a given code . . . . In the dissipative model, *all* the vacua are available for memory printing." [68–70] [53]

While I admit to general reservations about over-use of representationalist assumptions, there are two especially valuable insights to be found in dissipative brain studies. One is our mastery of vast quantities of information. The other is that organisms as complex as we call for the attention of more branches of scientific theory than are usually called upon. If, as may well be the case, that there is a place in this enquiry for quantum analogies, there could be analogies to be found in physics macro-regions. Such is in fact the case, as I will argue below, with thermodynamics.

It would be a large mistake and, I think, a silly one to overdo the experimental challenges that face the brain sciences. Experimental science is a stern taskmaster to which mankind owes large swaths of its cognitive prosperity. But holding all empirical enquiry into the human condition to its standards is an endless impoverishment of what how cognitive economics is actually good for. Two rich sources that cry out for further theoretical exploration are the large databases generated by the study of the brain and the fruitful hypotheses induced by their peculiarities and unanswered questions. As Christopher Mole points out, experimental scientists are less practised and less good at hypothesis-generation, without which progress in science would be considerable hobble. Philosophers, by contrast, are dab-hands at the arts of conjecture (Mole [58]). At the beginning, we announced our subscription to the What Actually Happens rule. It behooves us to shape our thinking in light of what appears to us all to actually happen. The focus of what happens is what happens when Fred makes his routine way in cognitive life. And since Fred stands in for everyone, we have been respecting the rule by attending to patterns of What Actually Happens in the cognitive lives of beings like us. For this to be possible, a good deal, at least, of What Actually Happens must be discernible in the behavioural regularities of the observably cognizant. In shorter words, what the rule requires of us is that we delay the gratification of normative epistemology until we have examined how Fred fares in

the cognitive economies to which he is party. The section that follows could be called a prolegomenon to any future epistemology.

1.8. Cognitive Economies

A money economy is a nature-based interactive multi-agent intelligent cooperative—an ecology, for short—for the generation and circulation of wealth. A cognitive economy is a nature-based interactive multi-agent intelligent cooperative for the generation and circulation of knowledge. It is both prior to and a precondition of any money economy within its midst. The natural resources of the cognitive economy include neurophysical and social embodiment, placement in the causal order of nature, information-processors, time, language and intelligence. Its input capital is information-flow. Its capital developments include conversation, communication and the sundry devices of systematic enquiry. We could liken a healthy cognitive market to a competitive system in wide reflective equilibrium, and a proposition with good market-share to a hypothesis security stabilized by such an arrangement. The expression "wide-reflection equilibrium" was conjured by John Rawls, but the concept it expresses originated with Nelson Goodman's classical paper on how logical principles are certified. Roughly put, a principle of logic is correct to the extent that regular practice comports with it, and reasoning is in no logical duress to the extent that it heeds the logical principles [71–73]. [54]

Parties to these economies are biological beings, residents of the natural world and the subjects of the causal order. They operate with the strength of their natural endowments and the advantages of their learned ones. They are also subject to the limitations that apply to beings of their nature and circumstance. They operate with limitations of relevant information, storage-capacity, memory and information retrieval, time, energy and whatever else. Herbert Simon is famous for having coined the term "bounded rationality" for beings subject to these natural conditions of operation. [55] "Bounded" is right, but "rationality" is wrong. What's bounded here are natural resources, an economic constant. Because knowledge is acquired and circulated by natural beings in patterns of behaviour whose regularities are empirically discernible, no account of its role in the cognitive economy could pretend to empirical adequacy if it neglected these behavioural patterns.

Fred is an organic information processing being—an embodied seeker, acquirer, absorber and transmitter of knowledge. He is nosy and helpful. He is a gregarious knower. As a natural being, he has a varied and multi-layered place in the causal nexi of natural life. As a social being, he owes his survival, endurance and prosperity to cooperative interaction with his fellows, in patterns of behaviour both causally enabled and underwritten by social convention. Fred embodies a distinction which has yet to receive adequate theoretical elucidation. It is the distinction between energy-to-energy transductions and energy-to-information conversions. Too often, theorists conceive of the natural limitations on what Fred is able to do and have done to him as disadvantages—setbacks—which deny him optimal outcomes. But optimality here is an unnatural attainment. It carries no more weight than the fact that, unlike the fabled cow, Fred could not jump over the moon if he set his mind to. [56] Simon was subject to this misconception when he coined the term "bounded rationality", as an informal metric of the distance of Fred's rationality from that of the deity, or some other logically omniscient being. Simon is closer to accuracy in pointing out that Fred is also a "satisficer", a being who knows the wisdom of the adage that perfect is the enemy of good; but Simon loses his edge by adding that satisficing is a suboptimal achievement, and in so doing confuses the optimal with the maximal. Sometimes we quit before getting a better result than the one we have decided to settle for. There is an overarching cost-benefit reason for this. The cost of the outcome is not offset by the betterness of its attainment. Again, there is nothing at all unnatural about what is plenty good enough. It is neither the heartache nor natural shock that flesh is heir to. [57]

In a nice turn of phrase from John Locke, Fred is a man of parts. Among his other endowments, he is a person of *judgement*, a person whose intelligence enables him to use his head—his common sense—in dealing with contested matters or matters that are otherwise

challenging [74]. [58] To the extent that he is so enabled, Fred is wise. The word's proper antonym is "unwise". Another is "stupid". Some of the smartest people we will ever run into are, in matters of judgement, simply too stupid for words. They may have ascended to the heights of algebraically stacked topological vector subspaces, but they could not lift a finger to assist in the resolution of a collective bargaining grievance [75]. [59] They lack judgement. They are not wise. Any person of sound judgement and requisite will have the intelligence to learn from experience to be an Appeals Court justice. Fred, I say, is just the ticket. Where, we wonder, would this leave Frederika? In [17], Brian Cantwell Smith presses hard against the idea that judgement is the sort of thing that presently or foreseeably one could find oneself ascribing to Frederika. In his usage as with mine, judgement is distinguished from reckoning, the sort of thing Leibniz's "Calculemus" called for and machines do so well at. The Smithian notion of judgement is a much more layered idea than mine is. Either way, it raises a question that is critical in more ways than one. For Smith and his like-minded, it is the question of whether the nonalgorithmic enablements of judgement can be matched by the algorithmic enablements of Frederika. For me, there is a prior question. It is whether the causal enablements of Fred's knowledge are themselves possessed of an algorithmic nature. If so, does the algorithmicity permeate Fred's cognitive involvements? Could we have it that the causal enablements of tacit and implicit knowledge are algorithmic, whereas the causal processes of articulated reason are differently structured? Are we, then, to say that Fred's causal ways have a sufficiently satisfactory match with Frederika's algorithmic ways to ground a favourable third-generation verdict on her ability to do what Fred does when he uses his best judgement? It all depends on how far down the manifestation-enablement chain Frederika must go to produce a good third-generation outcome.

All competitive free markets of ideas are dialectically structured. They are sites of trial by combat. As noted by Aristotle in Book A of *An. Post*. [37], it lies in the routine management of a cognitive economy, in which hypotheses gain ascendancy as known truth, that they are the products of dialectical exhaustion. They will have survived all efforts of refutation from the top experts in the field to show them incapable of consistent defence. Seen this way, a proposition that has traction in an economy of wide reflective equilibrium has earned its alethic spurs until, as may be, overthrown by future market-eruptions.

1.9. Jiu-Jitsu Advantage

We come now to one of the most distinctive features of Fred's cognitive make-up. Like all the rest of us, Fred knows lots and lots of things about lots and lots of different things. We are abundant and versatile knowers. He, and we, also make lots and lots of errors about lots and lots of different things. Here, too, our errors are numerous and widespread. But, like the rest of us, Fred has the large advantage of feedback mechanisms which enable their detection and correction. Like any active partner in an economy, Fred is both subject and alert to cost-benefit considerations. Beings like Fred weigh the costs of errors which get corrected after detection against the benefits of having avoided these errors in the first place, and for large ranges of cases favour correction over avoidance. Although he rarely puts these considerations into words, Fred is implicitly seized of the huge cost-impact of large-scale error-avoidance. Responsive to the other old adage that much of what we know is learned from experience, Fred calls on the corollary that often the best way of doing so is by learning from our mistakes. It turns out, then, that the detection and correction of error is a powerful positive resource in cognitive economics. Of equal moment is our capacity to make instructive mistakes, rather than the ones from which no recovery is possible. Our error-making ways are a net benefit for cognitive prosperity. They are an exercise in *jiu-jitsu* economics, in which the disadvantage of error is converted to the error-manager's greater net advantage. There lies a critical question for the software engineer. To what extent, if any at all, is there any call upon Frederika to achieve the bounty of Fred's management of his limitations by having them installed in her? I have already mentioned the challenge posed by Fred's management of very large measures of stored

information, and its evident dependency on his wherewithal for an also large knowledge of things, achieved implicitly and tacitly. With respect to the capacity for information storage, Fred is no match of Frederika. Less clear is whether in her case, the devices of retrieval hinge on the implicit and tacit grasp of its contents.

A considerable part of the literature of error management lays heavy odds in favour of what is called predication error minimization. On a face-value reading, predictive error minimization seems just the ticket for a well-made chap like Fred. There will be large ranges of cases in which Fred will forgo acting in a certain way on the strength of his expectations of an untoward outcome. In framing assessment of risk, predictions of outcome probabilities are a matter of course. What I find most striking about much of the leading work in this area has nothing to do with Fred and his action plan. It has everything to do with perception, such as seeing the redbreasted robin in the tree, and the neural enablements that bring it about. [60] Although in some respects rivals of one another, there are central assumptions on which they converge. It is nicely summed up in Michael Rescorla's [76] (pp. 2–3) discussion of Clark's *Surfing Uncertainty*. [61]

> "Recently, perceptual psychologists have developed Helmholtz's suggestion by modeling perception as unconscience Bayesian inference. On a Bayesian approach, the perceptual system maintains prior probabilities regarding the distal environment (e.g., certain retinal inputs are deemed likelier than others) and prior likelihoods that relate the distal environment to sensory input (e.g., certain retinal inputs are deemed likelier given certain distal shapes and certain lighting conditions). The perceptual system deploys these prior to transit from sensory input to a posterior-probability (e.g., the posterior may assign high probability to the perceived object having a convex shape. Based on the perceptual system chooses a privileged estimate of distal conditions."

Clark's prediction error minimization mode does not instantiate Bayesian modelling, but there is a basic idea that they share. The brain engages priors to form predictions about sensory inputs to compute prediction error. So the error in question here is misperception, and the means of avoiding it are supplied by the probabilistic workings of the perceptual neural system. As such, systems are highly representational and, hence, to some extent risk being at one remove from What Actually Happens. I mention this now not to make quarrelsome joist with brainiacs of representationalist persuasion, but rather to call attention to Helmholtz's notion of unconscious inference [77]. [62] Although Helmholtz is writing about vision, concerning which I would again say that in the general case inference has no role to play, it stands to reason that it might have a hand to play in Fred's acquisition of the knowledge that lies beyond the reach of voice and pen. Knowledge without access to inference seems to me an improbable arrangement. On the other hand, a further feature of Clark's predictive coding approach is the hierarchical structure it invokes. This suggests to me that Fred's cognitive system might also be hierarchically structured, each level of which is a reflection of the inferential variations that are obtained in goings-on of the spritely and alertly awake and the basement-dwelling enactments of the causally efficacious upshots of *information processing*.

1.10. Fallibilism and Epistemic Bubbles

To some observers, a related limitation has the look of a real liability. The cognitive economy is a bred-in-the-bone economy of fallible performance. In the only sense in which fallibilism is interesting, it is a doctrine to the effect that much of what Fred currently experiences himself as knowing, including what he thinks he knows now, he actually does not. There are exceptions, of course, perhaps the law of identity, for example, or Fred's knowledge of his sensorium and so on. Fallibilism induces a partition on what Fred thinks he knows. For some of what he thinks he knows he actually does not. For some other part of it, it remains open whether he knows it or not. For all the rest of what he thinks he knows, he does in fact know it. Of course, the partition is not mathematically exact.

One of the things currently in some doubt is where the lines are to be drawn and for what principled reasons. It is a wise fallibilist who resists the temptation to get too far into the weeds. Suppose now that fallibilism instantiated itself, that is, that fallibilism were false. Then fallibilism could be one of the things Fred thinks he knows to be true which he actually does not know to be true. [63]

If there were nothing Fred now experiences himself as knowing now that is concurrently something he actually knows, Fred would be incapable of recognizing any error in anything he experiences himself as knowing. This would not be a recognition limitation. It would be a limitation on Fred's error-making ways, indeed their outright preclusion. If, *per impossibile*, Fred were incapable of cognitive error, so would all of us be. We would all be successful Simon-optimality freaks. And there would be a new question to ask of Frederika. Would it be possible to build a machine with the infallible cognitive reach of Fred?

A principal factor in our down-to-earth fallibility is what I have called epistemic bubbles. They bear most tellingly on the correction of detected error. Consider a case. When an error you have committed has now come to your attention and you find yourself in a position to correct it, you now stand in an after-the-fact relation to your former self. In the original situation, you could not concurrently have been in a state of knowing for a fact that S is true and also in a state of detecting and correcting the error you have fallen prey to in thinking so. Correction requires a distancing from the original state you were in, to one in which you are now in a state of mind to make some assessment of it. Supposing now that at the distant point at which you think you know for a fact that your former belief was false and that the present replacement-belief S* is true, the state you are now in with respect to S* is the same state you were originally in with respect to S. You can be just as wrong about S* as you were about S. They both might be false, and S* might be false and S actually true. There is no principled or systematic way of evading these possibilities. You are therefore in an epistemic bubble, and for any proposition you could be wrong about, then you could be wrong now about its replacement. Sometimes you could be wrong each time. Escape from the bubble is causally impossible [78–80]. [64]

The no-escape constraint is not scepticism. It is fallibilism. If fallibilism were true in the ways we have been detailing, then in the general case the KK-hypothesis could have no purchase. It might be true that Fred has lots and lots of knowledge, but it cannot be true from simply being in a state of knowing that S logically implies in turn that Fred is in the second-order state of knowing and that he knows that S. Any epistemology that is fit for service in the actual cognitive economy will have to make allowance for this. This would be a good place to flag the ambiguities of "belief". It is generally accepted that its meaning varies as between really thinking that S is true and accepting that it is true [81–83]. [65] To some extent, differences in meaning reflect differences in Fred's state of mind, as measured by the appositeness or otherwise of concurrently held hedges in the form, "but I'm not sure", "I could be wrong" and the like. It strikes me that the distinctions that serve us best are those between *experiencing oneself as knowing* that S is true, *being wholly satisfied* that S is true and *acquiescing* to S's truth. Assuming our acceptance of the trichotomy, the meanings of "KK" multiply accordingly, thereby losing interest as they gain in plausibility. This gives us three different things that belief S can rightly be said to be. It is an epistemic orientation to the truth of S; it is, secondly, a doxastic orientation to the truth of S; and thirdly, it is an acquiescent orientation to S's truth. This tells us something important about knowledge. Fred's knowing that S is not a state of mind. Fred's knowledge is the joint product of his states of mind about S in conjuncture with the state the world is in with respect to S. Truth is the *outlier* here. Knowledge is the collateral benefit of well-produced, well-circulated belief in a stable and prosperous cognitive economy. Much of what Fred knows is brought about just by being alive, breathing, and in good health, and in the requisite alignment with what is truth, much in the way in which he is the producer of CO_2.

Bearing on this is one of the empirically discernible behavioural regularities of cognitive economic life. We met with it earlier. It is the widespread frequency with which people at large will exclaim, "Why, I knew that all along, but didn't actually realize that I

did!" When such utterances are made or such sentiments felt, the word "know" is used in a certain way. It is used in a way that is incompatible with any generalized form of the KK-hypothesis. It is also used in a way that is disapproved of by KK-boosters. As we have it now, we are in a stand-off, each party to which is at risk for *petitio principii.* This leads me to think that a condition on the adequacy of our description of Fred and the relations in which he stands to the cognitive economy is that we find a theory of knowledge among whose provisions lies a resolution of the stand-off with least damage to what we already know of Fred and his cognitive involvements with what we ourselves already know of the world.

There is no doubt, however, that error management is an enormous draw on Fred's intelligence, and a major and indispensable agent of the species' cognitive prosperity. Again, it marks one of the cognitive economy's most productive instrumentalities. It marks the cognitive economy of humanity as a jiu-jitsu economy in which, again, disadvantage is converted to value-adding advantage as a matter of course. Jiu-jitsu conversion is one of the most load-bearing pillars of our cognitive prosperity. If Frederika is to have the all-purpose intelligence of Fred, perhaps she will need an equivalent capacity for jiu-jitsu error-management. Since Fred must exercise his error-management intelligence in a state of perpetual epistemic embubblement, Frederika's intelligence would have to include the capacity for cognitive prosperity in a fallibilist economy. Fred and Frederika alike must attune themselves to the delights of highwire cognition without a net. Strange as it may seem, the last thing that Fred ever *routinely* experiences himself being is as a cognitive highwire trapeze artist without a net.

1.11. The Causal-Response Epistemology

In our observations of late on the cognitive economy to which Fred is party, I have taken cognitive economics, the study of such arrangements, to be a behavioural science and a purely descriptive one which derives its substance from the empirically discernible behavioural regularities of beings like Fred. It is perfectly true that among such regularities, there lie the unmissable marks of a self-regulating economy wholly at ease with the distinction between doing things in the right way rather than the wrong way. The conventions to which cognitive practice is subject operate in the way counterpart conventions provide for correct speech. In both cases, the right way of proceeding is the way in which things normally proceed. Seen in this way, we see in the cognitive regularities of human life the converge of the normative on the normal.

I see cognitive economics as the necessary prelude to any philosophical theory of knowledge that aspires to respect the regularities in view. If at some point the epistemologist thinks that he has just cause to override the convergence of the *normative* on the *normal*, it falls to him to show rightful cause and to make his amends with due regard to the fact that Fred is not a fool.

We have had good occasion above to note the large gap between reasoning that is valid and reasoning that is correct. As we saw, no reasoning is correct just because it is truth-preserving, and that tenability lies in the success or otherwise of reasoning validly in fulfillment of the goods sought by the reasoner and the characteristics possessed by the premisses available to him at that time. As used here and in common usage, "correct" is normatively tinged. When one reasons correctly, one is—all things considered—reasoning in the right way. Charles Peirce is but one of his subject's modern great practitioners to mark logic as the normative overseer of drawing necessary conclusions. He was not alone in not pressing the importance of the ambiguities of normative discourse. In one sense, the emphasis falls on the normal, and in another it falls on the moral. If one seeks to speak good French, one must learn to speak in the way that French is normally spoken. If one seeks to live a good life, one should bend one's ways to the biddings of the Ten Commandments. It reflects considerable credit on Peirce to have placed such *moral* weight on the necessity of intellectual honesty. An intellectually dishonest economy is headed for ruin, all the more acutely and speedily so, when dishonesty corrupts the ways and

means of circulating knowledge. Peirce's alarm is deeper than that of economic self-interest. His worry is that once he disabides the necessaries of intellectual honesty, one loses one's *capacity* for truthfulness [1] (1. 149). We might readily enough grant to Frederika her fidelity to truthful self-interest in her own cognitive economy and/or in its interaction with Fred's. But the question that really bites is whether Frederika can flourish in the absence of the intellectual *conscience* that Fred must have.

Meanwhile, there is the question of what a logician is to do when he thinks that reasoning in the normal way is reasoning incorrectly. A compact way of taking the measure of this question is by examining the track record of fallacy theorists in the interval from 1970 to the present. There is no space for this here. Interested readers could consult my *Errors* for a largely negative answer.

Viewed from all the usual angles, we should emphasize that Fred is two things at once. He is both actor and acted-on. He is as much a doer as a done-to. He is a goal-directed striver and the absorber of causal impacts. In some things, he is the master of his ship, but in most things he is advanced by what happens to him and how he is built. It is true that Fred is a tasker, a problem-solver, a planner and executor, but it is not true that his successes and failures are solely or even mainly the fruits of his own doings. Human knowledge is the gift of the reciprocity of opposites, of making things happen and having things happen to one. Much of what enriches Fred cognitively lies out of sight of his mind's eye, beyond the reach of the heart's command, and unamenable to engagement with tongue or pen (or keystroke). It lies not in the front of his mind, or in the back either. It lies in his cognitive down-below. The very fact of its fruitfulness for good cognitive health is something to pause over. It is not something that the human knower can do without. And it is the very key to the puzzle of *lex non scripta.* But first we must deal with information. To help set this up, we should lay down some markers for what would count as an empirically sensitive epistemological response to what I have been calling data for theory.

Knowledge according to the justified truth belief (JTB) model is a case-making achievement, an act we might say of forensic fulfillment. Fine as far as it goes, the JTB model lands well short of accounting for our present data for theory in a way that risks the odium of unearned scepticism. Beyond doubt, there are occasions and plenty when, in the absence of forensic pleading, knowledge is out of the question. But on any honest reading of the J-condition, its across-the-board application to knowledge as such is an open door to scepticism. To stay true to the data, we must disavow its generality and tread gently when occasion arises to give it consideration. In his classic paper of 1967, Alvin Goldman spotted the weakness of the JTB and, in an insightful move, sought to strip the J-condition of its forensic sense and replace it with a causal interpretation of the J-condition. On Goldman's reading, the J-condition for a proposition S would be met for an agent X when X's belief that S was caused by belief-forming mechanisms in good working order, and working here in the way that nature had intended. I consider the casualization of the J-condition to be the single most important contribution to the advancement of empirically sensitive epistemology since Mill. It broke the back of the forensic fallacy of supposing that human knowledge as such is the fruit of actively intelligent agency, and opened the ground for principled reflection on a cross-kind alternative hypothesis, in which knowledge itself is the by-product of fruitful matchings of *causally well-produced beliefs* that S is true when S actually is so as a matter of objective fact. Seen from this perspective, the J-condition loses its purchase as a general condition on knowledge-that, to be called upon only when context demands it. Required revisions of Goldman's ceiling-breaking insight were made by Gabbay and Woods's [46] and *Abduction*, and strongly advanced in Woods' *Error* and *Legal*, and put to work in 2018 in [84]. [66] To give some sense of this, the Causal Response Model of knowledge, we could write as a first pass,

- *The knowledge as causal thesis*: Fred knows that S on information I when S is true; in processing I, Fred's belief-forming devices were causally induced to produce the belief that S, Fred's devices are in good working order and operating here in the way that

nature designed them to, I is good information (up to date and accurate) and there is no interference caused by negative externalities (e.g., too many bourbon Manhattans).

This is a first stab at sufficiency for Fred's knowledge of S in his cognitive up-above, in which Fred is conscious and attentive and S is both propositionally structured and linguistically formulated. In such a case, Fred has articulated command of a semantically well expressed object of knowledge. The question now is what happens when S recedes from the front of Fred's mind to find cost-efficient harbour in his down-below? It is all, really, a matter of what can be made of information.

1.12. The Place of Information in Knowing

Since the workings of humanity are played out in an information economy, we must pay due heed to the very idea of it. In truth, information is one of philosophy's most unruly concepts. It has been well said that

> " . . . information is notoriously a polymorphic phenomenon and a polysemantic concept . . . ." [84–90] [67]

There are four especially influential members of the unruly family. I will call them the *epistemic* sense, the *probability* sense, the *complexity* sense, and the *military* sense. In its epistemic sense, information is something taken in by a cognitive agent. It conveys how things actually are. Information in this sense lies in the investigative domain of logic and epistemology [88,90–92]. [68] In its probability sense, information is what is channelled from a source to a receiver. The source emits signals with a certain frequency, and the information picked up by the receiver is conceived of as the expected reduction of probabilistic uncertainty. Agency is not a necessary factor in the transmission or reception of information in this sense. Its principal domains of investigation are probability theory and physics [15,87,93,94]. [69]

In its third sense, information has to do with codes. The informational value of a code-spring is the algorithmic of Kolmogorov complexity of the string, which is defined as "the shortest program that computer it on some fixed universal Turing machine." [85,95–97] [70] Complexity information is studied by theoretical computer science, probability theory, statistics and physics. A problem posed by the first two senses is that information in the first sense cannot not be true, yet information in the second sense need not be true, a problem well discussed in [46]. On the other hand, there is a route from the complexity conception to the probability conceptions via the set of all prefix-free programs under provisions of a technical result known as Kraft's inequality. [71]

The fourth notion of information derives from its use in intelligence and counter-intelligence work. According to the CIS's World Fact Book:

> "Information is raw data from any source, data that might be fragmentary, contradictory, unreliable, ambiguous, deceptive, or wrong. Intelligence I information that has been corrected, integrated, evaluated, analyzed and interpreted." [72]

This, the military sense of information, is widely used in computer science and informatics.

Far from a unitary concept, the information quartet is not a happily reconcilable family, leading Hintikka to observe that it is not at all 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." [73] [98]

This raises something of methodological importance. If information is a concept all tangled up in polysemy and internal inconsistencies of usage, why would it assume so important a role in a causal response theory of human knowledge? The short answer is that I do not know. What I do know is that great strides have been made in science and technology such that ill-understood concepts are the central focus. In the case of sets, the very concept was unruly and at most half-baked at birth and remained a trial and

tribulation awaiting the settled maturity that lay many decades ahead. We also know that in that case, the concept of set gripped its purveyors with *la force majeure* of an insistent idea. I think the same can be said now of an information-centred and causally responsive epistemology. It is an idea whose time has come, and well worth the effort to foster a prosperous upbringing. It was Hintikka who said in 1998 that abduction had become epistemology's central problem [99], [74] and it can be said now that real-world mathematical truth-making has become abduction's central problem. If that be so, we must take care to develop an epistemology that accommodates and supports that insight. Along the way, it would be ill-judged to suppose that those involved will be, in a manner of speaking, flying by the seat of their pants. Enough of such matters is already known to support the idea that, in large measure, they will responding to what is already implicitly and tacitly known to them.

One of the oddities of Fred's situation is that although endlessly awash in incoming information, how notable the frequency with which his cognitive advancement is slowed, and sometimes impeded, by its lack. In dire circumstances, Fred could be the information processing counterpart of the ancient mariner, who woefully observes "Water, water everywhere, and not a drop to drink". Information flow is so capacious and relentless as to raise serious consideration of overload. The masses of it that reach us each millisecond achieves *some* sort of footfall. But it is not clear that Fred's information receptors are as open to incoming traffic as our causal response mechanisms are to causal stimulation. [75] No one doubts that our absorption of causal forces is largely unconscious and automatic. It is also clear that when causal forces arrive, to a considerable extent they dissipate upon absorption. On the other hand, however, is a significant fact about its susceptibility to storage and retrieval, its use and re-use. On the intuitive face of it, a great deal of incoming information is not even processed; it simply dissipates upon arrival. We know something rather remarkable about information processing in thermodynamically closed systems. Consider the sensorium, the juncture of the five senses. Its rate of information processing is approximately 11 million bits per second. However, when the possessor of the processing devices is conscious, fewer than forty bits make their way into consciousness. Consciousness therefore is a very large suppressor of information. Being conscious is a highly entropic and thermodynamic state to be in. Our sociolinguistic practice is more expensive still, dropping from about 40 to about 16 bits per second [100]. [76] It is true that Fred's information system can be hardly said to be thermodynamically closed. But it would be the sheerest folly to discount the close qualitative similarities between the size of the gaps between conscious and unconscious processing rates. Fred knows from having one how little room he has for storing things in the front of his mind and for that matter, in the back of his mind as well. Fred also knows from its constant and beneficial place in his life that his memory operates with such efficacy and timeliness as to invite comparison to an assembly line connected to a just-in-time supply line. He knows how little hand he has in directing the usages of memory. As is now starting to be clear, Fred owes much of his cognitive prosperity to devices beyond his control. It would not surprise Fred to learn that much of his cognitive life is lived in his cognitive down-below, unconsciously, automatically, inattentively, involuntarily, effortlessly, non-semantically, computation-luxurity, parallel-wise and deeply. Fred also has first-hand experience with his own cognitive up-above, in which the goings-on are conscious, intentional, attentive, voluntary effortful, semantically loaded, linguistically expressible, computationally weak, linear and shallow [101]. [77] Much of the time, of course, it is not one or all of these clusters—indeed, rarely all of them at once. But they offer the right checkpoints for limning the cognitive state Fred might be in at any given time. It would take some doing, and more time than is available here, to upgrade these remarks in light of the present states of information theory and consciousness studies. But there is more than enough at hand to make two points of particular importance for our present interest.

- *Talk is not cheap*: Speech, especially conversation, is an expensive facility to exercise. Still, although its draw on information is astonishingly low, its draw on material relevance is remarkably high [102]. [78]

At the same time,

- *Talk is epistemically freighted*: Most of what Fred will ever know in life depends on someone having told someone something, and in doing so, causing him to know it. Most of what we know is the end product of a multi-agent manifold of epistemically productive tellings. In cost-benefit terms, telling is the highest yielding of one of our most expensive assets.

For the knowing-by-being-told-it thesis to be true, Fred must have a preternatural disposition to speak truthfully. There are exceptions, of course, both on the telling and being told side. It would be ill-considered of Fred to tell a nervous colleague the truth when asked whether her new slacks made her look a bit plump. Equally, on the being-told side, Fred is not likely to believe a moral proposition just by virtue of having been told it. He is unlikely to show the same resistance to a scientific proposition. [79] There is no more routine pass in Fred's cognitive life than when the multi-agent telling-matrices tell him something he did not know before in the absence of which, withal, there is no ken of human life, indeed no human life at all.

It is an understatement that there is much yet to be learned of the energy-to-information conversions of energy-to-energy transductions. Fred Dretske is famous for the hypothesis that belief is information, in completely digitized form (1981; Gabbay and I have reported reservations about whether this accurately characterizes belief [46] (chapter 7, section 7)), but we do not doubt the importance of the distinction with respect to the cognitive up-above and the cognitive down-below. It moves me to suggest that a more useful application of the digital-analogue distinction is that conscious belief carries information re analog form, whereas subconscious belief carries the same information in digitized form. Assuming this to be so for the sake of argument, it raises the earlier question in a new form: How can the analog belief that S and the digitized belief that S be the same belief? And is there any well individuated single object that plays the role of S each time? The answer I think may lie in what is known of thermodynamics:

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

In physics, a phase transition is a reversible 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. There is a link between the phase transitions of modern physics and Aristotle's concept of potentiality (*dunamis*), by virtue of which a substance is able to take on a new form without losing its identity [37] (*Metaphysics* 8, 1 1046^a 12, 1048^a 25, 27). An item's haecceity is that in virtue of which it is the very thing it is and not another thing. Its quiddity is that in virtue of which a thing is the very kind of thing it is. The phase-transition thesis tells us that there are ranges of cases in which a thing's haecceity is unmolested by a change *and restoration* of quiddity [103,104]. [80] The economic savings learned by such transitions speak for themselves. Some readers might pause to wonder whether Peirce himself would have any truck with this. The answer is that he virtually invited it:

> "The soul [= mind] then *certainly* does act dynamically on matter. It does not follow that it acts directly on matter, because there may be involved an *endless* series of transformations of energy from the motion of one fluid to another, all these fluids being spiritual [= mental], followed by the *beginningless* series of

> transformations of energy in one fluid to anergy in another. All *these* fluids being material." (NEM, 3, 897)

We have, for most of the space available to this essay, focused on Fred and his kind. The reason we began with Fred lay in our interest in how much like him Frederika could be made to be in the manifestation of her intelligence. Why, then, do we pause now with Madam Justice Flanagan, our imaginary exerciser of juridical thinking? The answer is that in all essentials, she is in matters of basic intelligence just like Fred and the rest of us. Yet in her professional life, she is routinely engaged in managing things in ways that we have not yet in this essay managed to get to the bottom of. So we turn now to what What Actually Happens bids us to do. We go to the place where this very capacity is on its fullest behavioural display.

2. Madam Justice Flanagan [81]

Relevant Similarity

We begin with a light sketch of the epistemological significance of the *lex non scripta* doctrine that will do for now. In the application of precedents, we are not to take it that Courts of Appeal have been visited by mystical insight or any other source of miraculous divination. Judges are not prophets. Rather, while the doctrine ascribes to the judges a full cognitive command of what the law provides, it also provides that their knowledge of it is intrinsically *tacit* and *implicit*. The question to which this leads is whether tacit and implicit knowledge peculiar to High Court findings is more widely distributed over matters of non-legal standing in various different contexts of engagement. It would repay us to pause briefly with logic's first provisions for the tacit and implicit in Aristotle's treatment of enthymemes in *Rhetoric*. Although rhetoric is the art of persuasion and is not treated in any of the *Organon*'s six volumes, the importance of its treatment of the enthymeme is clear for what concerns us here. An enthymeme is an uncompleted syllogism, made so by the fact that one of its premisses is left unstated. In the classic example, the one-premiss argument ‹"All Greeks are mortal" "Socrates is mortal"› is converted to a proper two-premiss syllogism upon instatement of the unspoken premiss "Socrates is a Greek". We have in this the distinction we need. The tacity and implicity of missing premisses exemplifies the voluntarily *unsaid*, whereupon the tacity and implicity of a High Court precedent exemplifies the structurally *unsayable*. There is something decidedly odd about this, for we have not yet said anything about the conditions in virtue of which the structurally unsayable is so. If Fred is capable of tacit and implicit knowledge in the way an appellate court judge is, then there are things that we all know of which we cannot make a record or to which give voice without erasing it as knowledge. More concretely, is the wherewithal for tacit and implicit knowledge standard equipment for Fred and, if so, what are the conditions of its enablement and the degree to which its implementation is intelligent? The follow-up questions ask themselves: Is it possible for a computer to be a tacit and implicit knower in the sense in which a High Court judge is; that is to say, without the means of bringing what it knows to express articulation? If so, what do we ask the software engineer to do to bring about the creation of a machine of such design? And with these questions an earlier one recurs. How like Fred must a machine be to perform in all the ways in which Fred manifests his intelligence? I think that we will not arrive at a satisfactory answer to that question until we spend a bit more time with how Flanagan J performs her functions.

An even more central place for the implicit and tacit is judge-made law, as when an appellate court establishes the legal precedent that arises in the case at hand from the court's reasons for judgement (*ratio decidendi*). A precedent is a finding that is binding (*stare decisis*). It binds future cases whose material facts have a sufficiency of relevant similarity to the material facts of the original case. According to the *stare decisis* doctrine, decisions of a higher court are binding on all courts below and can have "persuasive authority" for sister courts domestically and to some extent also abroad. Courts at all levels should not disoblige their own prior decisions. A finding is binding when it rests on something which

a capable judge would be able without unnatural effort to construe as "a general legal principle" or "rule of law" [18]. [82] By and large, courts apply precedents with confidence, and normally that confidence is respected by future courts. Although a court's reasons for decision are carefully written and sometimes rival the page count of a Russian novel, the *precedent* is not itself expressly articulated. The law gives no express reason for this, but it cannot be said to be unforthcoming about the rule of law concomitant with a finding that makes new law or strikes down an old law:

> *The rule of judge-made law*: When a precedent arises from a juridical finding in a case framed around some particular facts, it is easy to state the rule of law that attends it. The rule says that in any future case whose facts bear a sufficiency of relevant similarity to the facts on which the original finding was based, then in the subsequent case the court must find as in the originating case.

Let us break this down to the constituent elements:

i. The facts material to the court's original finding.
ii. The finding in the present case.
iii. The rule of law as just stated.

So now we have it thus:

- The facts of the original cases do not travel; that is to say, they are not the facts material to future cases.
- The finding, however, does travel; that is to say, the finding of the first case retains re-application in the future case.
- The rule of law has standing general authority. In any subsequent case when new law, the finding must be not discomply with the prior finding.

These three facts constitute the whole content of the doctrine of precedents in English Law. This situation in question here is fully explicable and submittable to paper without having to call upon some fourth element. Suppose the original finding were that section so-and-so of the criminal code contravened section such-and-such of the constitution and, therefore, is struck down. The finding—because so-and-so section of the code is *ultra vires* [83] the such-and-such section of the constitution, that section of the criminal code is null and void—is certainly expressible. The rule of judge-made law is also expressible, and provisions for future cases are now adequately laid out. Accordingly, when we speak of a precedent having been created in an original case and of its having authority in future cases that meet the requisite conditions, we are speaking a sort of shorthand for the very set-up we have just been describing. When legal scholars say that the precedent cannot be written down, they are mistaken. One writes it down by writing down what it refers to, the threefold constitution of appellate practice. One cannot, however, write it down if it serves as a fourth element. The reason why is that there *is* no fourth element. The problem, as we can now see, lies elsewhere, albeit nearby. Here it is.

> *The real source of the non scripta problem:* The sense in which precedents are unspeakable and unwriteable is that the concepts embedded in the applicable rule of judge-made law—the concepts of sufficiency and similarity—are not subject to biconditional definition in law. The law rues the very idea of biconditional hook on words in common use. Wisely so.
>
> *Corollary*: This is not a peculiarity of the law. Most words, expressions and sentential utterances in the working vocabulary of the common man is not only understood in the absence of a definitions, but efforts to give the meanings of those items full-voiced biconditional articulation would, with a striking frequency, misdescribe them. [84]

This raises a nice little question of its own:

> *The how-possible question*: How is it possible for someone to know something without being able to say what it is that he knows? (St. Augustine on time.) Because it lies in the nature of the subject matter for some of its aspects to be

known to their information processors, both expert and otherwise, and to be known to the non-expert tacitly and implicitly.

3. Frederika

How Things Are Now

Before reflecting on what Frederika might be capable of doing or of being made to do, it would be a mistake to overlook what her present-day predecessors are and are not able to do. Of course, I mean here by "do or not do" do or not do by the exercise of intelligence. When we bear in mind that the disputed military installations of the Spratly islands in the South China Sea can be obliterated by pushing buttons in Tampa, we know that bellicose military engagement has metamorphosed into something it could not have been beforehand. When we learn that much of the more routine paralegal work in the offices of our solicitors is being done robotically, more efficiently and at lower cost, we might wonder why this is not reflected in the legal fees we are asked to pay. Automated facial recognition technology is so accurate and widespread as to have driven personal privacy to the point of extinction. Inter-corporate and inter-state security breaching is a matter of course and a rising and presently unstoppable peril. Some of this is comparatively old hat, but more recent is the AI penetration of the learned professions beyond the roboticizing of routine clerical work. Only this week (I write this in April 2022), the AI chief of one of Canada's six chartered banks has left to establish Signal 1 AI Inc., whose function is to ascertain when the hospitalized sick or injured are at greatest risk of dying or in need of immediate intervention. Previously, the head of the new company had sold his AI system Layer 6 to the Toronto-Dominion Bank and joined as the bank's chief AI officer. Layer 6 had many functions, including predicting when the bank's customers might be interested in buying a house and detecting fraud in insurance claims. A serious problem for large institutions is what I call chronic *data-stray*. A datum or set thereof has strayed from a database when its relevance for data in other bases goes unnoticed or, more generally, when its significance is overlooked. In actual institutional settings, the larger the database, and the more varied, the greater likelihood of stray. When the incidence of stray is high, the likelihood of database coherence is low. Any big hospital's information-system is sluggish with data-stray, whose occlusions are a net loss for good health-outcomes. So the principal task of Signal 1 AI is to winnow out data-stray from the hospital's information system. [85] On a less happy note, also a Canadian example, we have the Phoenix payroll system for federal government employees provided by IBM using PeopleSoft software and operated by Public Services and Procurements, the federal ministry in charge. Initial funding began in 2009 and the system launched in 2010, and its implementation passed to the ministry in 2016. Phoenix was a failure at launch and has resisted repair with astonishing intractability. By the summer of 2018, Phoenix had mismanaged the payrolls of close to 80% of the government's 290,000 public employees, either through under-payment, over-payment, misdeposit or no payment at all. While the company was initially designed to produce annual savings of $70 million, it is estimated that the cost of Phoenix's still-unresolved snafu will reach $2.2 billion. [86] There is in this *débâcle* a lesson as old as the hills. Even the best-laid plans go awry when regulation and control pass to lazy and stupid civil servants and nervous impulsiveness of second-rate grandstanding politicians. Of course, it is not always this way.

In what we have covered so far, we have learned many valuable things about Fred, which is to say about us all. We have learned of our unseverable tether to the natural world and to one another. We have been reminded of our ceaseless exposure to the habitat's causal wash and the causal flows within that carry information. As we now see, we are not only knowledge-seeking beings, we are knowledge-dependent beings, beings whose drive for knowing things rivals the need for air. Of very considerable importance is the insight that we individuals are not, and could not be, on our cognitive own. We are all parties to various of our life's cognitive economies, as we are to the cognitive economy that is common to us all. A condition of admittance and continued membership is a powerful disposition to

tell things to people, underlain by a like desire to be right in what we say. Should those impulses not have arisen or were they to abate, there would be no told knowledge, in whose absence the cognitive economy would crash into oblivion. For this not to have happened, the evolution of language was a foregone necessity. That in a nutshell is the story of Fred. He is *homo economicus* of biosocial embodiment.

In all the respects in which Fred is the way he is, we could ask whether there are ways in which Frederika might be made to be or could. If anything is clear about such matters it is that, up to now, the naysayers have been mainly wrong and the software engineers mainly right. All the same, if they are to be right about the Deep Learning third wave, there are some questions that still await conclusive answer.

> *The enablement-manifestation question:* When a being operates or responds in a way that manifests intelligence, must the intelligence in question inherit the peculiarities of its means of enablement?

As has been said, one of the chief characteristics of a cognitive system is the success it has achieved in liberating itself from the cognitive setback imposed by the limitations on its wherewithal for knowledge of its participating agents. In mature cognitive economies, the conversion of disadvantage to advantage has a dominant place in the advancement of cognitive agendas. The classic case is the conversion of the setback imposed by ignorance to cognitive rebound by the *jiu-jitsu* mechanism of abductive reasoning. It is a process in which a limitation imposed by ignorantae is compensated to net advantage by such uses of one's head as to occasion contact with a hypothesis which not only solves the abducer's ignorance-problem but also in ways that spur greater advancement in science and policy formation. If, as averred by Peirce and strongly suggested by What Actually Happens, it is true in these cases that it matters utterly for scientific prosperity, the hypothesis has to be the truth-making fruit of the hypothesizer's own creative making-up. In mathematical circles, especially in its higher-up precincts, creative out-of-the-box truth-making sometimes raises the whole domain of enquiry to new heights. Figures as differrent from one another as Frege, Peirce, Dedekind and Hilbert were all at one mind about this.

This leaves us critical matters to consider in relation to Frederika. Can we build into Frederika the *instinct* to fasten onto the right hypothesis? Earlier, we likened the making of mathematical truth to the making of fictional truth. We recommend retaining the creative-making aspect and disposing of the fictional aspect. In each case, the enablement conditions lie in the semantic conventions of human speech and thought. Their provisions are indeed quite remarkable. In the case of Sherlock, they endow the world with truths about Sherlock, and do so in a way that what is true of him is a real-world truth. Although similar, our semantic conventions endow the world with truths about mathematical objects, and do so in a way that what is true of them are real-world truths. What matters at bottom in each case is that the truth-makings of Sherlock by the scribbling of Sir Arthur Conan Doyle are in turn enabled by the real-world facts about how fictional truths come to pass. Equally, the mathematical truth-makings of Riemann and Cantor could not have come to be save for the real-world facts of human mathematical creativity. Again, the common factor is that intentionality is not blocked by impalpability, and ultimately what is true of Sherlock and true of the hyperreals is underwritten by the language-using realities of the world itself. Could the likes of Frederika join the ranks of Riemann, Frege, Peirce, Dedekind and Hilbert, and by exercising her instinct for hypothesis selection in ways that bring to pass large advances in our scientific understanding? Suppose that we conceded that, given the sheer size of her knowledge base, Frederika evades the seduction and necessity of abductive thinking. Could she, even so, do what Fred in his own abductive circumstances can do to such overall cognitive advantage, independently of the *wherewithal* for abductive assuagement? In some of the writings listed above, the key question is whether Frederika is capable of judgement. In both its common and legal meaning, judgement is the rational reconciliation of conflicting interests or some other way in which their presence might be reconciled to the economy's wide reflective equilibrium. Given the promise and the triagic nature of Signal 1. AI's doing its thing at St. Michael's Hospital, I can only think that the

answer is in the affirmative. And since the system's capacity for improved quality of care outcomes does not depend on enabling conditions that match the nature and quality of the predecessor arrangements, I can only say that it is not in general the case that anything of Fred's doing that Frederika can do better depends on enabling conditions that match Fred's own.

It is time to bring these reflections to a close, beginning with a review of some of the pending questions:

- Can Frederika go it alone in cognitive life independently of her fellow counterparts?
- Is Frederika susceptible to intuitively subpar performance for the good of cost-benefit compensatory payoff?
- Is Frederika bound by conditions under which TI-knowledge is the best outcome?
- In what does the larger virtuosity of intelligence inhere? In the skills possessed by the agent or in the agent's natural resources?
- Judged in one way, Frederika would match or better the outcomes redounding to Fred's application of his skills to the resources at his command.
- Judged in the other, the outcomes achieved by Fred would have the greater intelligence quotient.
- So is there any overall benefit to be achieved by Frederika by building into her the productive limitations of Fred and his *jiu-jitsu* manoeuvres—his cognitive *sweeps* so to speak?
- All these are questions of sufficient interest and importance, but the truly core question is: can Frederika be built with the instinct for truth-creation about beings impalpable to us and also to her?

My answer is that the question must be tabled until more is known of how Fred himself is able to do this. [87] A last word, then. One of the most cherished products of cognitive economies like Fred's are its creative measures for making new truths in mathematics. A question for Frederika is whether she can bring to pass the likes of the following:

- When Conan Doyle thinks up Sherlock and writes down things about him, he makes Sherlock, and makes things true of him, as a matter of real-world fact. Sherlock is a man, a real man, a real man of fiction, just as π is an actual number of mathematics.
- When an out-of-the-box mathematician is seized of a new idea and incorporates it into a hypothesis, then given the ways of the world, the hypothesis, if well-abduced, is gravid with young truth.
- This is fully half the semantic reality of mathematical truth. The other half is provided by the conditions, should they obtain, by which a young truth matures into the real thing.
- For the rites of passage to prevail, the young truth will have shown a steady theoretical fruitfulness [105], [88] but will also have withstood all the pressures against it by the free market of the dialectical structure of expert opinion.
- When all opposition subsides, we are left with an abduction problem. What hypothesis, if true, would account for such placement in a mathematical economy in wide reflective equilibrium? Since there is in such cases no question of testing the matter by direct observation of the real truth of the matter, what would Fred recommend as the answer?
- And could Frederika match this or better it?

Funding: This research required no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Acknowledgments: For support and adroit editing, as well as the example set by his own scientific work, I warmly thank this issue's SI editor, Woosuk Park. For criticisms and constructive suggestions in preparing this essay, in addition to Woosuk Park and his referees, I am indebted to Selene Arfini, Richard Atkin, Katie Atkinson, Jean-Yves Beziau, Lorenzo Magnani, Peter McBurney, Ahti-Veikko Pietarinen and Alirio Rosales. Dedication: I have the honour of dedicating these reflections to an old friend and gifted pioneer: Arthur W. Burks (1915–2008).

Conflicts of Interest: The author declares no conflict of interest.

Notes

1 The correct spelling is "unwritable", as with "undecidable". I misspell it on purpose to bring the inscription more in line with the spoken word.

2 While the usual expression is "cognitive agent", it suggests agency when there often is none, and overlooks the frequency with which knowledge just happens to one. You check your inbox and find: "Nancy just called. Could you give her a ring at her office when you get a chance?" Now you know that Nancy called, and you did not lift a finger to know it. All you needed were some functioning eyes and a basic command of English. The more accurate alternative to "cognitive agent" would be "cognant". Still, since "cognant" is unlikely to have much take-up in the research community, I'll go with the flow and stick with "cognitive agent", and leave it to context to determine when agency is involved.

3 My ascription of intelligence to Fred is without prejudice to the disputed questions of nonhuman and nonlinguistic cognition. See here Bermúdez [2], and "Animal reasoning and proto-logic" in Bermúdez [3] (pp. 127–137). For a good examination of future prospects, readers are invited to consult Park [4] (pp. 1–32). My own position can be put this way: Spend a fortnight with a social ecologist examining the activities at a bumblebee colony. It will soon be noticeable how causal forces dominate bumblebee economies. See here Heinrich [5].

4 In its most widely circulated form, here is Peirce's schematization: "The surprising fact *C* is observed. But if *A* were true, *C* would be a matter of course./Hence there is reason to suspect that *A* is true." [1] (5. 189). The line-breaks are mine. The schema gives a sample of Peirce's thinking, but is not to be considered the final word.

5 In routine diagnostic settings, symptoms present themselves and are already known to be of this or that or some other malady in the general case. But, as in law, so in medicine, facts of a given case often differ from facts of like kind in the general case. The abducing physician must select this or that or the other, or do some fresh thinking.

6 These days, it is impossible to be unmindful of the large pockets of resistance to these ways of speaking. If amends could be made by speaking instead of the woman in the street or the common woman, I would contemplate adopting them. But since they do not work, I shall not adopt them, and will not deign to feign acquiescence to the imagined singularity of the plural pronoun "they".

7 "Abduction is the process of forming an explanatory hypothesis. It is the only logical operation which introduces any new idea; for induction does nothing but determine a value, and deduction merely evolves the necessary consequences of a pure hypothesis." [1] (5. 171)

8 The point is argued in my "Abduction and inference to the best explanation", in [7].

9 Note the resemblance to certain readings of verificationist criterion of meaning.

10 Inference is "the conscious and controlled adoption of a belief as a consequence of other knowledge." [1] (2. 442)

11 Note that these considerations cleanse Peirce's understanding of guessing of any connotation of shots in the dark. Peirce's guesses are educated guesses.

12 For consciousness without attention, see Mole [8].

13 Details can be found in Woods [9].

14 Details can be found in [9].

15 See [1] (2. 58).

16 Burks [14]; Pople [12]; Poole et al. [13]; and Kakas et al. [11].

17 See [15].

18 In the jurisdictions of my acquaintance, it is customary to refer to members of the appellate bench as "justices" rather than "judges". We can be more flexible in our usage here. It is customary to refer to a high court judge as, for example, Madam Justice A. T. Flanagan or, more briefly, as Flanagan J.

19 On the strongly Nay-side, among others, is Cantwell Smith [17]. Similar but not quite equivalent views can be found in Clark, [16].

20 Lower court decisions from other common law jurisdictions can have "persuasive authority" for sister courts at home. Such decisions are not binding. Also, in some jurisdictions, for example Canada and the United States, the Supreme Court can sometimes find a way to reverse some of its own findings. This is a constitutionally tricky thing to do, since the Supreme Courts are the tribunals of last resort.

21 Considerable credit is due computer scientists for having spotted the logico-epistemic importance of the administration of justice. It is a matter of regret that the same cannot be said for epistemologists or logicians. To this day in English-speaking countries,

university courses on the philosophy of law are wholly absorbed with issues in ethics as well as social and political theory. Journals of logico-epistemic note include *Artificial Intelligence and Law and JURIX*.

22 Woods [18].

23 Hewitt [19], and Hewitt and Woods [20].

24 The meaning here is that from any logically false proposition, every statement of the language in which the inconsistency arose is thereby logically entailed.

25 Think here of the Erdös discrepancy problem, some attempted proofs of which require more gigabytes than that of Wikipedia in its entirety.

26 Classical sources include Goldman [22], Quine [23], Kripke [24] and Benacerraf [21].

27 Morris [25].

28 Lewis [26], Austin [27] and Grice [28].

29 See, for example, Priest et al., [30]; Batens et al. [29]; Priest [31]; and Schotch et al. [32].

30 See Frege [33], and Frege [34]. The difference between the paraconsistent and Fregean response to inconsistency lies in the fact that they were responding to different things. For paraconsistents, the *bête noir* was not a given inconsistency but rather what any given inconsistency logically entails, namely the inconsistency of every statement in the system, proved by the theorem called *ex falso quodlibet* (from a [logical] falsehood everything follows). So any system thus afflicted simply blows up ("explodes", "detonates"). *Ex falso* is derivable in Frege's system, but this is damage to a system that had already collapsed. The trouble was local, Frege's comprehension axiom for sets (value ranges of functions). Under Russell's prodding, Frege was able to prove that his axiom entailed a contradiction. With the axiom now disabled, every theory of the system whose proof depends on that axiom is also disabled, as in turn any further theorem depending on those theorems, and so on. In short, the system simply falls down. It does not blow up.

31 Hewitt [20]. See also Woods [35].

32 Let us keep in mind that it was Tarski, after all, in his classic and ground-breaking paper of 1936 on logical consequence, who insisted that when adapted to the special needs of the metamathematics of classes, the results must preserve as much of this meaning as possible. The sole exception was the theory's use of the uniform substitutivity of nonlogical terms for other nonlogical terms of the grammatically admissible type. See here, in excellent translation, Alfred Tarski, "On the concept of following logically" [36] translated by Magda Stroińska and David Hitchcock, with an introduction by Hitchcock.

33 Making such a logic is the primary object of Woods [9].

34 Developments in computational abduction include Poole et al. [13] and Kakas et al. [11].

35 For example, von Neumann computers cannot handle higher order logics, but Vladimir Lifschitz has pointed out that McCarthy's Circumscription can be computed for a class of separable formulas which contains only quantifier-free formulas. See, for example, Gershan et al. [40], Adams et al. [38], Clark [16], and Park [42]. An important exception is Bruza and Gibson [43]. My own view is that Fred's cognition is the unmediated causal outcome of largely subconscious information processing devices. See here Busemeyer et al. [39], and Kirkhof and Robertson [41].

36 Representationist theories are thick on the ground and often at significant variance with one another. See, for example [16,38,40,42]. An important alternative is [43]. My own view is that Fred's cognition is the unmediated causal outcome of largely subconscious information processing devices. See here [39,41].

37 As its students will know, as the Second World War was coming to an end in Europe, and Allied soldiers were more widely dispersed than ever, there would appear on walls and fences without apparent surcease the drawn figure of two hands clutching the top of the structure, betwixt a pair of large eyes and overhanging large beak, and just below the inscription "Kilroy was here". No one reading those words with understanding will not have known that the string inscribed under the cartoon asserted the recent presence of someone called "Kilroy".

38 A good account of how this came about is Everett [44].

39 In announcing "l'état, c'est moi", Louis XIV was more venturesome than accurate. Louis XVI was spot on with the dire prediction, "Après moi, le déluge".

40 There are notable resemblances between Peirce and Schopenhauer, author of the famous *The World as Will and Representation*. One day during a pleasant Ann Arbor lunch, Arthur Burks quipped to the present author that Peirce's semiotics could be entitled *The World as Willful Misrepresentation.*

41 Gabbay and Woods [46], Woods [48], and Morton [47].

42 Haugeland and Smith are two exceptions. See Haugeland [49] and Cantwell Smith [17], Clark [16], and Thagard [50].

43 Perconte and Plebe [52], Hinton et al. [51], and Rumelhart et al. [53]

44 In the late 1990s, during a visit from the University of Toronto to Dov Gabbay's Logic and Computation Group at King's College London, Ray Reiter, the mathematician who founded the computational logic of default reasoning in 1980, mentioned to the two of us that he, Reiter, would not then make the long-list for appointment to a computer science department even at a second-tier university.

45 A signature theme of [46] and [48]. See also Woods [56]. A classic paper is Harman [54]; see also Harman [55].

46 Cobb [57].

47 I have drawn these quotations from Mole [58] in a review of Cobb [57].

48 There were exceptions of course. See, for example, Fukushima [59].

49 Automated taxi-drivers are a bit trickier to sort out. Although here, too, there were some early successes, as witness the VaMoRs van of automatic mobiling based on the work of Ernst Dickmanns. As anyone familiar with London's black cab fleet, drivers are held to a very high expectations of cartographic memory. There are also the sundry obligations of law and social convention to observe. But no London cabbie will prosper if he does not master the art of breaking the traffic laws systematically, safely and without too much notice. In our plans for Frederika, is it contemplated that she would be thus accoutered? For more on taxi-driving robots, see Gabbay et al. [60].

50 See also for example, Doya et al. [61], Hohwy [62], and Hooker [63].

51 Magnani [64–66].

52 In this paragraph and the next I draw upon Woods [67].

53 Vitiello [70], Bruza et al. [68], and Bruza and Woods [69].

54 Goodman's solution has been nicely adapted by William Lycan for enquiry into the management of philosophical disputes. See here Goodman [71], Rawls [73] and Lycan [72].

55 Simon [8].

56 From the nursery rhyme "Hey diddle-diddle, the cat and the fiddle", and not really a fable at all.

57 With a respectful nod to *Hamlet*, Act 3, scene 2.

58 For more on the importance of commonness for Fred's cognitive well-being, readers could consult Woods [74].

59 I hope that it needs no saying that no such indictment is intended of Alexander Grothendiek, author of the splendid [75]. In a way, it is regrettable that "stupid" also means unintelligent, not-smart or, as in the vulgate, dumb. The usage proposed here bears some likeness to "foolish" which, in some contexts, carries connotations of the comic or amusing. "Unwise" is best here, meaning "of unsound judgement".

60 See again Hohwy [62] and Clark [16].

61 *Notre Dame Philosophical Reviews*, 2017. 01. 15, 1–10.

62 Helmholtz [77].

63 Not to be overlooked is the threefold ambiguity of "fallibilism". In its epistemic sense, it is the doctrine according to which it is possible that some given proposition, which presently I experience myself as knowing, I actually do not. In its logically independent alethic sense, it is possible that the thing I presently experience myself as knowing is actually false. (Note that instantiation of sense one does not entail the instantiation of sense, but the entailment does hold in reverse. In its synechist sense endorsed by Peirce, fallibilism is the doctrine that no proposition of whatever modality currently known to be true will stay true no matter what.)

64 Woods [80], Arfini [78], Arfini and Bertolotti [79].

65 A classic work is Cohen [81]. In some writings "assent" is a third meaning. In a fourth sense, belief is "commitment", concerning which see Walton and Krabbe [82], and for a different perspective Woods [83]. As the book was ready for production, Hermes was shutting down its English-language operations, and it took some doing to get the book published. When the time came for a second and revised edition with College Publications in 2014, I unwisely omitted this chapter, the first edition being long out of print and virtually impossible to find. Time permitting, I will refresh the chapter on induction and republish it on its own.

66 Woods [84].

67 Floridi [90]. Perhaps the best single-volume treatments of the current state and future prospects of the philosophy of information are Adriaans and van Benthem [85] and Floridi [89]. Also important is Dretske [88]. Older but still in fine fettle is Cover and Thomas [87]. For early applications to psychology, see Attneave [86].

68 See, for example, Bar-Hillel [91], Dretske [88,92] and Floridi [90].

69 The classical works are Hartley [93], Shannon [94] and Shannon and Weaver [15]. See also Cover and Thomas [87].

70 Adriaans and van Benthem [85], Chaitin [95], Li and Itánge [96] and Solomonoff [97].

71 Adriaans and van Benthem [85] (p. 12).

72 See here Adriaans and van Benthem [85] (p. 7).

73 Hintikka [98].

74 Jaakko Hintikka [99] "What is abduction?" The fundamental problem of contemporary epistemology", *Transactions of the Charles Peirce Society*, 34 (1998), 503–533. Reissued with revisions in *Socratic Epistemology* (2007), as "Abduction—inference, conjecture, or answer", pages 38–60.

75 Note here the resemblance to the endless wash of entailment-inconsistency, and the dab hand we are in making highly selective and hugely productive uses of it.

[76] Zimmerman [100].

[77] See, for example, Schiffrin [101].

[78] In "Four-grades of ignorance-involvement and how they nourish the cognitive economy" [102], I launched the idea of *anselmian* knowledge, St. Anselm is well-known for his declaration "Credo, ut illelligam", which I translate as saying that he believes in order that he might be assisting in coming to understand. As striking as the incidence of told-knowledge is Fred's cognitive economy; not far behind is the incidence of Anselmian knowledge. Consider what most people know about the constitution of energy. It is conveyed by the sentence "$E = mc^2$". Some people have a deep understanding of what the equation provides. But what the man in the street knows of it is that it expresses a law of physics about energy and mass. People who know just that of it are not in the same boat as someone who has no idea of what it says. Two things, beyond its great frequency, stand out about Anselmian sentences. One is their semantic lightness and, with it, the propensity to travel cheaply. The other is that in the passage from the Anselmian to be more fully grasped, the way is eased by what is implicitly and tacitly known. It also travels cheaply.

[79] Excluding those that arise from partisan-ridden branches of enquiry or, worse, sectors of enquiry routinely hobbled by government agencies that cook the books and issue lying reports of the known facts.

[80] For physics see Callen [103], and for kinds see Pelletier [104].

[81] Flanagan J is a creature of imaginary convenience, standing in for the bench of which she is a member. Justice Flanagan won a substantial reputation as corporate crimes trial counsel, followed by a six-year tenure as provincial court judge. She is currently the Chief Justice of her court. Here initials stand for "Anne Theresa", and she is known as Theresa.

[82] The doctrine of *stare decisis* is under current duress in Canadian courts by activist judges in courts *below*. See here Woods [18], appendix H.

[83] Beyond the intention of the legislators (roughly speaking).

[84] Think here of the classical definition of deductive implication: Propositions $S_1, \ldots, S_n$ deductively imply proposition S′ if and only if it is in no sense possible for the S_i to be jointly true and S′ concurrently not. When read from left to right, the definition is met with widespread communal approval. But when read from right to left, the approval rate goes down; and if attention is called to *ex falso*, approval rates tumble. *Ex falso* is the doctrine that a contradiction deductively implies every proposition whatever. This does not, of course, falsify the definition. But it does suggest that when people make the correct consequence-attributions, they are implementing the left to right reading.

[85] I have drawn these remarks from the Report on Business section of Toronto's *Globe and Mail*, 25 April 2022, section B, pp. 1 and 8.

[86] With a population one-tenth the size of the American population, this represents a cost in US terms of ≈ 220 USD p/a. It is not chump-change, and the thing does not work.

[87] More of this is on view in the closing chapter of Woods [9].

[88] See, for example, Tappenden [105].

References

1. Peirce, C.S. *Collected Papers*; Hartshorne, C., Weiss, P., Burks, A.W., Eds.; Harvard University Press: Cambridge, MA, USA, 1958.
2. Bermúdez, J.L. *Thinking Without Words*; Oxford University Press: New York, NY, USA, 2003.
3. Bermúdez, J.L. Animal reasoning and proto-logic. In *Rational Animals?* Hurley, S., Nudds, M., Eds.; Oxford University Press: Oxford, UK, 2006; pp. 127–137.
4. Park, W. What proto-logic could not be. *Axiomathes* **2021**, 1–32. [CrossRef]
5. Heirich, B. *Bumblebee Economics*; Harvard University Press: Cambridge, MA, USA, 1979.
6. Peirce, C.S. *Philosophy of Mathematics: Selected Writings of Charles Sanders Peirce*; Moore, M.E., Ed.; Indiana University Press: Bloomington, IN, USA, 2010.
7. Woods, J. Abduction and inference to the best explanation. In *The Critical Thinking Book*; Blair, J.A., Ed.; WSIA: Windsor, ON, USA, 2018; pp. 405–430.
8. Mole, C. Attention and consciousness. *Conscious. Stud.* **2008**, *15*, 86–104.
9. Woods, J. *The Three Dimensions of Deductive Involvement and How They Matter for Proof Theory*; forthcoming in 2023; Department of Philosophy, University of British Columbia: Vancouver, Canada.
10. Peirce, C.S. *The Charles S. Peirce Papers*; Harvard University Photographic Service: Cambridge, MA, USA, 1966.
11. Kakas, A.C.; Kowlaski, R.A.; Toni, F. Abductive logic programming. *J. Log. Comput.* **1992**, *2*, 719–770. [CrossRef]
12. Pople, H.E., Jr. On the methanization of abductive logic. In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-73), Stanford, CA, USA, 20–23 August 1973; pp. 147–152.
13. Poole, D.; Goebel, R.; Aleixunas, R. Theorist: A logical reasoning system for defaults and diagnosis. In *The Knowledge Frontier: Essays in the Representation of Knowledge*; Cercone, N.J., McCalla, G., Eds.; Springer: Berlin/Heidelberg, Germany, 1987; pp. 331–353.
14. Burks, A.W. Peirce's theory of abduction. *Philos.Sci.* **1946**, *13*, 301–306. [CrossRef]
15. Shannon, C.E.; Weaver, W. *The Mathematical Theory of Communication*; University of Illinois Press: Urbana, IL, USA, 1949.
16. Clark, A. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*; Oxford University Press: Oxford, UK, 2016.
17. Cantwell Smith, B. *The Promise of Artificial Intelligence: Reckoning and Judgment*; MIT Press: Cambridge, MA, USA, 2019.

18. Woods, J. *Is Legal Reasoning Irrational? An Introduction to the Epistemology of Law*; College Publications: London, UK, 2018.
19. Hewitt, C. Formalizing common sense reasoning for scalable inconsistency-robust information coordination using Direct Logic Reasoning and the Actor Model. In *Inconsistency Robustness*; College Publications: Rickmansworth, UK, 2015; pp. 3–103.
20. Hewitt, C.; Woods, J. *Inconsistency Robustness*; Studies in Logic; College Publications: London, UK, 2015; Volume 52.
21. Benaceraff, P. Mathematical truth. *J. Philos.* **1973**, *70*, 661–679. [CrossRef]
22. Goldman, A.I. A causal theory of knowing. *J. Philos.* **1967**, *64*, 357–372. [CrossRef]
23. Quine, W.V. *Epistemology naturalized*; (pre-printed in Quine, Relativity and Other Essays); Columbia University Press: New York, NY, USA, 1969; pp. 69–90.
24. Kripke, S. *Naming and Necessity*; Harvard University Press: Cambridge, MA, USA, 1982.
25. Morris, C.W. *Writings on the General Theory of Signs*; Mouton: Hague, Netherlands, 1971.
26. Lewis, D.K. *Convention: A Philosophical Study*; Harvard University Press: Cambridge, MA, USA, 1969.
27. Austin, J.L. *How to do Things With Words*; Oxford University Press: Oxford, UK, 1961.
28. Grice, H.P. *Studies in the Way of Words*; Harvard University Press: Cambridge, MA, USA, 1989.
29. Batens, D.; van Bendegem, J.P.; Priest, G. *Frontiers of Paraconsistency*; Research Studies Press: Baldock, UK, 2000.
30. Priest, G.; Routley, R.; Norman, J. *Paraconsistent Logic: Essays on the Inconsistent*; Philosophia: Munich, Germany, 1989.
31. Priest, G. Paraconsistency and dialetheism. In *The Many Valued and Nonmonotonic Turn in Logic*; Gabbay, D.M., Woods, J., Eds.; Elsevier: Amsterdam, The Netherlands, 2007; pp. 129–204.
32. Schotch, P.; Brown, B.; Jennings, R. *On preserving: Essays on Preservationism and Paraconsistent Logic*; University of Toronto Press: Toronto, ON, Canada, 2009.
33. Frege, G. Letter to Russell. In *From Frege to Gödel. A Source Book in Mathematical Logic, 1879–1931*; van Heijenoort, J., Ed.; Harvard University Press: Cambridge, MA, USA, 1902; pp. 127–128.
34. Frege, G. *Philosophical and Mathematical Correspondence*; Basil Blackwell: Oxford, UK, 1980; pp. 130–170.
35. Woods, J. How paradox fares in inconsistency-robust logic and beyond: Computational and naturalized approaches. *J. Appl. Log.* **2021**, *8*, 2119–2172.
36. Tarski, A. On the concept of following logically (trans. Magda Stroińska and David Hitchcock). *Hist. Philos.Log.* **2002**, *23*, 155–196. [CrossRef]
37. Aristotle. *The Complete Works of Aristotle. The Revised English Translation*; Barnes, J., Ed.; Princeton University Press: Princeton, NJ, USA, 1984.
38. Adams, R.; Shipp, S.; Friston, K. Predictions, not commands: Active infeeds in the motor system. *Brain Struct. Funct.* **2013**, *218*, 611–643. [CrossRef]
39. Busemeyer, J.R.; Bruza, P.D. *Quantum Cognition and Decision*; Cambridge University Press: Cambridge, UK, 2012.
40. Gershman, S.; Vul, E.; Tanenbaum, J. Multistability and perceptive inference. *Neural Comput.* **2012**, *24*, 1–24. [CrossRef] [PubMed]
41. Kirkhof, M.D.; Robertson, I. Enactivism and predictive Processing: A nonrepresentational view. *Philos. Explor.* **2018**, *21*, 264–281. [CrossRef]
42. Park, W. *Abduction in Context: The Conjectural Dynamics of Reasoning*; Springer: Cham, Switzerland, 2019.
43. Bruza, P.D.; Gibson, A. Abduction beyond representation. In *Handbook of Abductive Cognition*; Magnani, L., Ed.; Springer: Berlin/Heidelberg, Germany, (Forthcoming).
44. Everett, D.L. *How Language Began: The Story of Humanity's Greatest Invention*; Liveright Publishing: New York, NY, USA, 2017.
45. Russell, B. *The Principles of Mathematics*; Cambridge University Press: Cambridge, MA, USA, 1903.
46. Gabbay, D.M.; Woods, J. *Agenda Relevance: A Study in Formal Pragmatics*; Elsevier: Amsterdam, The Netherlands, 2003.
47. Morton, A. *Bounded Thinking: Epistemic Virtues for Limited Agents*; Oxford University Press: New York, NY, USA, 2013.
48. Woods, J. *Errors of Reasoning: Naturalizing the Logic of Inference*; Studies in Logic; College Publications: London, UK, 2013; Volume 45.
49. Haugeland, J. *Artificial Intelligence: The Very Idea*; MIT Press: Cambridge, MA, USA, 1985.
50. Thagard, P. *Bots and Beasts: What Makes Machines, Animals and People Smart?* MIT Press: Cambridge, MA, USA, 2021.
51. Hinton, G.E.; McClelland, J.L.; Rumelhart, D.E. Distributed representations. In *Parallel Distributed Processing: Exploring the Microstructure of Cognition*; Rumelhart, D.E., McClelland, J.L., Eds.; MIT Press: Cambridge, MA, USA, 1986; pp. 77–109.
52. Perconti, P.; Plebe, A. Deep learning and cognitive science. *Cognition* **2020**, *203*, 104365. [CrossRef] [PubMed]
53. Rumelhart, D.E.; McClelland, J.L.; PDP Research Group. *Parallel Distributed Processing: Explorations of the Microstructure of Cognition*; MIT Press: Cambridge, MA, USA, 1986; Volume 2.
54. Harman, G. Induction: A discussion of the relevance of Knowledge to the theory of induction. In *Induction, Acceptance and Rational Belief*; Swain, M., Ed.; Reidel: Dordrecht, The Netherlands, 1970; pp. 83–99.
55. Harman, G. *Change in View: Principles of Reasoning*; MIT Press: Cambridge, MA, USA, 1986.
56. Woods, J. Epistemology mathematicised. *Informal Log.* **2013**, *33*, 292–331. [CrossRef]
57. Cobb, M. *The Idea of the Brain: The Past and Future of Neuroscience*; Basic Books: London, UK, 2020.
58. Mole, C. Mind the gaps. *Times Lit. Suppl.* **2021**, *27*, 6153.
59. Fukushima, K. Neocognition: A self-organizing neural network model for a mechanism of pattern recognition unaffected by a shift in position. *Biol. Cybern.* **1980**, *36*, 93–202. [CrossRef] [PubMed]

60. Gabbay, D.; Cramer, M.; Dauphin, J.; Farjami, A.; Rivlin, L.; van der Toore, L. Machine argumentation: Can we replace taxi drivers by robots? In *Natural Arguments*; Gabbay, D., Magnani, L., Park, W., Pietarinen, A.-V., Eds.; Tributes series; College Publications: London, UK, 2019; Volume 40, pp. 177–194.
61. Doya, K.; Ishii, S.; Pouget, A.; Rajesh Rao, P.N. *Bayesian Brain: Probabilistic Approaches to Neural Coding*; MIT Press: Cambridge, MA, USA, 2007.
62. Hohwy, J. *The Predictive Mind*; Oxford University Press: New York, NY, USA, 2013.
63. Hooker, C.A. *Handbook of the Philosophy of Science*; Elsevier: Amsterdam, The Netherlands, 2011; Volume 10.
64. Magnani, L. *Abduction, Reason and Science: Processes of Discovery and Explanation*; Kluwer: Dordrecht, The Netherlands; Plenum: New York, NY, USA, 2001.
65. Magnani, L. *Abductive Cognition: The Epistemological and Eco-Cognitive Dimensions of Hypothetical Reasoning*; Springer: Berlin/Heidelberg, Germany, 2009.
66. Magnani, L. Human abductive cognition vindicated: Computational locked strategies, dissipative brains, and eco-cognitive openness. *Philosophies* **2022**, *7*, 15. [CrossRef]
67. Woods, J. How abduction fares in mathematical space. In *Handbook of Abductive Cognition*; Magnani, L., Ed.; Springer: Cham, Switzerland, 2022.
68. Bruza, P.; Widdows, D.; Woods, J. A quantum logic of down below. In *Handbook of Quantum Logic and Quantum Structures*; Engesser, K., Gabbay, D., Lehman, D.L., Eds.; Elsevier: Amsterdam, The Netherlands, 2009; Volume 2, pp. 625–660.
69. Bruza, P.; Woods, J. Quantum collapse in semantic space: Interpreting natural language argumentation. In *Proceedings of Quantum Integration 09*; College Publications: London, UK, 2009; pp. 59–80.
70. Vitiello, G. The dissipative brain. In *Brain and Being: At the Boundary Between Science, Philosophy, Language and Arts*; Globus, G., Pribram, K.H., Vitello, G., Eds.; John Benjamins: Amsterdam, The Netherlands, 2012; pp. 317–338.
71. Goodman, N. The new riddle of induction. In *Fact, Fiction and Forecast*; Harvard University Press: Cambridge, MA, USA, 1953; pp. 59–83.
72. Lycan, W.G. *On Evidence in Philosophy*; Oxford University Press: New York, NY, USA, 2019.
73. Rawls, J. *A Theory of Justice*; Harvard University Press: Cambridge, MA, USA, 1971.
74. Woods, J. The role of the common in cognitive prosperity: Our command of the unspeakable and the unwriteable. *Log. Univers.* **2021**, *15*, 399–443. [CrossRef]
75. Grothendiek, A. *Topological Vector Subspaces*; Gordon and Breach: New York, NY, USA, 1973.
76. Rescorla, M. Clark's Surfing Uncertainty. *Notre Dame Philos. Rev.* **2017**, *15*, 1–10.
77. Von Helmholtz, H. *Handbuch der Physiologischen Optik*; Voss: Leipzig, Germany, 1867.
78. Arfini, S. *Ignorant Cognition: A Philosophical Investigation of the Cognitive Features of Not-Knowing*; Springer: Cham, Switzerland, 2019.
79. Arfini, S.; Bertolotti, T. Self-knowledge and self-knowledge: Further cognitive implications of epistemic and other bubbles. In *Natural Arguments: A Tribute to John Woods*; Gabbay, D., Magnani, L., Park, W., Pietarinen, A.-V., Eds.; Tribute Series; College Publications: London, UK, 2019; Volume 40, pp. 665–692.
80. Woods, J. Epistemic bubbles. In *We Will Show Them! Essays in Honour of Dov Gabbay on his 60th Birthday*; Artemov, S., Barringer, H., d'Avila Garcez, A., Lamb, L.C., Woods, J., Eds.; College Publications: London, UK, 2005; Volume 2, pp. 731–774.
81. Cohen, L.J. *An Essay on Belief and Acceptance*; Clarendon Press: Oxford, UK, 1992.
82. Walton, D.; Krabbe, E.C.W. *Commitment in Dialogue*; SUNY Press: Albany, NY, USA, 1995.
83. Woods, J. *Aristotle's Earlier Logic*; Hermes Science: Oxford, UK, 2001.
84. Woods, J. *Truth in Fiction: Rethinking its Logic*; Springer: Cham, Switzerland, 2018.
85. Adriaans, P.; van Benthem, J. Philosophy of Information. In *Handbook of the Philosophy of Science*; Gabbay, D., Thagard, P., Woods, J., Eds.; Elsevier: Amsterdam, The Netherlands, 2008.
86. Attneave, F. *Applications of Information Theory to Psychology: A Summary of Basic Concepts*; Holt: New York, NY, USA, 1959.
87. Cover, T.M.; Thomas, J.A. *Elements of Information Theory*; John Wiley & Sons: Hoboken, NJ, USA, 1991.
88. Dretske, F. Epistemology and information. In *Handbook of the Philosophy of Science, Volume 8. Philosophy of Information*; Adriaans, P., van Benthem, J., Eds.; Elsevier: Amsterdam, The Netherlands, 2008; pp. 29–47.
89. Floridi, L. *The Blackwell Guide to the Philosophy of Computing and information*; Blackwell: Oxford, UK, 2003.
90. Floridi, L. Philosophy of information and information ethics: Critical reflections on the state of the art. *Ethics Inf. Technol.* **2008**, *10*, 89–96.
91. Bar-Hillel, Y. *Language and Information: Selected Essays on Their Theory and Application*; Addison-Wesley: Reading, MA, USA, 1964.
92. Dretske, F. *Knowledge and the Flow of Information*; MIT Press: Cambridge MA, USA, 1981.
93. Hartley, R.V.L. Transmission of information. *Bell Tech. J.* **1928**, *7*, 535–563. [CrossRef]
94. Shannon, C.E. A mathematical theory of communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423, 623–656. [CrossRef]
95. Chaitin, G.J. *Algorithmic Information Theory*; Cambridge University Press: New York, NY, USA, 1987.
96. Li, M.; Itánge, P. *An Introduction to Kolmogorov Complexity and Its Applications*; Springer: Berlin/Heidelberg, Germany, 1997.
97. Solomonoff, R.J. The discovery of algorithmic probability. *J. Comput. Syst. Sci.* **1997**, *55*, 73–88. [CrossRef]
98. Hintikka, J. *Socratic Epistemology: Explorations of Knowledge-Seeking by Questioning*; Cambridge University Press: Cambridge, UK, 2007.

99. Hintikka, J. Abduction-inference, conjecture, or an answer to a question? In *Socratic Epistemology: Explorations of Knowledge-Seeking by Questioning*; Cambridge University Press: New York, NY, USA, 2007; pp. 38–60.
100. Zimmerman, M. The nervous system and the context of information theory. In *Human Physiology*; Schmidt, R.F., Thews, G., Eds.; Springer: Berlin/Heidelberg, Germany, 1989; pp. 166–175.
101. Schiffrin, R.M. Automatism and consciousness. In *Scientific Approaches to Consciousness*; Cohen, J.D., Schooler, J.W., Eds.; Erlbaum: Mahwah, NJ, USA, 1997; pp. 49–64.
102. Woods, J. Four grades of ignorance-involvement and how they nourish the cognitive economy. *Synthese* **2021**, *198*, 3339–3368. [CrossRef]
103. Callen, H.B. *Thermodynamics and an Introduction to Thermostatistics*, 2nd ed.; Wiley: Hoboken, NJ, USA, 1985.
104. Pelletier, F.J. (Ed.) *Kinds, Things and Stuff*; Oxford University Press: Oxford, UK, 2010.
105. Tappenden, J. Mathematical concepts: Fruitfulness and naturalness. In *The Philosophy of Mathematics and Mathematical Practice*; Mancosu, P., Ed.; Oxford University Press: New York, NY, USA, 2008; pp. 276–301.