Information, Possible Worlds and the Cooptation of Scepticism

Luciano Floridi[1, 2, 3]

[1]Research Chair in Philosophy of Information and GPI, University of Hertfordshire; [2]Faculty of Philosophy and IEG, University of Oxford; [3]UNESCO Chair in Information and Computer Ethics. Address for correspondence: Department of Philosophy, University of Hertfordshire, de Havilland Campus, Hatfield, Hertfordshire AL10 9AB, UK; l.floridi@herts.ac.uk

Abstract

The article investigates the sceptical challenge from an information-theoretic perspective. Its main goal is to articulate and defend the view that either informational scepticism is radical, but then it is epistemologically innocuous because redundant; or it is moderate, but then epistemologically beneficial because useful. In order to pursue this cooptation strategy, the article is divided into seven sections. Section one sets up the problem. Section two introduces Borel numbers as a convenient way to refer uniformly to (the data that individuate) different possible worlds. Section three adopts the Hamming distance between Borel numbers as a metric to calculate the distance between possible worlds. In sections four and five, radical and moderate informational scepticism are analysed using Borel numbers and Hamming distances, and shown to be either harmless (extreme form) or actually fruitful (moderate form). Section six further clarifies the approach by replying to some potential objections. In the conclusion, the Peircean nature of the overall approach is briefly discussed.

Keywords

Borel numbers; Hamming distance; informational scepticism; David Lewis; Levenshtein distance; modal metrics; philosophy of information; possible worlds; scepticism; semantic information.

1. Introduction

Sceptics have never enjoyed a good press. Indeed, only sophists seem to have been bashed by philosophers more regularly and systematically. How unfair. Sceptical challenges provide the starting points of many philosophical investigations; young epistemologists sharpen their dialectical teeth on sceptical debates; more than one academic career has been made out of alleged refutations of some sceptical position; and teaching philosophy without dealing with sceptical doubts would be like cooking without spices.

There are intuitive reasons why no brand of scepticism appears to be welcome. Some are more historical. They will not occupy us in the following pages, although I shall briefly return to them in the conclusion. Here, suffice to recall that, insofar as scepticism seeks to undermine all certainties – what Sextus Empiricus referred to as *dogma*[1] – it is a disruptive force, which is at loggerheads with intellectual conservatism and intolerant religious doctrines.[2] Crusaders and fundamentalists of all sorts inevitably loathe serious doubters.

Some other reasons are more theoretical. They usually concern the *degrees[3] of certainty* to which an epistemic agent *a* (whether an individual or a group of agents, i.e., a

---

[1] In the Greek tradition, dogma are fixed and untouchable beliefs and philosophical dogmatism is any non-Pyrrhonian philosophy advancing positive theories.

[2] True, some uses of scepticism have sometimes represented a background condition for a fideistic defence of religious faith (Floridi (2002), Popkin (2003)). But these are exceptions, which bear only a negative relation to the epistemological problems originally raised by scepticism. They stressed the weakness of human knowledge only in order to redirect the interest of the audience towards Christian faith, and hence to foster behaviour consistent with it. The emphasis was not on the nature of knowledge, but on the ethical attitude to be taken with respect to revelation and Christ's teachings.

[3] The qualification is introduced here in order to take into account epistemologically fuzzy statements such as "I am almost certain (sure) that *p*" or "I am pretty certain (sure) that *p*", or "I am not completely certain (sure) that

multi-agent system, it does not matter, so I shall not differentiate between the two cases in the rest of this article) may be justifiably entitled, when assessing *a*'s *knowledge of the external world*.[4] Thus, most of the anti-sceptical literature, at least from Augustine onwards, has focused on the possibility of showing that *a* can indeed reach irrefutable certainties, and then on arguing, from Descartes onwards, that *a* can successfully exploit at least some of these certainties in order to discriminate between bad beliefs that must be discarded (Descartes' 'rotten apples', see Descartes (1984), 7:481), and good beliefs (the 'good apples') that may constitute knowledge, given the right circumstances and proper further qualifications. This is the classic perspective from which I shall approach the sceptical challenge. There are, however, three points where my contribution will depart from other, more orthodox lines of reasoning.

First, I shall focus on information, rather than on beliefs or knowledge. Second, I will be concerned with informativeness, rather than with certainty. And finally, I will seek to co-opt moderate scepticism as a valuable ally in the search for information about the world, rather than refute all kinds of scepticism as a disease, or endorse it wholeheartedly and unreservedly as a panacea. Each point deserves a brief clarification.

Information – in the factual, semantic sense of well-formed, meaningful and truthful data about some topic (Floridi (2010)) – is a weaker concept than knowledge. This is so not least because information does not require what is a necessary feature of knowledge, namely a sufficient dose of the relevant kind of rational support, be this, for example, a satisfactory degree of justification, warranty, explanation, or account. An agent *a* might be informed that *i* – e.g., *a* might hold the information that a playing card is indeed the King of Spades –

---

p". The reader who objects to the possibility of degrees of certainty may replace them with only two binary values without any consequence for the rest of the article.

[4] I have discussed mathematical scepticism elsewhere, see Floridi (1998b), Floridi (2000), and Floridi (2004a).

without yet knowing that *i* because *a* lacks any further support in favour of *i*. You are still informed that *i* – to the extent that you can successfully use *i*, convey *i*, sell *i*, etc.– even if you merely guessed correctly that *i* and cannot explain or justify why *i* for the life of you. Information is not subject to the problems notoriously caused by epistemic luck.

Regarding the second point, it is useful to focus on the more impoverished concept of information, rather than on the richer one of knowledge, because this clarifies that the sceptical challenge concerns the empirical truth, and hence the informativeness of *i*, not the kinds or degrees of support that *a* might enjoy in holding that *i* (this soon turns into some version of the Gettier problem), or the information network that may account for *i*. Let me explain. Information goes hand in hand with unpredictability. More precisely, the Inverse Relationship Principle (IRP), as Barwise labelled it, states that there is an inverse relation between the probability of *i* and how much semantic information is carried by *i*. Nowadays, one often translates IRP modally, by stating that the semantic information conveyed by *i* is equivalent to the set of all possible worlds – or, more cautiously, to the set of all the descriptions of the relevant possible states of the actual world – that are excluded by, or are inconsistent with, *i*. So, if *a* holds some content (i.e., some well-formed and meaningful data) *c* about some topic *t*, the informativeness of *c* is inversely related to the number of possible worlds in which *c* is true. If *c* is true in all possible worlds (necessary truth) then it follows from IRP that *c* is not informative.[5] If *c* is true in no possible world (necessary falsehood) then I have argued (Floridi (2004b)) that, in this case too, *c* is not informative. In either case, the sceptical challenge regarding our information about the external world does not apply. The scope of IRP is limited to only possibly-true content: if *c* is true only in some possible worlds, then establishing in which ones makes *c* either informative or disinformative. It

---

[5] See Bar-Hillel and Carnap (1953), reprinted in Bar-Hillel (1964). Note that this is the source of Hintikka's scandal of deduction, see D'Agostino and Floridi (2009).

follows that the sceptical challenge is about whether one can treat some content that seems to be information as actually being information. Is $c = i$? The question is rephrased by the sceptic in the following way: how can you establish whether c is informative or disinformative with respect to t in a possible world W? That is, how can you establish whether c tells you something true or false about t in W? We shall see that the sceptic suggests that you cannot – mind, *not* that it is not – and hence that you should withhold your commitment regarding the informativeness of c. Assuming that $c = i$ means assuming not only that ordinary epistemic problems regarding the c in question are not occurring, or have been solved (mistakes happen, but can be corrected), but also that the more radical, sceptical challenge about the truthfulness of c has been met. This bracketing of the sceptical problem is fine (Husserl was right about this), as long as our philosophical concerns are of a different nature – for example, if we wish to clarify what it means for some semantic information to be relevant – but it would be a silly way of begging the question in this context.

Finally, there are different strategies to deal with informational scepticism. In what follows, I shall argue that either the sceptical challenge is radical, but then it is epistemologically innocuous, because informationally redundant; or it is moderate, but then epistemologically beneficial, because informationally useful. The conclusion is that, either way, there is no epistemologically harmful scepticism.

In order to pursue this cooptation strategy, I will need to make much more precise and clearer what informational scepticism consists in. This task will require the use of two sharper tools, as previously mentioned in the introduction. For we will need to be able to refer uniformly to (the data that individuate) different possible worlds (section two), and to adopt a metric that enables us to compare, with some precision, the distances between such possible worlds, or better, their data (section three). Once so equipped, it will be easy to analyse the two forms of scepticism (section four), and then show how the radical one can be tolerated,

while the moderate one should be promoted (section five). To repeat, the fundamental thesis is simple: scepticism is either a harmless exercise to be enjoyed intellectually, but not to be taken at face value epistemologically, or it is actually an essential component of any decent epistemic work, in academia as well as in real life. Before the conclusion, I shall consider some objections to the general strategy developed in this paper, in order to dispel a few potential misunderstandings, for the confused critic, and reinforce the cogency of the approach, for the unconvinced reader (section six). Finally, the overall approach to scepticism pursued in this paper is clearly much more Peircean than Cartesian (Floridi (1998a)), so it will be useful to say just a few words about this methodological aspect in the conclusion (section seven). Let us now start by sharpening our tools.

2. Possible Worlds and Borel Numbers

Suppose we toss a fair coin. The outcome may be described (modelled) by the following information *i*: the coin is heads. Using a common convention, we shall refer to anything like the actual coin being tossed and landing heads on a surface after some time as our *targeted system S*, while *i* will be said to provide a *model M of S* (M_S).

Still following a standard approach in information theory, we can disassemble *i* into a combination of a Boolean question *Q* and a Boolean answer *A*. The question is a message, requesting data, sent by the sender, our epistemic agent *a*, to the receiver, in our case the system represented by *S*. The answer is also a message, providing the requested data, sent by the queried *S* to the querying *a*. In short, we have:

$$i = Q + A \qquad\qquad [1]$$

The polarization of *i* into $Q + A$ makes evident that the role of *A* is to *saturate Q*, to adapt a Fregean idea lately borrowed by information theory.[6] Although it is trivial to apply [1] to any piece of information *i*, like 'the coin is heads', in order to obtain:

Question   'Is the coin heads?' +

Answer   <u>   'yes'</u> =

Information   'The coin is heads'

it is important to keep in mind that the correct interpretation of *Q* in [1] is not as (i) a request for *confirmation* or (ii) a *test*, but as (iii) a genuine request by *a* to *S* to erase a data deficit on *a*'s side through *saturation*. The difference is that, in (i) and (ii), *a* already holds the information that *i*, but wishes to double-check it, or to check whether the receiver also holds that information; whereas in (iii), *a* lacks the information that *i* and wishes to acquire it from the receiver, by obtaining the missing data. Since we cannot assume that *a* already holds *a priori* the information that *i*, the polarization of *i* requires a channel of communication between *a* and *S*.

Of course, in real life, questions and answers share, in variable proportions, the amount of semantic content that is to be found in the corresponding semantic information. However, in [1], the full semantic content, to be found in 'the coin is heads', is entirely allocated to *Q*. Now, one cannot expect *a* to be able always to maximise the content of *a*'s questions, for *a* often lacks much more than just a positive or negative saturation of a Boolean question by *S*. Frequently, *a* wishes to know, e.g., *what* the capital of France is, not just *whether* Paris is the capital of France. Still, given some idealization, typical of controlled

---

[6] In information theory *saturation* is the condition at which a communications system reaches its maximum capacity of traffic-handling.

experiments, it is perfectly reasonable, for our present purposes and for the rest of this article, to assume that all non-Boolean questions can be translated into Boolean equivalents.[7]

As I have argued elsewhere (Floridi (2010)), [1] requires us to specify the *context* (C) in which, the *level of abstraction* (LoA) at which, and the *purpose* (P) for which the question is formulated, and hence it is expected to be satisfied by the answer. In our case, for example, we are in a ideal context (the coin is perfectly fair); we are able to ask only about heads and tails (this is the LoA at which data are obtained from *S*) but not, for instance, about the kind of coin tossed, who is tossing it, or the length of time the coin takes to be tossed and to land on one side; and the purpose may be assumed to be merely illustrative. For the sake of simplicity, I shall refer to the combination of these three parameters by means of the acronym CLP. To recall another Fregean point, questions cannot acquire their specific meaning in isolation or independently of their CLP parameters. It is a bit of a pain, but we need to keep these variables in mind, lest the conceptual mess becomes unmanageable. They will turn out to be rather useful in section five. So, as a memory aid, let me revise [1] by adding a combined index, thus:

$$i^{\text{CLP}} = Q^{\text{CLP}} + A \qquad\qquad [2]$$

Note that it is now easy to appreciate that it is really *Q* and not *A* that sets the scope of the CLP parameters. A Boolean answer can only endorse (or reject) the *context* (C) in which, the *level of abstraction* (LoA) at which, and the *purpose* (P) for which the Boolean question is formulated; it can neither change nor challenge them.

Two more advantages offered by [2] are related and they finally lead us to the conclusion of this section. In [2], the relation 'is correctly saturated by' is a function from a

---

[7] The reader interested in an explicit defence of such translatability will find it in Floridi (2010).

countable domain of Boolean questions to a co-domain of only two possible Boolean answers {Yes, No} or better {1, 0}. This is an information-theoretic way of interpreting Frege's idea that declarative sentences (propositions) are terms that denote one of only two truth-values. Unfortunately, having only two values without any possibility of reiteration would be too restrictive, because it would make it hard to take into account the complexity of, and the possible state-transitions in, the system *S*. Luckily, [2] allows for a small improvement.

Suppose we have two fair coins, which are tossed repeatedly. [2] may still be sufficient to capture the more complex system *S* and its dynamics if we allow strings of binary values, as in Table 1.

Context, Level of Abstraction, Purpose

	Q1	Q2
	is the first coin heads?	is the second coin heads?
A1	1	0
A2	0	1

Table 1 Elementary example of polarised information.

Each numeric row in Table 1, e.g. <1, 0>, provides data on the synchronic state of *S*, obtained through a set of questions, asked relatively to some given CLP parameters. The columns provide the data on the transition state of *S*, i.e. from 1 to 0, in the first column, and from 0 to 1, in the second. The CLP must be kept invariant for the state transition or indeed any comparison (see below) to be, and remain, meaningful. Once again, this will turn out to be crucial in section five. The result is that the number of answers is equivalent to the number of time-steps in the state transition, whereas the number of questions indicates the complexity of the (analysis of the) system. Simplifying, all these data can be contained in one *n*-tuple: <1, 0

| 0, 1>, where occurrences of ',' separate synchronic data, and occurrences of '|' separate diachronic data. We are now ready to appreciate Borel numbers.

In 1927, Émile Borel, in a brief article defending mathematical finitism,[8] presented the following "odd problem", as he called it:

> One could define [a] number by saying that each of [the] successive digits [of its decimal expansion] is equal to 0 or 1 according to whether the answer to some question or other is affirmative or negative. Moreover, it would be possible to order all the questions that can be asked in the French language by sorting them [...] as is done in dictionaries. Only those questions for which the answers is *yes* or *no* would be retained. The mere knowledge of the number thus defined would give answers to all past, present and future enigmas of science, history and curiosity.[9]

Change French into English or indeed into any other sufficiently expressive language *L*; sort questions not "as is done in dictionaries" but into numbered questions asked relative to some given CLP parameters; make the set of questions finite; add a dynamic dimension (Borel's world is static), and you will recognise our *n*-tuple, expressed as a binary string.

Borel thought that his odd problem was a difficulty not in finite cases, like our elementary two-coin system, but whenever one admits real numbers that are inaccessible, or even in the case of extremely large numbers that are incomputable. Recently, Chaitin (2006) has revived Borel's criticism in order to defend a finite and digital ontology.[10] For our present purposes, however, I hope that any reader who is happy to talk about possible worlds and the boundless number of their properties will also allow Borel numbers (henceforth simply β) to

---

[8] See Marion (1998) for an interesting discussion of Borel's position.

[9] Borel (1927), p. 271, English translation in Mancosu (1998), chapter 21. I use here the slightly revised text provided by Tasić (2001), p. 52.

[10] I owe to Greg Chaitin my first introduction to what he calls Borel "know-it-all number" during the 30th International Wittgenstein Symposium in Kirchberg, in August 2007. Note that both Borel and Chaitin talk only about one Borel number, and they do not associate Borel numbers to possible worlds or any modal analysis, let alone edit distances.

be as long as we may need them to be, and hence possibly infinite, in order to capture such properties. Consider that they are no more problematic than real numbers, and they are a very efficient and elegant device to describe the data characterising a system *S* relative to some given questions and CLP parameters, in the following way (in what follows, the reader acquainted with Carnap-like state-descriptions might wish to compare them to Borel numbers; the comparison is explicitly discussed in section seven, objection four).

If one simplifies the analysis and disregards any dynamic aspect (for we now know that such a feature could easily be added anyway), one can see that the data produced by static possible worlds are Borel numbers. Alternatively, Borel numbers are snapshots of possible worlds. Thus, our one-coin, static system has a very simple β, which is either 1 or 0. The four states of the two-coin, dynamic system are described by four β, and so forth. Transitions systems are described by well-ordered sets of Borel numbers, where the ordering relation ≤ is suitably interpreted in chronological terms, e.g. 'is not a state subsequent to'. If this is unclear, here is another elementary illustration.

Imagine that our system *S* is a micro-world consisting of two cards. Our LoA consists of the following observables: a card can be either a King or not a King, and it can be either Spades or Clubs. *S* can be observed to be in sixteen possible states. Suppose the actual state of *S* from an imaginary God's eye perspective is the following:
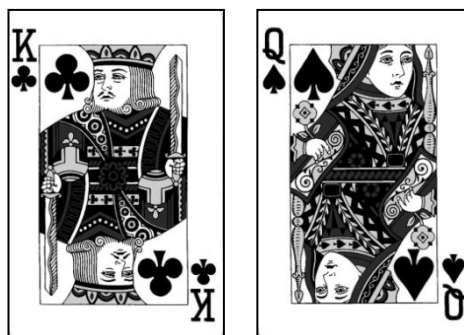

Figure 1 Example of a state of a system *S* consisting of two cards.

Depending on the CLP parameters, there are many ways of expressing the information that the first card is a King of Clubs and the second card is Spades but not a King, and this further clarifies the crucial importance of the CLP parameters. We could transform the information [the first card is a King of Clubs and the second card Spades but not a King] into the question: is the first card a King of Clubs and the second card not a King but Spades? The Borel number would simply be 1. This alternative is always available in principle, since, no matter how complex (in the sense introduced above) the system is, one can always generate such a conjunction of all the relevant questions as to require only a single yes/no answer. In light of what has been said above, we may agree to call one-digit Borel numbers Fregean numbers. Fregean numbers[11] are usually uninteresting. Our information about the world is obtained and managed at a much more finely grained level. Of course, we could ask two questions instead of only one ('is the first card such and such?' and then 'is the second so and so?'), but let us assume, just for the sake of illustration, that we work at what seems to be a more intuitive and information-theoretically fruitful level of abstraction, one at which we ask four questions: Q1: is the first card a King? Q2: is the first card Clubs? Q3: is the second card King? Q4: is the second card Spades? This way, we obtain $\beta = 1101$. Now, 1101 is the shortest Borel number that characterises the possible world S at a given time and given the specified questions and CLP parameters. It is also the Borel number that needs to be obtained by a in order to be informed about S relatively to the given questions and CLP parameters. Figure 2 provides a graphic illustration.

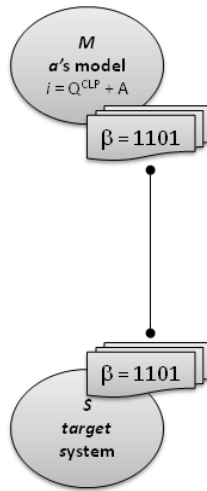[11] According to *The Hitchhiker's Guide to the Galaxy* there is a third Fregean number, 42.

Figure 2 System and Model share the same Borel number.

We now have a simple way to describe possible worlds as their shortest Borel numbers, according to a given set of questions and CLP parameters, no matter how complex a possible world is and hence how long its β needs to be. To put it in terms of Kolmogorov or algorithmic complexity: Borel numbers are the minimal computational resources needed by an informational agent such as *a* to specify a possible world. Each possible world is characterised by its Borel number. There are as many Borel numbers as (states of) possible worlds, and as many Borel numbers as possible sets of CFP parameters relative to which sets of questions are asked. Transitions are captured by well-ordered sets of Borel numbers, which can be compressed, if one disregards all the numbers that remain unchanged when moving from β_m to β_n (think of this as data refreshing cycles on β).

Of course, in order to be able to compare Borel numbers of possible worlds, both the questions asked and the relative CLP parameters must remain invariant across the comparison so, from now on, we shall simply assume that they are specified. I shall return to this important simplification below. For the moment, let us just say that Borel numbers are the

variables that make the difference. In order to analyse and compare such difference we need a metric.

3. The Edit Distance as a Modal Metrics

Since Lewis' work on counterfactuals (Lewis (1973)), it has become common to talk about degrees of *closeness* between possible worlds. Nowadays, the widespread impression seems to be that we know exactly what we are talking about. The truth is that often we do not.

The false impression may be due to the fact that we rely on simplistic assumptions and implausible scenarios in order to illustrate the notion of closeness. This is misleading. It goes without saying that a possible world in which my wife is late because of some traffic is closer to the possible world in which she is on time than a possible world in which she is late because she has been momentarily kidnapped by some extra-terrestrials and then released. But it is hardly intelligible to ask whether a possible world in which Japan never bombed Pearl Harbour is more or less close to the one in which you are reading these words, when compared to another possible world in which Hitler honoured the Treaty of Non-Aggression between Germany and the Soviet Union. It has to be said that Lewis himself was aware of the difficulty, and virtuously bit the bullet:

> Overall similarity among worlds is some sort of resultant of similarities and differences of many different kinds, and I have not said what system of weights or priorities should be used to squeeze these down into a single relation of overall similarity. I count that a virtue. Counterfactuals are both vague and various. Different resolutions of the vagueness of overall similarity are appropriate in different contexts. Lewis (1979), 465.

In truth, Lewis' comment reminds one of the notorious disclaimer in IT, according to which a problem 'it's not a bug, it's a feature'. The vagueness of the concept of modal closeness is

due to the fact that it was introduced to deal with the semantics of counterfactuals and "counterfactuals are infected with vagueness, as everyone agrees" (Lewis (1979), 457). We basically use common sense (e.g. Lewis' suggestions concerning the absence of large miracles) and some elementary logic relations (e.g. asymmetry and Lewis' back-tracking) in order to extricate ourselves from the mess they introduce. So I agree that

> we must use what we know about counterfactuals to find out about the appropriate similarity relation – not the other way around. Lewis (1979), 467.

However, that is exactly why, in the end, it becomes a matter of more or less compatible intuitions, an issue that Lewis' analysis was meant to help to resolve in the first place. So far, the conceptual difficulty. There is a second difficulty that is more formal in nature.[12]

Contrary to expectations, the formal relations of *closeness* (Lewis (1973), 29), *nearness* or *approximation* – as well as that of *similarity* often used to make sense of them[13] – do not provide a metric *d* on a set of elements, including sets of possible worlds. This because they fail to satisfy the second of the four axioms for a metric:

1. $d\,(x, y) \geq 0$ *non-negativity*

2. $d\,(x, y) = 0 \Leftrightarrow x = y$ *identity of indiscernibles*

3. $d\,(x, y) = d\,(y, x)$ *symmetry*

4. $d\,(x, z) \leq d\,(x, y) + d\,(y, z)$ *subadditivity* or *triangle inequality*

If the closeness, nearness, approximation or similarity between two possible worlds is null, then they patently cannot be the same possible world.

---

[12] But not only formal, see Kutach (2006).

[13] "The right general analysis of counterfactuals, in my opinion, is one based on comparative similarity of possible worlds", Lewis (1979), 464.

There are at least two strategies to deal with the difficulties highlighted above, short of abandoning the whole project. One is to maintain the choice of the semantic relation and decide on, for example, a *proximity space*, which axiomatises relations of 'nearness' (Naimpally and Warrack (1970)). A proximity space (X, δ) is a set X with a binary relation δ (which can be read as 'is close to') of proximity (proximity structure) on the *power set P(X)* of all its subsets, satisfying the following conditions (for all subsets A, B, C):[14]

     1. $A \,\delta\, A \Leftrightarrow A \neq 0$ *reflexivity*

     2. $A \,\delta\, B \Rightarrow B \,\delta\, A$ *symmetry*

     3. $A \,\delta\, (B \cup C) \Leftrightarrow A \,\delta\, B$ or $A \,\delta\, C$ *binary additivity*

It is easy to see that the three axioms are satisfied by Lewis' modal closeness, which is also a tolerance relation.[15] Since δ holds between sets, these could be Borel numbers.

The alternative strategy, which I shall adopt here, is to modify the semantic notion, used to compare possible worlds, in such a way as to satisfy the identity of indiscernibles axiom. Here is the rationale. The properties of proximity spaces are a generalization of the properties of a metric space. The proximity spaces which satisfy the metric axioms are known as metrizable, and this is the kind of metric space that we need here in order to calculate the distance between possible worlds. For the goal is to introduce the 'edit distance' between Borel numbers as the possibility of correcting substrings in one Borel number when this is compared to another. So, a very simple option is to choose the dual notion of *close*, i.e., *distant*, and hence of *similarity*, i.e., *dissimilarity*. Dissimilarity relations (including, for

---

[14] There are many equivalent sets of axioms for a proximity space, see for example the one provided here tends to be more common and is only slightly modified from Deza and Deza (2009), 70, in order to match more closely the previous list of metric axioms.

[15] A tolerance relation is a reflexive and symmetric relation without transitivity property, i.e. a weaker form of the equivalence relation.

example, information divergence) are still retro-compatible with Lewis' analysis of counterfactuals, since every metric space is a proximity space. But they also satisfy all four axioms for a metric space, giving us a more refined and flexible analysis of distance. Dissimilarity works at a finer LoA (points not just sets) and, above all, as I anticipated, since it is a metric, it makes possible to use the edit distance to compare Borel numbers. The idea is not entirely new. In order to introduce it, let us return to our example.

Let us assume that the system S is still in the same state as illustrated above, but that a's information i, expressed in some language L, is equivalent to:
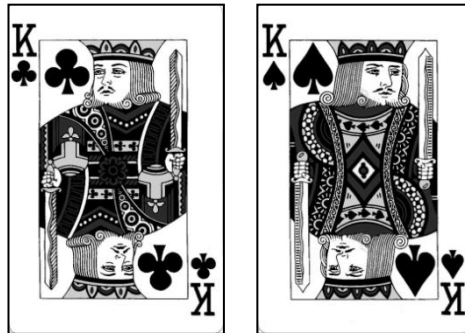

Figure 3 Example of *a*'s misinformation about the state of system *S* consisting of two cards (see Figure 1).

There is now a mismatch between $\beta_S= 1101$ and β_M 1111. In order to correct a's misinformation, β_M needs to be edited. In this elementary illustration, the editing distance consists in only one step, a replacement of the third digit (see Figure 4). This is known as the Hamming distance (*hd*), which measures the minimum number of substitutions required to change one string into another.

The Hamming distance is a standard tool in any context where two strings of data *of equal length* need to be compared for similarities (dissimilarities, in our revised approach) and correctible errors, including biological sequences, automatic spell-checkers, or pattern recognition systems. Unsurprisingly, it is a popular metric in the literature on belief- and

knowledge-base revision (Papini (2000)), or agent-based systems (Gabbay and Schlechta (2009)). In modal logic, it is one of the staple options to calculate the distance between possible worlds:

> The spheres around the worlds are assumed by Lewis to be given. Epistemologically, one would like to be able to say how to find this structure, that is, to construe it from properties of the worlds themselves. A primitive solution is that the spheres around w are the worlds [our Borel numbers] of Hamming-distance $\leq n$ for every given n. (This means that at most n values of the primitive letters can be changed in going from w to a world in that sphere.) We could also say that the worlds in the spheres result in the revision of the theory at w. Kracht and Kutz (2007), 958-959.

The Hamming distance between two binary strings, such as two Borel numbers β_n and β_m, has the nice property of being equal to the sum of 1s in their exclusive disjunction, so it can be expressed as $hd\,(\beta_n, \beta_m) = \beta_n \oplus \beta_m$.

The Hamming distance is not the only edit distance. There are others, depending, for example, on what sort and combination of edit operations, such as substitution, insertion, deletion or transposition, are allowed. In particular, it is well-known that the Hamming distance is criticisable for being too rigid (Lafage and Lang (2001)) and, since it is syntactically inflexible and defined only for inputs of the same length, in some cases it might not be the best option. If S grows or shrinks (suppose we add or remove a card in our example), or a has asked too few or too many questions, it would be useful to be able to rely on the other two operations of deletion and insertion in order to correct β_M. This give us the Levenshtein distance, which can be further improved through the Needleman-Wunch distance by adding a 'gap cost', when there are more or less 'expensive' ways of correcting a

string. These and other[16] alternatives are available to any reader interested in refining the approach. But since all such distances share the same axioms, in this paper their choice will not make a difference and, for the sake of simplicity, I suggest we adopt the simpler, Hamming distance. It is now time to analyse the sceptical challenge.
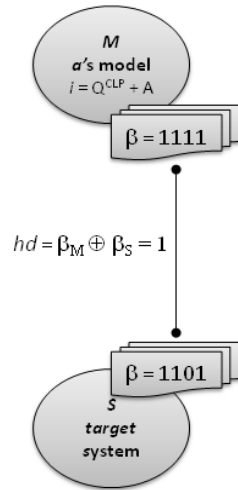

Figure 4 An example of Hamming distance between two Borel numbers.

4. Informational Scepticism or the Sceptical Challenge Reconstructed

According to the anti-sceptic philosopher, an informational agent a embedded in a possible world enjoys a boundless number of informative, zero-Hamming-distance relations with it. This means that a's empirical data about a's world are, normally, indiscernible from the data a receives from it, and when they are discernible they are either corrigible, at least in principle, or at least they are transparently incorrigible, again, in principle. More formally, given β_M, either $\beta_M \oplus \beta_S = 0$ or, if $\beta_M \oplus \beta_S \neq 0$ then, either $\beta_M \oplus \beta_S$ can, in principle, be reduced to zero by further editing, or at least a has in principle access to the meta-information

---

[16] See for example the Dalal distance, popular in model-based belief-revision analyses, which compares sets of models according to their minimal Hamming distance to other sets of models (Dalal (1988), Satoh (1988)).

that $\beta_M \oplus \beta_S \neq 0$ (*a* can be informed that *hd* should be edited but also that it is not improvable). Either way, scepticism may be dismissed as unjustified.

The sceptic, of course, begs to differ. He does not argue that, given *a*'s β_M, the distance between *a*'s model and its targeted system is greater than zero, but that the length of such distance cannot be established. More formally, he argues that $hd_1 = \beta_M \oplus \beta_S = ?$, not that $hd_1 = \beta_M \oplus \beta_S > 0$. Such challenge is motivated by the fact that *a* might be radically misinformed. Circumstances could be such as to give *a* the impression that $hd_1 = \beta_M \oplus \beta_S = 0$, when actually $hd_1 = \beta_M \oplus \beta_S \neq 0$. Such circumstances are well known: *a* might be dreaming, or be a brain in a vat, or an avatar in *Second Life* who believes to be a human being in *First Life*, or a biological body in a *Matrix* scenario, or a citizen constantly and secretly misled by some Big Brother. The examples can easily be multiplied but, at least since Descartes, they all share the following format (see Figure 5). Perhaps (*a* does not have this bit of information), *a* is not informed about the real world (the system *S*), despite all appearances that $hd_1 = 0$, because the following is a logical possibility: *a*'s model *M*, which allegedly carries information about *S*, is actually indistinguishable[17] ($hd_2 = 0$) from another model (call it *D* from dreamt), which carries information ($hd_3 = 0$) about a possible world (call it *V* for virtual), and *V* and *S* might (*a* does not have this bit of information either) be very different from each other, in which case *a* might be utterly misinformed. More formally, informational scepticism argues that, since $hd_2 = 0$ and $hd_3 = 0$ might be the cases, then $hd_4 = ?$ and therefore $hd_1 = ?$

---

[17] "Quasi scilicet non recorder a similibus etiam cogitationibus me alias in somnis fuisse delusum; quae dum cogito attentius, *tam plane video nunquam certis indiciis vigiliam a somno posse distingui* [my italics] ut obstupescam, et fere hic ipse stupor mihi opinionem somni confirmet", Descartes, *First Mediation*, 1641.
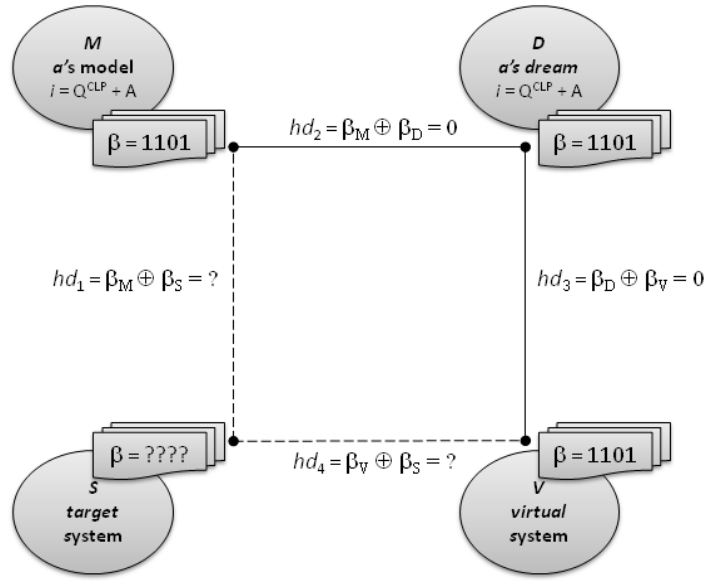
Figure 5 Example of Informational Scepticism

5. The Redundancy of Radical Informational Scepticism

Two quick clarifications are now in order, before showing why informational scepticism is harmless. First, the sceptic must (and is also happy to) concede that, because of the triangle inequality then $hd_5 \leq hd_2 + hd_3$, but since $hd_2 + hd_3 = 0$, then $hd_5 = 0$ (see Figure 6). This is exactly as it ought to be because the disruptive hypothesis, which the sceptic is inviting us to entertain, is that a's information and its corresponding $hd = 0$ might not concern S but V.

Second, the challenge posed by informational scepticism has been formulated in terms of $(hd_4 = \beta_V \oplus \beta_S = ?) \rightarrow (hd_1 = \beta_M \oplus \beta_S = ?)$, but it could be formulated equally well in terms of: $(hd_6 = \beta_D \oplus \beta_S = ?) \rightarrow (hd_1 = \beta_M \oplus \beta_S = ?)$.

At this point we know that:

1. $hd_2 = 0$ sceptical hypothesis

2. $hd_3 = 0$ sceptical hypothesis

3. $hd_5 = 0$ from 1, 2 and the triangle inequality axiom

4. $hd_1 = hd_4$ from 3 and the triangle inequality axiom

It follows that, in order to meet the sceptical challenge, it is sufficient to show that we can establish one of the following distances: hd_1, hd_4, or hd_6.

One approach, the direct analysis of hd_6, is precluded by informational scepticism itself. The hypothesis – to be taken seriously by the anti-sceptic who wishes to engage with the sceptic – is that there is no way of inspecting whether D is informative about S. The dreaming argument, for instance, presupposes that you cannot just wake up and realise you were dreaming, because you may then be dreaming that you are no longer dreaming, and so forth. Pascal summarises the point nicely:

> The main strengths of the Pyrrhonists – I shall leave aside the lesser ones – are that we can be in no way sure of the truths of these principles apart from faith and revelation, except that we feel them to be natural to us. Now this natural feeling is not a convincing proof of their truth, since, having no certainty, apart from faith, about whether we were created by a benevolent God, or an evil demon, or by chance, it is open to doubt whether the principles given to us are true, or false, or uncertain, depending on our origin. In addition, since none of us can be certain, apart from faith, whether we are awake or asleep, given that while asleep we believe as firmly as we do that we are awake. We think we see space, figures, movement. We feel the passage of time, we measure it; in short we behave just as we do when awake. The result is that, spending half our lives asleep, by our own admission and whatever it seems like to us, we have not the slightest conception of the truth, as all our feelings during that period are illusions. Who knows if that other half of our lives when we assume we are awake is not another form of sleep, slightly different from the first kind, which we awaken when we think we are sleeping? As we do often dream that we are dreaming, piling one dream on top of another, is it not perfectly feasible that the half of our life when we think we are awake is just itself a dream on to which the others are grafted and from which we will awaken at our death? Pascal (1995), n. 164.

In the sceptical *Matrix*, to use a different example, there is no red pill. So the disagreement boils down to diverging views about hd_1 and hd_4. Again, this is consistent with the classic debate about scepticism. The anti-sceptic holds that there are many cases of genuine information, i.e., $hd_1 = 0$, and therefore that, in such cases, $hd_4 = 0$, and hence that $hd_6 = 0$. The sceptic holds that nobody can tell whether $hd_4 = 0$, and therefore that nobody can tell whether $hd_1 = 0$ and hence whether $hd_6 = 0$. The anti-sceptic defends a position in the philosophy of information which the sceptic attacks metaphysically.
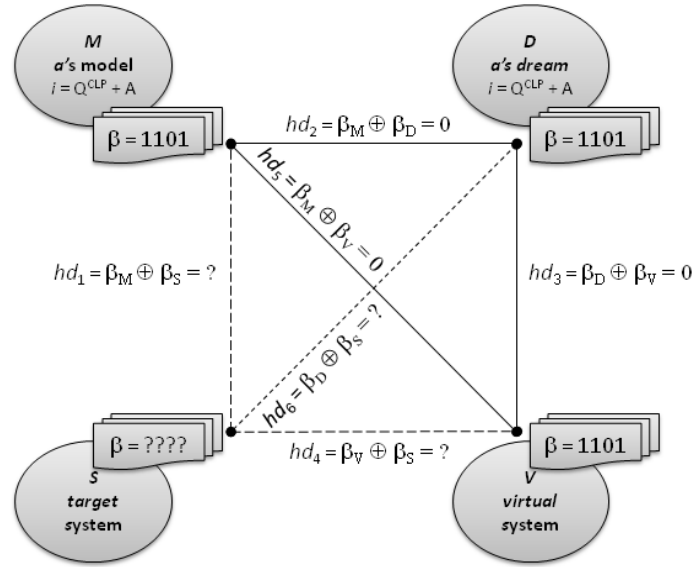
Figure 6 Some Implications of Informational Scepticism

Luckily, the situation is much better than it looks. As we know, $hd_1 \geq 0$. Trivially, if $hd_1 = 0$ then there is no sceptical challenge, so let us assume that $hd_1 > 0$. In this case, we have two alternatives:

    i) β_M cannot be edited in order to decrease hd_1, so a stops immediately; or

    ii) β_M can be edited, in which case a keeps editing it until hd_1 cannot be further decreased.

Following the sceptical challenge, the possibility of editing β_M must be understood, in both (i) and (ii), to be 'in principle': when a stops, this is because, no matter how much more editing a might still be able to do, this would not improve the distance between β_M and β_S. At this point, a cannot assume to have any direct information about the exact hd_1 between β_M and β_S. However, a does have the meta-information that there is not a single bit of information, concerning β_S, that could make an informative difference to the editing of β_M, for if there were, a would carry on the editing and stop only after the required steps (recall

24

that the sceptical challenge does not allow any 'red pill' or 'weak-up' solution). This is crucial. For consider now the following two ways of describing informational equivalence or co-informativeness:

   i) p and q are co-informative if and only if all the information in p is also inferable from q and vice versa (Larkin and Simon (1987)), i.e., if and only if they both exclude exactly the same possible worlds.

β_M and β_S are *co-informative* in (i) sense, since all the information obtainable from one is also obtainable from the other and vice versa.[18] If they were not, a would edit β_M until they become co-informative. But (i) is equivalent to (ii):

   ii) p and q are co-informative if and only if p can be transformed into q without any loss of information and vice versa. (Simon (1978))

And since β_M and β_S are *co-informative*, then they can be transformed into each other without any further editing. The result is that β_M and β_S are identical, in the following, qualified sense.

   Following Leibniz's principle of the indiscernibility of identicals, it is obvious that, if p and q qualify as information, and $p = q$, then p and q are co-informative. The identity in question, in such a *co-informativeness of identicals*, is the very strong identity one obtains only when there is no single property that distinguishes p from q, no matter what the context, LoA and purpose (CLP) might be. In other words, $p = q$ is the case if and only if one can safely ignore the CLP parameters at which the identity is predicated. Consider now Leibniz's more controversial principle of the identity of indiscernibles. At first sight, this would translate into the *identity of co-informatives*, which is obviously false. It takes only a moment to realise that p and q may be perfectly co-informative without being the same (Larkin and

---

[18] For a similar approach see Williamson (1987), who, generalising a suggestion by Geach, proves that, in the more restricted case of two relations, these are informationally equivalent if and only if each can be defined in terms of the other without the use of quantifiers.

Simon (1987)). Just imagine two fire alarms, one acoustic and the other visual, conveying exactly the same information. What is wrong with the identity of co-informatives is that it tries to adopt the same approach to CLP parameters that one finds at work in the principle of the indiscernibility of identicals. In other words, one makes the mistake of thinking that, in this case as well, it does not matter which context, LoA or purpose are in question. However, in the simple counterexample above, we can distinguish between the two co-informatives precisely because we take into account not only the identical information they convey, but also their physical features, namely a piercing noise and a flashing light. So the correct way of stating the second principle is by constraining it within the given CLP parameters according to which the relata are considered to be co-informative in the first place: if *p* and *q* are co-informative, they are so at a given LoA, in a specific context and for a particular purpose, so they are also identical, but at *that* given LoA, in *that* specific context and for *that* particular purpose. Wittgenstein's famous two issues of the same newspaper are co-informative, and therefore identical, exactly in this qualified sense. Since this is what we mean by saying that *p* and *q* are informationally equivalent, let me refer to this revised Leibnizian principle as to the *constrained identity of co-informatives*.

Let us now return to the conclusion of the previous reasoning. Since β_M and β_S are co-informative according to the CLP parameters specified, then they are identical, still according to those CLP parameters. But we said at the beginning of our investigation that we were going to ignore the CLP parameters only for the sake of simplicity. They have been there all along. All Borel numbers are always parameterised, and they must all be equally parameterised for their comparison to make any sense. So all the distances identified so far are valid at some given CLP parameters, just as the constrained identity of co-informatives is. It follows that $\beta_M = \beta_S$ according to the same CLP parameters implicitly employed when discussing all the other distances. So, contrary to our initial assumption, *hd*$_1$ cannot be greater

than 0, but given that $hd_1 = hd_4$, it follows that $hd_4 = 0$ and that $hd_6 = 0$ because of the triangle inequality axiom. Summarising, we have:

1. $hd_2 = 0$ sceptical hypothesis

2. $hd_3 = 0$ sceptical hypothesis

3. $hd_5 = 0$ from 1,2

4. $hd_1 = hd_4$ from 3 and the triangle inequality axiom

5. $hd_1 = 0$ from the constrained identity of co-informatives

6. $hd_4 = 0$ from 4 and 5

7. $hd_6 = 0$ from 2, 6 and the triangle inequality axiom

All distances are actually zero (see Figure 7). Once made fully explicit and clarified in detail, radical informational scepticism, with its fanciful scenarios of possible worlds, can be proved to be entirely redundant informationally, so it can be disregarded as harmless. Wondering whether we might be dreaming, or living in a Matrix, or might be butterflies who think they are humans, or might be characters in a sci-fi simulation created by some future civilization, and so forth, are pointless speculations that may be amusing, but make no significant difference to the serious problem of how we acquire, manage, and refine our information about the world.
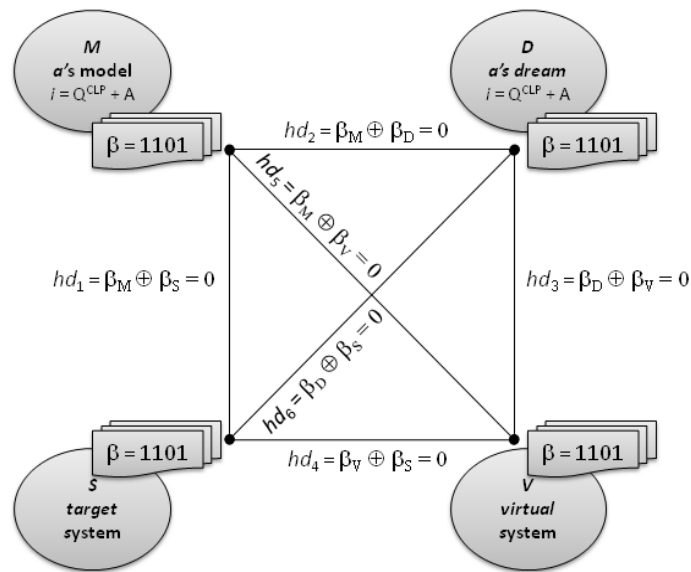
M
a's model
$i = Q^{CLP} + A$

$\beta = 1101$

D
a's dream
$i = Q^{CLP} + A$

$\beta = 1101$

$hd_2 = \beta_M \oplus \beta_D = 0$

$hd_5 = \beta_M \otimes \beta_V = 0$

$hd_1 = \beta_M \oplus \beta_S = 0$

$hd_3 = \beta_D \oplus \beta_V = 0$

$hd_6 = \beta_D \otimes \beta_S = 0$

$\beta = 1101$

S
target system

$hd_4 = \beta_V \oplus \beta_S = 0$

$\beta = 1101$

V
virtual system

Figure 7 An Example of the Redundancy of the Sceptical Challenge.

6. The Usefulness of Moderate Informational Scepticism

Absolute, radical or total forms of informational scepticism are cases of friendly fire: they seem to be attempts to investigate and test our information about the world without prejudice, but they end up promoting disastrous politics of epistemic *laissez-faire*, which scientologists and creationists, astrologists and fundamentalists as well as gurus of all kinds are way too happy to exploit. If anything goes, if there are no firewalls against idiocy and irrationality, if we create an informational vacuum, then any bogus belief has an equal right to be sold in the market of ideas. This conclusion should not be read as a dismissal of the useful role that moderate forms of informational scepticism can and must play in our intellectual lives. On the contrary, given the nature of our information-based society, there is an increasing need for a sceptical attitude towards the overabundant mass of contents (in the technical sense introduced above) to which we are exposed. Mis- and dis-information[19] are rampant, as well

---

[19] Disinformation is understood here as the intentional dissemination of misinformation.

as dogmatism and intolerance, so we need more rather than less scepticism, but it must be of the right kind. As Cajal insightfully remarks in his beautiful *Advice for a Young Investigator*, "The investigator's motto will always be Cicero's phrase: *Dubitando ad veritatem pervenimus*".[20] So where is the constructive role for sceptical doubts in the previous analysis?

Recall how we saw that *a* is supposed to edit the relevant Borel number until the distance between β_M and β_S cannot be further decreased. In less technical terms, *a* "is distinguished by a rigid worship of truth, and by a sound and genuine skepticism" (Ramón y Cajal (1999), 145), which allows *a* to question and test *a*'s information thoroughly, without prejudice but with an open mind, in order to refine it as much as this is feasible. Good scepticism offers a robust benchmark for testing our information under extreme, but still plausible, conditions. You would not dismiss a car because you cannot drive it on the Moon. This role differs from that of radical informational scepticism insofar as it is not *metaphysical* – it does not deal with logically possible worlds in which we might be deluded into thinking that we are informed about the real world – but *epistemological* or *methodological*, since it deals with actual errors, usual mistakes, potential biases, and the inevitably fallible and often probabilistic nature of our informative transactions with the world, thus exercising an open and tolerant attitude towards other views. As Abelard writes: "By doubt indeed we come to questioning; by questioning, we perceive the truth".[21] Sceptical doubts of an epistemological kind are essential for the acquisition, refinement and updating of information. They play a very useful role in the evolution of our understanding of the world.

---

[20] "By way of doubting we arrive at the truth" (Ramón y Cajal (1999), p. 55). The common attribution to Cicero is actually based on a paraphrase of a passage in his *Tusculanae Disputationes* 1.30.73.

[21] Abelard (1976), Prologue: "*Dubitando quippe ad inquisitionem venimus; inquirendo veritatem percipimus*".

7. Objections and Replies

By way of further clarification, let me now address a few potential objections that might help to dissipate some doubts and criticisms.

1) Objection: infinite Borel numbers may be incommensurable. It might be objected that, if Borel numbers can be infinite, as they might have to be if they are to encode the binary strings that answer all questions about the state of a possible worlds, then it is unclear how one might be able to compare them at all.

Reply: the difficulty is real, but irrelevant. The sceptic asks us to consider a whole possible world in which we are, for example, sleeping. So the possibility of an infinite β is something we are conceding to the sceptic, not introducing ourselves in the first place. If a strict finitist reader remains dissatisfied, we could agree to work on meaningful substrings of Borel numbers, those that, for instance, represent the Boolean answers to Descartes' questions about "the opening of the eyes, the motion of the head, the forth-putting of the hands [... the possession of] an entire body [or] hands" (*First Meditation*).

2) Objection: distances between Borel Numbers can be very counterintuitive. Suppose we have three possible worlds W_1, W_2 and W_3, such that W_1 is the empirical world in which your wife is late because of some traffic, W_2 is a possible world in which your wife is late because her car had a flat tyre that had to be replaced, and W_3 is a possible world in which your wife is late because she was momentarily abducted by some extra-terrestrials. Both the Hamming distance between W_1 and W_2 and the Hamming distance between W_1 and W_3 are 1, since, in both cases, we are assuming that it is only one question that receives a different answer, but surely this must be a *reductio ad absurdum*, because intuitively W_1 and W_2 are much less

distant (much closer, in Lewis' sense) than W_1 and W_3, at least because we have decided to conceive them to be so. It follows that the analysis is flawed: either Borel numbers do not provide a satisfactory description of possible worlds, or the Hamming edit distance between them provides the wrong metric, or indeed both, and the same holds true for other metric distances as well. In any case, the issuing analysis of radical informational scepticism must be deficient.

Reply: the objection raises a very serious problem, but this is actually different from the one that it seems to present. What is indeed untenable is the idea that possible worlds might be comparable 'ceteris paribus' without being extremely careful regarding the scope of such clause and what it really implies. Let me explain by using the same example. The mistake lies in pretending to be able to conceive three possible worlds such as W_1, W_2 and W_3, that, all other things being equal, differ only for one single feature, the cause of your wife's delay. This is absurd. A world in which your wife is late because she is spending some time on an extra-terrestrial spaceship is very distant from W_1 exactly because that event requires the coherent modification of an enormous number of other answers to other questions about the existence of aliens, of their visit to earth, of life in the universe, and so forth. Once again, and at the risk of becoming tedious, it is a question of context, levels of abstraction and purpose. Anything can be 'similar' to anything else, given some CLP parameters, but that is why we should be careful about the choice of the latter. If all we can observe, when analysing three systems such as W_1, W_2 and W_3, is determined by only one question about the cause of your wife's delay, then it is correct to say that the distance between W_1 and W_2 and W_1 and W_3 must be 1. If this is felt to be unsatisfactory, the problem does not lie with Borel numbers and their Hamming distances and the issuing analysis, but with the selection of the right CLP parameters. Make them more inclusive and you will soon see that W_1 and W_2 are much closer

than W_1 and W_3, since, if we choose the right questions, progressively less will need to be edited in W_2 than in W_3. Borel numbers and metric distances are good conceptual tools, but there is no blind and mechanical *computemus à la* Leibinz: their adequate and productive use requires intelligence and insightfulness in the handling of the CLP framework.[22]

3) Objection: edit distances might be computationally intractable. Even regular languages (the set of words accepted by a given finite automaton) are computationally complex, when it comes to calculating their Hamming (Manthey and Reischuk (2005)) or Levenshtein distances (Konstantinidis (2007)). The relevant and interesting result for the analysis proposed in this paper was already obtained by Ernvall et al. (1985), who proved that the Travelling Salesman Problem, when cities are digital strings (our Borel numbers) with Hamming distances, is NP-complete. Of course, strings may be finite, short and easily computable, but given the intractability of the worst scenarios, and the fact that complete descriptions of whole possible worlds do represent such worst scenarios, what are we gaining from the introduction of Borel numbers and Hamming distances? Aren't these just pointless technicalities, which, in all sufficiently interesting cases, merely complicate rather than facilitate the problem of describing and comparing possible worlds in terms of their distances? It seems that either the descriptions and the comparisons are trivial, in which case the formal tools are unnecessary, or they are complex, in which case the same tools are useless.

Reply: this is a fair objection, but it must be properly addressed, for it does not affect the line of reasoning developed in this paper, but rather the Lewisian approach, in the following sense. In the analysis provided in this article, I deal with potentially infinite descriptions

---

[22] *The Art of Discovery* (1685), now Leibniz (1951), 51.

(Borel numbers) only when their Hamming distances are assumed to be zero *ex hypothesi* (informational scepticism), or different from, but reducible to zero, by further investigation (informational anti-scepticism). I have not attempted to show how to calculate the distance between infinite β (two full descriptions of two possible worlds) when their *hd* might be equally infinite. For I agree that, when we have sufficiently complex (in the sense introduced above) Borel numbers, such comparisons are computationally unrealistic. Yet this does not take anything away from the project of making more rigorous the possibility of having a modal metrics. We are indeed better off by understanding what exactly is required through any serious attempt to develop it. This is partly because we can then come to realise how inane it is to speak of whole possible worlds comparisons in general, and partly because we can rely on such formal analysis to deal with degenerate cases in which infinite strings have *hd* = 0. So, the objection is correct but it does not affect the argument developed in this paper, which deals with a special case, and it underestimates the importance of providing a negative result. The objection should rather be addressed to anyone who seeks to compare infinite descriptions potentially different at an infinite number of data points.[23]

4) Objection: the analysis lacks originality. We do not really need Borel numbers and Hamming distances. Possible worlds may be described by assigning truth-values to maximal sets of atomic propositions, in some sufficiently expressive language L, which completely describe a possible world, so that modal distances can be calculated in terms of the number of atomic propositions that differ in their assigned truth values. Indeed, Borel numbers might be compared to (possibly infinite) Carnapian state-descriptions:

---

[23] Holm (2003) provides the generalization needed to deal with possibly infinite descriptions in state description semantics. I am grateful to Giuseppe Primiero for calling my attention to this point.

A state-description is defined as a conjunction or class of basic sentences (i.e., atomic sentences and negations of such) which for every atomic sentence S contains either S or non-S but not both and no other sentences. A state-description is intended to represent a possible state of affairs of the universe of discourse. Carnap (1951), 75.

Since Carnap, the idea of replacing possible worlds by their descriptions in some language L (such as constituents or model sets) has been rather popular (see for example Hintikka (1973); Niiniluoto (1987), p. 481 note 3 explicitly addresses this point). Borel numbers are just another variation of the same approach.

Reply: this is more a clarification than an objection. If the comparison between Carnap-like state-descriptions and Borel numbers helps one to understand the latter, it is very welcome. I also agree that, for finite and simple systems, the difference might not be significant. However, despite the similarity of the two approaches, there are some features that make Borel numbers preferable. First, the Q&A framework avoids the construction of the informed agent *a* as a passive recipient of messages (whole propositions) sent by *S*. Rather, the language in which Galileo's book of nature is written in that of data (Borel numbers), not of sentences, and it is up to *a* to extract the data by questioning nature in a Baconian way. Second, note that state-descriptions *à la* Carnap are bottom-up propositional constructs from a formal language, whereas Borel numbers are top-down data, extracted from contents that we consider information, so the latter do not encounter the same limitations that undermine the former (see the debate between Bar-Hillel (1951) and Carnap (1951)): Carnap-like state-descriptions seem to facilitate and invite an ontological commitment that Borel numbers make much more difficult, if not impossible. Third, Borel numbers are more economical: we use bits of information and polarise them into questions and answers, whereas Carnapian state-descriptions require propositions (our Q + A) *and* corresponding truth values as extra

keys to determine their fitness to a particular target, so that the ontological status and location of such values becomes a further issue that requires more theorising, somewhat like Frege's position about the Truth and the False. Fourth, Borel numbers, being the result of an informational analysis, are better suited than Carnapian state-descriptions to provide a full and detailed analysis of the CLP parameters at which they are obtained. A final reason worth listing here is connected with the sort of correctness theory of truth that Borel number and the Q&A framework can elicit, but I shall say more about this in the reply to the next objection and in the conclusion.

5) Objection: there is a confusion between truth-values and probability distributions. Borel numbers and their Hamming distances are introduced from an informational-theoretic perspective. However, a Borel number and its corresponding question form a propositional artefact in which 1s or 0s are equivalent to truth values and have nothing to do with the probability that the content in the question might be true or false, respectively; yet, this is exactly what is relevant in information theory, where we deal with probability distributions instead. In the one-coin system example, the 1 bit of information provided by heads (or tails) corresponds to the 50% probability that the coin might land heads (or tails), not to the binary value of the truth (or falsehood) of the assertion that it landed heads (or tails). Thus, if the coin were biased, the information provided by its tossing would be less than 1, but 1 (or 0) would still be the correct answer to the relevant question. All this means that there is at the root of the analysis a confusion between truth-values and probability distributions that undermines the whole analysis.

Reply: there is no confusion and hence no problem. One could construct the whole analysis by treating Borel numbers as probability values (Yue et al. (2008)), use, for example, the so-

called 'earth mover's distance' or Wasserstein metric as the distance function defined between probability distributions on a given metric space, and then treat the co-informativeness of β_M and β_S in terms of zero Kullback–Leibler divergence, i.e., $D_{KL} = 0$, since $D_{KL}(\beta_M \parallel \beta_S) = 0 \Leftrightarrow \beta_M = \beta_S$. Yet the exercise would be rather pointless, since this is not the issue at stake. We are treating Borel numbers as unlocking keys of specific contents, exactly as discussed in the previous objection and the consequent reply. So although the warning of a potential confusion is welcome, the actual error is nowhere committed.

6) Objection: the approach is too friendly in its rejection of radical informational scepticism. Let us assume that the analysis and the arguments presented in this paper are accepted. The problem remains that too much has been conceded to the radical sceptic. For if all $hd = 0$, as argued, then there is really no way of showing whether a lives in a virtual or the real world. The distances might all be equal to zero because there is no real world, only a virtual one in which we all live.

Reply: the objection is trivially correct, insofar as it is a matter of elementary logic that, if the Hamming distance between two relata is zero, then one may as well say that the first relatum is identical to the second. The objection, however, is more interestingly mistaken when it infers from such a platitude that the radical sceptic might be satisfied by it. What the objection is missing is that there is no room left to argue that the virtual world, in which we might live, is not the real world. Yes, it might be true that "All the world's a stage, And all the men and women merely players",[24] but this does not make any informational difference, as long as we accept that there is only this stage anyway. There is nothing to be epistemically

---

[24] William Shakespeare, *As You Like It*, II.7.

worried about calling the real virtual, or the virtual real, if the two are identical. It is only a matter of poetic taste.

7) Objection: the approach is inconclusive in its defence of moderate informational scepticism. Once again, let us assume that the analysis and the arguments presented in this paper are accepted. There is a difficulty about which we are offered no clue, namely when and how the moderate form of informational scepticism should stop. In theory, the sceptic might argue that his testing and probing could go on interminably, leaving forever undetermined whether we have reach a point when β_M and β_S are really co-informative or merely appear to be so. But if this *regressus ad infinitum* is the case, then the whole strategy is inconclusive and the sceptic has the last word.

Reply: the objection is correct in stressing that forms of moderate informational scepticism need to be handled with care. They are powerful acids that can not only clean and polish our information about the world, but also corrode it irreparably. So how can Cajal's young Investigator decide when to stop 'scepticising'? The answer is twofold. First, the approach presented here is fallibilist in Peirce's sense (see next section), but not sceptical. This means that, far from assuming that obtaining information about the world is impossible, one keeps more or less open the possibility that what appears as information might turn out to be disinformation, i.e. not information at all, but mere content, that applies to a different possible world. This is just another way of saying that Cajal's young Investigator should keep in mind that, although he is right, he might have been wrong, and has been wrong in the past. Our unconvinced opponent might still retort that this merely shifts the problem: how do we know that this is the time when we are right? And this is the other part of the answer. The objection is really asking for a way of understanding whether there is anything, in the best information

we can gather about the world, that is a clear hallmark of its truth, i.e. of its zero-Hamming distance with the targeted system. A full explanation would take us too far away, but I have argued in Floridi (2010) that such hallmark is indeed available, and it is represented by the commutative relation (in the category theory's sense of 'commutation') between the model under discussion and its target, i.e., in the vocabulary of this paper, between the proximal access to β_M and the distal access to β_S. Theories of truth often seem to be developed with passive viewers of an outside world in mind, detached observers, whether inside or outside Plato's cave, TV watchers, radio listeners, movie goers, in short, *systems users*, according to the computer science terminology favoured in this paper. The correctness theory of truth, defended in Floridi (2010), is an attempt to cater for a different sort of customer, namely embodied and embedded, creative agents, who interact with reality, who shape and build it, Plato's artisans, writers not just readers, Shakespeare's players not mere audience, in short *systems designers*. For these agents, truth is about constructing and handling informational artefacts and interacting with them *successfully*, not merely experiencing them passively. It is such successful interactions that provide the much-sought-after hallmarks of truth to our young Investigator. When they occur, then it is time to stop scepticising.

8. Conclusion: From Descartes to Peirce

The attentive reader might have noticed that, in my treatment of radical informational scepticism, I have not followed a Cartesian strategy. Instead, it is rather Peirce that has influenced the approach developed in this paper. For Descartes, radical scepticism is a means of clearing the ground for the static and permanent foundations of a new 'dogmatism', a vital element in the internal monologue of the single mind, and an essential step towards individualism and the subject's epistemic responsibility. For Peirce, a genuine form of doubt is a falsificationist means to keep the road of inquiry constantly open, a vital element in the

deontology of scientific communication, and an essential step towards the construction of a community of scientific inquirers less fallible than any of its members. The dynamic and multi-agent process of investigation, which permeates Peirce's whole philosophy, makes him aware of the importance and utility of a constructive form of scepticism of the sort I have defended above. Thus, it is thanks to a process of doubting that in *The Fixation of Belief* (Peirce (1877)) we can move from

- the method of tenacity (dogmatically holding fast to one's beliefs); to

- the method of authority (deferring to someone else the right to assess the epistemic value of a belief); to

- the *a priori* method (the intra-subjective way of coming to the acceptance of a belief without taking into account either reality or other people's minds); to

- the scientific method (the inter-subjective way of coming to an agreement about the acceptability of a belief, further constrained by reality).

So I agree with Peirce that inquiry is really prompted only by further genuine doubts of an external origin, and that a constructive scepticism shows the importance of being earnest in the pursuit of knowledge. We should follow Peirce in rejecting absolute scepticism as an anthropology (Pyrrhonian blessed state of ignorance) and as an ontology (irreconcilable dualism, nominalism, anti-realism), while appreciating it as a deontological stance in how to conduct our search for information.

aforementioned people helped me to improve the paper substantially but they are not responsible for any remaining mistakes.

References

Abelard, P. 1976, *Sic Et Non* (Chicago; London: University of Chicago Press). Critical edition by Blanche B. Boyer and Richard McKeon.

Bar-Hillel, Y. 1951, "A Note on State-Descriptions", *Philosophical Studies*, 2(5), 72-75.

Bar-Hillel, Y. 1964, *Language and Information: Selected Essays on Their Theory and Application* (Reading, Mass ; London: Addison-Wesley).

Bar-Hillel, Y., and Carnap, R. 1953, "An Outline of a Theory of Semantic Information" repr. in Bar-Hillel 1964, 221-74.

Borel, É. 1927, "À Propos De La Recente Discussion Entre M. R. Wavre Et M. P. Levy", *Revue de Métaphysique et de Morale*, 34, 271-276.

Carnap, R. 1951, "The Problem of Relations in Inductive Logic", *Philosophical Studies*, 2(5), 75-80.

Chaitin, G. J. 2006, *Meta Maths: The Quest for Omega* (London: Atlantic).

D'Agostino, M., and Floridi, L. 2009, "The Enduring Scandal of Deduction. Is Propositional Logic Really Uninformative?", *Synthese*, 167(2), 271-315.

Dalal, M. 1988, "Investigations into a Theory of Knowledge Base Revision: Preliminary Report", *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI-88)*, Minnesota: St. Paul, 475-479.

Descartes, R. 1984, *The Philosophical Writings of Descartes.* 3 vols. translated by John Cottingham, Robert Stoothoff, and Dugald Murdoch (Cambridge: Cambridge University Press).

Deza, M. M., and Deza, E. 2009, *Encyclopedia of Distances* (New York: Springer).

Ernvall, J., Katajainen, J., and Penttonen, M. 1985, "Np-Completeness of the Hamming Salesman Problem", *BIT Numerical Mathematics*, 25(1), 289-292.

Floridi, L. 1998a, "The Importance of Being Earnest: Peirce's Interpretation of Scepticism" in *C. S. Peirce Categories to Constantinople*, edited by J. van Brakel and M. van Heerden (Leuven: Leuven University Press). 47-60.

Floridi, L. 1998b, "Mathematical Scepticism: A Sketch with Historian in Foreground" in *The Skeptical Tradition around 1800*, edited by Johan van der Zande and Richard Popkin (Dordrecht: Kluwer). 41-60.

Floridi, L. 2000, "Mathematical Skepticism: The Cartesian Approach" in *Proceedings of the Twentieth World Congress of Philosophy*, edited by Akihiro Kanamori (Bowling Green: Philosophy Doc Ctr). 217-265.

Floridi, L. 2002, *Sextus Empiricus, the Recovery and Transmission of Pyrrhonism* (Oxford – New York: Oxford University Press).

Floridi, L. 2004a, "Mathematical Skepticism: The Debate between Hobbes and Wallis" in *Skepticism in Renaissance and Post-Renaissance Thought: New Interpretations*, edited by Jose Raimundo Maia Neto and Richard H. Popkin (New York: Prometheus), 143-183.

Floridi, L. 2004b, "Outline of a Theory of Strongly Semantic Information", *Minds and Machines*, 14(2), 197-222.

Floridi, L. 2010, *The Philosophy of Information* (Oxford: Oxford University Press).

Gabbay, D. M., and Schlechta, K. 2009, *Logical Tools for Handling Change in Agent-Based Systems* (New York: Springer).

Holm, R. 2003, "A Constructive Approach to State Description Semantics", *Journal of Applied Logic*, 1(1-2), 13-46.

Konstantinidis, S. 2007, "Computing the Edit Distance of a Regular Language", *Information and Computation*, 205(9), 1307-1316.

Kracht, M., and Kutz, O. 2007, "Logically Possible Worlds and Counterpart Semantics for Modal Logic" in *Philosophy of Logic*, edited by D. Jacquette (Amsterdam: Elsevier), 943-996.

Kutach, D. (ed.) 2006, Special issue on "Similarity Is a Bad Guide to Counterfactual Truth". Available online at http://www.brown.edu/Departments/Philosophy/Douglas_Kutach/Kutach_Similarity.pdf.

Lafage, C., and Lang, J. 2001, "Propositional Distances and Preference Representation", *Lecture Notes In Computer Science*, 2143, 48-59.

Larkin, J. H., and Simon, H. A. 1987, "Why a Diagram Is (Sometimes) Worth Ten Thousand Words", *Cognitive Science*, 11(1), 65-100.

Leibniz, G. W. 1951, *Selections*. Edited by Philip P. Wiener. (New York: Charles Scribner's Sons)

Lewis, D. 1979, "Counterfactual Dependence and Time's Arrow", *Noûs*, 13(4), 455-476.

Lewis, D. K. 1973, *Counterfactuals* (Oxford: Basil Blackwell).

Mancosu, P. 1998, *From Brouwer to Hilbert: The Debate on the Foundations of Mathematics in the 1920s* (New York; Oxford: Oxford University Press).

Manthey, B., and Reischuk, R. 2005, "The Intractability of Computing the Hamming Distance", *Theoretical Computer Science*, 337(1-3), 331-346.

Marion, M. 1998, *Wittgenstein, Finitism, and the Foundations of Mathematics* (Oxford: Clarendon).

Naimpally, S. A., and Warrack, B. 1970, *Proximity Spaces* (Cambridge: Cambridge University Press).

Papini, O. 2000, "Knowledge-Base Revision", *The Knowledge Engineering Review*, 15(04), 339-370.

Pascal, B. 1995, *Penseés and Other Writings* Translated by Honor Levi, with an introduction and notes by Anthony Levi. (Oxford: Oxford University Press).

Peirce, C. S. 1877, "The Fixation of Belief", *Popular Science Monthly*, 12 (November), 1-15.

Popkin, R. H. 2003, *The History of Scepticism: From Savonarola to Bayle* revised and expanded edition (New York; Oxford: Oxford University Press).

Ramón y Cajal, S. 1999, *Advice for a Young Investigator Reglas y consejos sobre investigación cientifica* (4th ed. 1916) translated by Neely Swanson and Larry W. Swanson. (Cambridge, Mass.; London: MIT Press).

Satoh, K. 1988, "Nonmonotonic Reasoning by Minimal Belief Revision", *Proceedings of the International Conference on Fifth Generation Computer Systems*, Tokyo, Japan, 455-462.

Simon, H. A. 1978, "On the Forms of Mental Representation" in *Minnesota Studies in the Philosophy of Science, Vol. Ix: Perception and Cognition: Issues in the Foundations of Psychology*, edited by C. W. Savage (Minneapolis: University of Minnesota Press), 3–18.

Tasić, V. 2001, *Mathematics and the Roots of Postmodern Thought* (Oxford: Oxford University Press).

Williamson, T. 1987, "Invertible Definitions", *Notre Dame Journal of Formal Logic*, 28(2), 244-258.

Yue, A., Liu, W., and Hunter, A. 2008, "Measuring the Ignorance and Degree of Satisfaction for Answering Queries in Imprecise Probabilistic Logic Programs", *Lecture Notes In Computer Science*, 5291, 386-400.