

	M E T A	1 6 4 7	B	Dispatch: 12.3.10	Journal: META	CE: Mahendra
	Journal Name	Manuscript No.		Author Received:	No. of pages: 18	PE: Deepa/ananth

© 2010 The Author

Journal compilation © 2010 Metaphilosophy LLC and Blackwell Publishing Ltd

Published by Blackwell Publishing Ltd, 9600 Garsington Road, Oxford OX4 2DQ, UK, and
350 Main Street, Malden, MA 02148, USA

METAPHILOSOPHY

Vol. ■■, No. ■■, ■■ 2010

0026-1068

THE PHILOSOPHY OF INFORMATION: TEN YEARS LATER

LUCIANO FLORIDI

Abstract: This article provides replies to, and comments on, the contributions to the special issue on the philosophy of information. It seeks to highlight convergences and points of potential agreement, while offering clarifications and further details. It also answers some criticisms and replies to some objections articulated in the special issue.

Keywords: information ethics, levels of abstraction, philosophy of information, semantic information, structural realism, veridicality thesis.

This is a collection of very fine articles. Their scope, depth, and insightfulness are testimonies not only to the brilliance and scholarship of their authors but also to the remarkable maturity reached by the philosophy of information (PI) during the past decade. In the late 1990s, I was searching for an approach to some key philosophical questions: the nature of knowledge, the structure of reality, the uniqueness of human consciousness, a satisfactory way of dealing with the new ethical challenges posed by information and communication technologies, to list some of the topics discussed in this issue. I had in mind a way of doing philosophy that could be rigorous, rational, and conversant with our scientific knowledge, in line with the best examples set by the analytic tradition; non-psychologistic, in a Fregean sense; capable of dealing with contemporary and lively issues about which we really care; and less prone to metaphysical armchair speculations and idiosyncratic intuitions. I was looking for a constructive philosophy which would provide answers, not just analyses, that would be as free as possible from a self-indulgent, anthropocentric obsession with us and our super-duper role in the whole universe, and respectfully sceptical of commonsensical introspections and Indo-European linguistic biases. It was a recipe for disaster, but then, sometimes, fortune favours the irresponsible. During that period of intellectual struggle and confusion, I realised one day that what I had in mind was really quite simple: a philosophy grounded on the concept of information. I was not on my way to Damascus but in Oxford, at Wolfson College, sitting on the bank of the river Cherwell, when I

© 2010 The Author

Journal compilation © 2010 Metaphilosophy LLC and Blackwell Publishing Ltd

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40

(BWUK META 1647 Webpdf:)

1 discovered that the spectacles I was looking for were on my nose. It was
 2 the summer of 1998. Six months later, I gave a talk in London, at King's
 3 College, entitled "Should There Be a Philosophy of Information?" The
 4 question was of course rhetorical, and I soon started working on the essay
 5 that became "What Is the Philosophy of Information?" It was published
 6 in this journal in 2002 (Floridi 2002).

7 Almost a decade later, it is reassuring to see that the project for a
 8 philosophy of information as a discipline in its own right was not ill
 9 conceived. Witness the fact that it would be impossible to do full justice
 10 both to the quality of the contributions in this special issue and to the
 11 value of the new and exciting area of research to which they belong. For
 12 this reason, in the following pages I shall not try to summarise or discuss
 13 every interesting issue raised by the contributors. Rather, I shall briefly
 14 highlight and seek to clarify some critical points, with the goal of at least
 15 reducing our disagreement, if not achieving a full convergence of views.
 16 As I remarked in a comparable context (Floridi 2008a), philosophy deals
 17 with problems that are intrinsically open. Intellectual disagreement is
 18 therefore an essential part of its conceptual explorations. When it is
 19 informed and reasonable, it should be welcome, not eradicated. It is a
 20 very sad restaurant, soon to be out of business, in which you can order
 21 only the same dish, no matter how delicious it is. At the same time, the
 22 controversies contained in this special issue should not eclipse the fact
 23 that there is a great deal about which we all agree, in terms of importance
 24 of topics, priority of problems, and choice of the best methods to be
 25 employed to address them. We would not be engaged in this lively
 26 dialogue if we did not share so much intellectually.

27 Let me now add a final word before closing this introductory section.
 28 There will be no space below to summarise the main lines of my research.
 29 So the reader interested in knowing more about my work might wish to
 30 have a look first at a very gentle introduction to the nature of informa-
 31 tion, written for the educated public (Floridi 2010a). The more adventur-
 32 ous reader might be interested in knowing that most of the essays referred
 33 to in this special issue were part of an overarching project and have now
 34 found their proper place, as revised chapters, in a single, more technical
 35 book (Floridi 2010b).

36 **Comments and Replies**

37 I shall be extremely brief in my comments on Hendricks's and Roush's
 38 interesting essays. Hendricks shows how *minimalism* might be useful in
 39 tackling the problem of "pluralistic ignorance," where ignorance is to be
 40 understood as lack of information, rather than lack of knowledge. I might
 41 have overlooked some error, but I must confess that I fully agree with
 42 both the analysis and the proposed solution. Roush follows a similar
 43 pattern, although in her case it is the *knowledge game* that she finds

1 valuable in order to approach the “swamping problem,” that is, the
 2 question of what added value knowledge might have, over and above
 3 mere true belief, or, I would add, mere information. What both essays
 4 share, methodologically, is a careful approach to formal details; a
 5 problem-solving orientation that allows the selection of the right infor-
 6 mation-theoretical tools; and a treatment of informational agents, their
 7 environments and the issuing processes free from psychologistic features.
 8 This is where PI most fruitfully joins forces with recent trends in formal
 9 epistemology. In both essays, for example, the agents involved could be
 10 companies, political parties, or individual human beings; it does not
 11 matter. In terms of contents, there is one more feature that I would like to
 12 stress: the central role played by *equilibria*, both negatively, when in
 13 Hendricks’s essay we need to disrupt pluralistic ignorance, and positively,
 14 when in Roush’s article we need to discover what stable conditions lead to
 15 an optimal epistemic relation with the world. Impeccable.

16
 17 The next essay, by Bringsjord, also focuses on the *knowledge game* (KG).
 18 As often when reading Bringsjord, my temptation is to treasure his
 19 insights and keep quiet. Given the circumstances, I will have to resist it.
 20 So let me start by saying that Bringsjord’s mastery of KG is not only
 21 flawless but impressive. I might have invented the game, but he certainly
 22 knows how to play it elegantly. Bringsjord and I agree that—given the
 23 *current* state and understanding of computer science—the best artefacts
 24 that artificial intelligence (AI) will be capable of engineering will be, at
 25 most, zombies: artificial agents capable of imitating an increasing number
 26 of human behaviours. In most cases (I am sure Bringsjord agrees), such
 27 agents will also be better than the fragile and fallible humans who provide
 28 the original templates. In some cases—here Bringsjord might disagree,
 29 but see below—human capacities will remain unmatched. I have in mind
 30 primarily our semantic abilities. In any case, and for a variety of
 31 converging reasons, neither of us is convinced that human-like minds
 32 might be engineered artificially. We also agree that current, off-the-shelf,
 33 artificial agents as we know them nowadays cannot answer self-answering
 34 questions and pass the test.

35 Where we seem to part company is in deciding whether this holds true
 36 also for *foreseeable* artificial agents. I think so, but Bringsjord offers a
 37 proof that this is not the case. In other words, he can foresee and forecast
 38 artificial agents that will pass the KG test. I am happy to concede the
 39 point. He might be right. Or maybe not. *Partly*, it is a matter of *details*,
 40 which Bringsjord, absolutely rightly, could not fully provide in his essay.
 41 Not his fault, but they remain the preferred hiding places for devilish
 42 disappointments. *Partly*, it is a matter of *implementation*. Sometimes what
 43 looks plausible on paper turns out to be unfeasible on the ground, thus
 44 proving to be only a logical and not an empirical possibility. *Partly*, it is a
 45 matter of *interpretation*. Passing a test means being able to pass it

1 regularly and consistently, according to qualified judges, not occasionally
 2 or only thanks to a favourable setting. Think of a reading test: if you
 3 could read only sometimes, and only when looking at a text you
 4 previously memorised, you would not qualify as a reader. Now, Bring-
 5 sjord acknowledges that “current logic-based AI is able to handle some
 6 self-answering questions. Notice, though, that I say ‘some’ self-answering
 7 questions. There is indeed a currently insurmountable obstacle facing
 8 logic-based AI that is related, at least marginally, to self-answering
 9 questions: it is simply that current AI, and indeed even foreseeable AI,
 10 is undeniably at-out impotent in the face of any arbitrary natural-
 11 language question—whether or not that question is self-answering.” So
 12 it seems that my cautious attitude is vindicated. However, let us forget
 13 about all these inclusions of “partly.” Let us assume, for a moment, that
 14 all possible reservations turn out to be de facto unjustified: details are
 15 provided, the implementation works, and there is no hermeneutic
 16 disagreement about what is going on. Test passed. This is the real point
 17 at which **I may** not be able to follow Bringsjord any further. For, as I
 18 wrote in “Consciousness, Agents and the Knowledge Game,” the KG
 19 was never meant to provide a “Floridi challenge” for AI. Let me quote
 20 the relevant passage (emphasis added):

21
 22 What is logically possible for (or achievable at some distant time in the future
 23 by) a single artificial agent or for an artificial multiagent system *is not in question*.
 24 I explained at the outset that we are not assuming some science fiction scenario.
 25 “Never” is a long time, and I would not like to commit myself to any statement
 26 like “*artificial agents will never be able (or are in principle unable) to answer*
 27 *self-answering questions*”. *The knowledge game cannot be used to argue that AI*
 28 *or AC (artificial consciousness) is impossible in principle*. In particular, its
 29 present format cannot be used to answer the question “How do you know
 30 you are not a futuristic, intelligent and conscious artificial agent of the kind
 31 envisaged in Blade Runner or Natural City?” As far as artificial agents are
 32 concerned, the knowledge game is a test to check whether AI and AC have
 33 been achieved.

(Floridi 2005a, 431)

34 So much so that the article ends by showing a slightly puzzling result: the
 35 test can also be passed by a multi-agent system made of zombies. This is
 36 one more reason why I would not be surprised if Bringsjord were
 37 completely right. A challenge is usually a negative modal statement of
 38 the form (c) “ x cannot ϕ ”, for example, “John cannot run the marathon”
 39 A test is usually a conditional (possibly modal) statement of the form (t)
 40 “if x (can) ϕ then x qualifies as y ,” for example, “if John (can) runs the
 41 marathon then John is fit.” It is true that the same ϕ , for example,
 42 running the marathon, enables one to meet the challenge and pass the
 43 test. This is where Bringsjord’s and my line of reasoning run parallel. But
 44 then, he seems to believe that I am arguing in favour of a specific
 45 interpretation of (c), whereas I am interested in (t). The purpose of the

1 KG is that of providing a test whereby you may show, given that you are
 2 conscious (not a zombie), how you know that you are (not a zombie). If
 3 that is achieved, that was the challenge I wished to meet.
 4

5 The essay by Scarantino and Piccinini and the essay by Adams form a
 6 perfect diptych in which the former provide some criticisms that are well
 7 answered by the latter. Adams himself, of course, has his own reserva-
 8 tions and, as we shall see presently, they are to be taken seriously, but let
 9 me comment first on Scarantino and Piccinini's contribution.

10 Scarantino and Piccinini are right in stressing the need for a *pluralist*
 11 *approach* to the many different uses of "information." If I may phrase
 12 such pluralism in my own words: "Information 'can be said in many
 13 ways,' just as being can (Aristotle, *Metaphysics* Γ.2), and the correlation
 14 is probably not accidental. Information, with its cognate concepts like
 15 computation, data, communication, etc., plays a key role in the ways we
 16 have come to understand, model, and transform reality. Quite naturally,
 17 information has adapted to some of Being's contours. Because informa-
 18 tion is a multifaceted and polyvalent concept, the question 'what is
 19 information?' is misleadingly simple, exactly like 'what is being?'" (Floridi
 20 2003, 40). As a consequence, any thesis about the nature of information,
 21 including the one about the *veridicality* of *semantic* information, should
 22 be handled with extra care. It would be daft, for example, to identify a
 23 piece of software as information—as we ordinarily do in IT and computer
 24 science—and then argue that, since information must be true, so must be
 25 that piece of software. "True about what?" would be the right sceptical
 26 question. Likewise, it would be unduly pedantic to insist that, given the
 27 veridicality thesis, cognitive scientists should stop speaking about infor-
 28 mation processes. Sometimes, they may be talking about information in a
 29 non-semantic sense; some other times, they may just be using a familiar
 30 synecdoche, in which the part (semantic information) stands for the whole
 31 (semantic information and misinformation), as when we speak in logic of
 32 the truth-value of a formula, really meaning its truth or falsehood. Often,
 33 they are using information as synonymous for data, or representations, or
 34 contents, or signals, or messages, or neurophysiologic patterns, depend-
 35 ing on the context, without any loss of clarity or precision. The reader
 36 looking for an initial map of the varieties of senses in which we speak of
 37 "information" can find it in Floridi 2010a.

38 Since I agree that information "can be said in many ways," I also
 39 subscribe wholeheartedly to Scarantino and Piccinini's invitation to
 40 adopt a *tolerant attitude* towards the uses to which the concept of
 41 information can be put. Bananas are not fruit, and tomatoes are not
 42 vegetables, but we know where to find them in the supermarket, and not
 43 even a philosopher should complain about their taxonomically wrong
 44 locations. So why, given their pluralism and tolerance, are philosophers
 45 so keen on rejecting the veridicality thesis?

© 2010 The Author

Journal compilation © 2010 Metaphilosophy LLC and Blackwell Publishing Ltd

1 The thesis in itself seems to be fairly harmless and most reasonable
 2 (Floridi 2005b, 2007). When you ask for some semantic information—
 3 about when the supermarket is open, for example—you never specify that
 4 you wish to receive truthful information. That goes without saying
 5 because that is what semantic information is. So if you get “false
 6 information” and go to the supermarket when it is closed, you may
 7 rightly complain about your source for having provided you with no
 8 information at all. Some semantic content *c* qualifies as semantic
 9 information *i* only if *c* is truthful. If it is not, then *c* is misinformation
 10 at best, disinformation at worst (this being misinformation wilfully
 11 disseminated for malicious purposes). Simple.

12 The veridicality thesis is also hardly original. It has been treated as
 13 obvious by several philosophers who have handled semantic information
 14 with all the required care, including Grice, Dretske, and Adams. It does
 15 make life easier when dealing with difficult and controversial issues such
 16 as some paradoxes about the alleged informativeness of contradictions
 17 (they are not informative now because they are false [Floridi 2004]); the
 18 link between semantic information and knowledge (knowledge encapsu-
 19 lates truth because it encapsulates semantic information, which, in turn,
 20 encapsulates truth, as in a three-doll matryoshka [Floridi 2006]); or the
 21 nature of relevant information [Floridi 2008b]). Despite all this, it would
 22 be ungenerous to dismiss the contribution by Scarantino and Piccinini as
 23 a fruitless mistake. Let me try to explain why.

24 Imagine that Mary is told by John that “the battery of the car is flat.”
 25 This is the sort of semantic information that one needs to consider when
 26 arguing in favour of the veridicality thesis. The difficulty is that such
 27 semantic information is the result of a complex process of elaboration,
 28 which ends with truth but certainly does not start from it. Indeed, one of
 29 the great informational puzzles is how physical signals, transduced by the
 30 nervous system, give rise to high-level, truthful semantic information.
 31 When John sees the red light ashing, there is a chain of data-processing
 32 events that begins with an electromagnetic radiation in the environment,
 33 in the wavelength range of roughly 625–740 nanometres, goes through
 34 John’s eyes and nervous system, is elaborated by him in terms of a red
 35 light ashing in front of him, is combined with regular associations on the
 36 physical side (the light being red is coupled to the battery being flat by the
 37 engineering of the system) as well as with background knowledge on
 38 John’s side (e.g., concerning signals of malfunctioning in cars). It all ends
 39 with Mary receiving John’s message that “the battery is flat.” Some
 40 segments of this extraordinary journey are known (again see Floridi
 41 2010a for a simple introduction to it), but large parts of it are still
 42 mysterious. Now, if one wishes to talk rather loosely of information from
 43 the beginning to the end of this journey and all the way through it, that is
 44 fine. We know our way in the supermarket, we can certainly handle loose
 45 talk about information. There is no need to be so fussy about words: the

1 tomatoes will be found next to the salad, and the bananas next to the
 2 apples. So I am fully convinced by Scarantino and Piccinini: from such a
 3 “supermarket approach,” the veridicality thesis is untenable, since truth
 4 or falsehood plays absolutely no role in “information” for a long while
 5 during the journey from electromagnetic radiation to “Sorry, dear, the
 6 battery is flat.” Of course, this leaves open the option of being concep-
 7 tually careful when dealing with semantic information itself, the end
 8 product of the whole process. Botanically, tomatoes and courgettes are
 9 fruit, and bananas are female flowers of a giant herbaceous plant.
 10 Likewise, in philosophy of information semantic information is well
 11 formed, meaningful and truthful data. If you still find the veridicality
 12 thesis as counterintuitive as the fruity tomatoes, just assume that
 13 Grice, Dretske, Sequoiah-Grayson, Adams, I and anyone else who
 14 endorses it are being botanically precise and talking about *premium*
 15 *semantic information*. As we saw above, some pluralism and tolerance
 16 will help.

17 A final comment before turning to Adams’s essay. I believe Scarantino
 18 and Piccinini might be on to something interesting, namely, a project
 19 concerning the various uses and meanings of information in cognitive
 20 science. If I am not mistaken, then this is most welcome, as their
 21 investigations will provide a much-needed insight into an area still
 22 under-investigated. Of course, it is to be hoped that such a project will
 23 complement and build upon the fundamental research by Barwise and
 24 Seligman on the logic of distributed systems and the analysis of informa-
 25 tion flow, and hence be consistent with their results, including the fact
 26 that “the proposal agrees with Dretske’s in that it makes information
 27 veridical. That is, if a is of type α and this carries the information that b is
 28 of type β , then b is of type β ” (Barwise and Seligman 1997, 36). I shall
 29 return to this point in a moment, in connection with Adams’s rejection of
 30 the *distributive thesis* (information closure).

31 Adams’s elegant and insightful essay deserves to be studied carefully.
 32 There is much about which I completely agree, and more that I have
 33 learnt. When I wrote above that Adams provides a useful answer to issues
 34 raised by Scarantino and Piccinini, I had in mind things like his clear and
 35 correct distinction between “a semantic notion of information . . . [under-
 36 stood as] information that p or about state of affairs f that exists in one’s
 37 cognitive system (one’s beliefs, or perceptions or knowledge)” and
 38 “information in the sense of natural sign or nomic regularity, where
 39 information can exist outside cognitive agents.” This is only an example,
 40 and the essay merits close analysis. Here I shall deal only with Adams’s
 41 two objections.

42 The first concerns the possibility of having a system acquire some
 43 semantics (ground at least some of its symbols) through supervision.
 44 According to Adams, helping a system (whether human or artificial, it

1 does not matter) to acquire the meaning of a symbol *s* might be a case of
 2 *causal derivation*, not of what I would call semantic cheating, that is, a
 3 case in which the trainer has the meaning pre-packaged and transfers it to
 4 the trainee, who not only would have been unable to acquire it except for
 5 such a present but above all still does not master it. I must admit that I
 6 remain unconvinced. We know that artificial systems with semantic
 7 capabilities are not yet available. Leave to one side whether they will
 8 ever be. The fact that currently they are not shows that the symbol-
 9 grounding problem remains unresolved. And I suspect that this is so
 10 because causally inducing a system to behave as if it had acquired and
 11 mastered the meaning of a symbol *s* is useful but still insufficient to
 12 guarantee that that system actually has gained a semantics for *s* and the
 13 capacity to use it efficiently. I agree with Adams that causal derivation
 14 might be *sufficient to teach* a potentially semantic system (e.g., a dog or a
 15 human being) the meaning of *s*, given the right circumstances (but see
 16 below the case of the slave boy). I also agree that the same causal
 17 derivation might be *sufficient to transmit* the meaning of *s*, from a system
 18 that enjoys its semantics to a system that might acquire it. I even agree on
 19 the fact that causal derivation might play a *role, perhaps crucial, in*
 20 *creating* the meaning of *s*, as when “nature teaches” a system the meaning
 21 of *s*. Where Adams and I might (but I am not sure) disagree is that causal
 22 derivation is sufficient in generating the meaning of *s* *ex nihilo*, without
 23 presupposing any previous semantics of *s* or another system proficient in
 24 handling its meaning. The problem is as old as Plato: does the slave boy
 25 (the close equivalent to a robot or “animated instrument”)¹ really
 26 understand that the diagonal of any square is the base of a square of
 27 double the area? Or is it only because Socrates knows Pythagoras’s
 28 theorem that he is able to induce the slave boy, gently but firmly, through
 29 causal derivation, to say what Socrates wishes him to say, like a dumb but
 30 well-trained parrot? I am sceptical about the slave boy’s actual acquisition
 31 of the right semantics. But even if you do not share my scepticism, and
 32 side with Plato, note that ultimately the Platonic solution is to make the
 33 slave boy’s semantics innate. Either way—the semantics is not really there
 34 or the semantics has been there all along, because pre-implanted—I
 35 remain unconvinced by Adams’s position and his comments regarding the
 36 zero-semantic commitment (see his footnote 3).

The second objection is a different case. I believe here Adams is right,
 and that he is so in an important sense. He argues that we should reject

¹ “So a slave is an animated instrument, but every one that can minister of himself is more valuable than any other instrument; for if every instrument, at command, or from a preconception of its master’s will, could accomplish its work (as the story goes of the statues of Daedalus; or what the poet [Homer] tells us of the tripods of Vulcan, ‘that they moved of their own accord into the assembly of the gods’), the shuttle would then weave, and the lyre play of itself; nor would the architect want servants, or the [1254a] master slaves.” Aristotle, *Politics*, book I, chapter V, trans. First Last Names (City: Publisher Name, 0000).

1 information closure, that is, the *distributive thesis* according to which, if a
 2 is informed both that p and that $p \rightarrow q$ then a is also informed that q . As
 3 Patrick Allo kindly reminded me, in general the problem is (at least)
 4 twofold (see Allo forthcoming). One might wish to reject information
 5 closure

- 6
 7 (i) either because a may not be informed that p ;
 8 (ii) or because p does not count as semantic information;
 9 (iii) or both, of course.

10
 11 I understand Adams as being concerned only with (i) in his contribution.
 12 His support of a tracking theory of knowledge, his references to Dretske,
 13 and his examples involving Chris all point in this direction unambigu-
 14 ously. So I shall dedicate much more attention to (i). On (ii) (and hence
 15 [iii]), let me just add that it is sufficiently open to interpretation to allow
 16 different views on the value of information closure, but that Adams and I
 17 appear to agree on the following. Recall the quotation above from
 18 Barwise and Seligman 1997. Suppose two systems a and b are coupled
 19 in such a way that a 's being (of type, or in state) F is correlated to b being
 20 (of type, or in state) G , so that $F(a)$ carries (for the observer of a) the
 21 information that $G(b)$. Information closure in this case means that, if
 22 $F(a) \rightarrow G(b)$ qualifies as information and so does $F(a)$, then $G(b)$ qualifies
 23 as well: if the low-battery indicator (a) flashing (F) indicates that the
 24 battery (b) is flat (G) qualifies as information, and if the battery indicator
 25 flashing also counts as information, then so does the battery being flat.
 26 Where Adams and I might disagree (but see below) is in relation to (i). As
 27 Adams acknowledges, I reject the *distributive thesis* in cases in which the
 28 kind of information in question is empirical (seeing, hearing, and so on),
 29 but not when it is semantic. He would like to see a more uniform
 30 approach, I resist it, but we might not be at variance. Consider the
 31 following case.

32 In the left pocket of your jacket you hold the information that, if it is
 33 Sunday, then the supermarket is closed. Your watch indicates that today
 34 is Sunday. Do you hold the information that the supermarket is closed
 35 today? The unexciting answer is maybe. Perhaps, as a *matter of fact*, you
 36 do not, so Adams is right. You might fail to make the note in the pocket
 37 and the date on the watch "click." Nevertheless, I would like to argue that
 38 you should, that is, that as a matter of *normative analysis*, you did have
 39 the information that the supermarket was closed. So much so that you
 40 will feel silly when you are in front of its closed doors and realise that, if
 41 you had been more careful, you had all the information necessary to save
 42 you the trip. You should have known better, as the phrase goes. Now, I
 43 take logic to be a normative discipline. From this perspective, the
 44 *distributive thesis* seems to me to be perfectly fine. Still from the same
 45 perspective, the *distributive thesis* is not always applicable to empirical

1 information. Adams is talking about the performance of actual players,
 2 I am talking about the rules of the game. Consider the same example.
 3 This time you *read* the following e-mail, sent by the supermarket: “The
 4 shop will be closed every Sunday.” You also read the date on your
 5 computer, which correctly indicates that today is Sunday. Have you read
 6 that the supermarket is closed today? Of course not, as we assume that
 7 there were no further messages. Should you have read that it was?
 8 Obviously not, for where was the text that you should have read? Should
 9 you have inferred that the supermarket was closed today? Surely, for that
 10 was the information that could easily be inferred from the two texts that
 11 you read. If Adams’s thesis is that information closure is at best only a
 12 matter of normative logic and certainly not an empirical fact, I am
 13 convinced.

14
 15 Like Adams’s, Colburn and Shute’s essay is another contribution from
 16 which I have learnt much. I believe Colburn, Shute and I fully converge
 17 on roughly the same conclusions, even if coming from rather different
 18 perspectives. I find this most reassuring, as evidence of a robust and
 19 convincing indication of a sound methodology. In light of this agreement,
 20 I would like to take this opportunity to stress two aspects of the method
 21 of levels of abstraction (LoAs).

22 First, it is true that once it is imported into a philosophical context, it is
 23 harder to re-apply the method in computer science in precisely the same
 24 new format. Some of the exact formalisms are inevitably lost, in order to
 25 make room for conceptual and qualitative flexibility, and this justifies the
 26 fact that it is a one-way adaptation. However, computer science must be
 27 credited for providing the conceptual resources that have led to the
 28 philosophical approach based on LoAs. We can do better philosophy by
 29 learning such an important lesson from our colleagues in the computer
 30 science department, and this much is to be acknowledged. It is a matter
 31 not of pedigree but of giving to Turing what is Turing’s.

32 The other aspect concerns exactly the roots of the method of abstraction
 33 in computer science, a science that, in so far as it is also a branch of
 34 engineering, builds and modifies its own objects, exactly as economics,
 35 law, and the social sciences may build and modify social life and
 36 interactions. LoAs do not represent only a hermeneutic device; like
 37 Dennett’s stances, they are the conditions of possibility of informational
 38 access to systems and hence what determine our models of them. In this
 39 sense, they have an ontic value that stances or other forms of “perspecti-
 40 vism” cannot but lack.

41
 42 Colburn and Shute’s essay ends with some considerations on the
 43 possibility of an ontology based on informational structures. This is
 44 where the thread of the conversation is picked up by Bueno’s essay. His
 45 contribution is a welcome expansion in the number of options favourable

1 to some form of structural realism. It seems to contain, however, just a
 2 couple of unfortunate misunderstandings of my position, which might be
 3 worth dissipating here in order to leave the reader with a set of clearer
 4 alternatives.

5 Bueno seems to accept the veridicality thesis. Excellent news. However,
 6 he is a bit cautious in some cases. He need not be. Everyone defending the
 7 veridicality thesis would agree with him when he writes that “information
 8 can be used successfully, but it need not be true for it to play a successful
 9 role. Truth is not required for empirical success, not even novel empirical
 10 success involved in the discovery of a new planet,” for the following
 11 reason: “It is easy to be confused about both ‘relevance’ and ‘misinfor-
 12 mation’ [‘false information’]. . . . That misinformation may turn out to be
 13 useful in some serendipitous way is also a red herring. False (counterfeit)
 14 banknotes may be used to buy some goods, but they would not, for this
 15 reason, qualify as legal tender. Likewise, astrological data may, acciden-
 16 tally, lead to a scientific discovery but they are not, for this reason,
 17 epistemically relevant information. Of course, there are many ways in
 18 which misinformation may be indirectly, inferentially or metatheoretically
 19 relevant, yet this is not what is in question here” (Floridi 2008b,
 20 84–85).

21 Bueno speaks of Bode’s law instead of astrological data, but the reader
 22 can see where the problem lies: understanding the veridicality thesis as if it
 23 were inconsistent with the usefulness of false content (misinformation)
 24 would be a mistake, not least because the usefulness of misinformation
 25 can be presented as a case of *ex falso quodlibet*. In a different context
 26 (Floridi 2005b), I also explained why some misinformation—for example,
 27 to be told that there will be three guests for dinner tonight when in fact
 28 only a couple is coming—might be much preferable to vacuous semantic
 29 information—for example, to be told that there will be fewer than a
 30 hundred guests tonight. Again, it is a simple exercise left to the reader to
 31 draw a similar conclusion about Newtonian physics. All this is really not
 32 a serious issue, and could be disposed of as just marginal details that can
 33 easily be rectified. Much more interesting is to try to understand what sort
 34 of theory of truth would allow us a robust approach to strictly speaking
 35 false but still valuable theories (or semantic misinformation). If I am told
 36 that the train leaves at 12.15 when in fact it leaves at 12.25, I might be
 37 slightly annoyed, but I can still catch it. In other words, an informee
 38 should prefer semantic information, but, short of that, she could still
 39 settle for, and exercise plenty of tolerance towards, valuable misinforma-
 40 tion, that is, false content that still allows her to interact with the targeted
 41 system successfully. I must confess that I am not keen on “quasi-truth,”
 42 since it seems to me to be a label under which one might hide the difficulty
 43 rather than clarify it. But I take the problem as seriously as Bueno, and I
 44 have tried to solve it through a correctness theory of truth that might deal
 45 with such cases (Floridi submitted). Even a short outline of it would take

1 us too far away here, but let me just say that the basic idea is to analyse
 2 truth in terms of correctness, and correctness in terms of a commutative
 3 relation (in the category theory's sense of "commutation") between the
 4 model under discussion and its target system, that is, between the
 5 proximal access to some semantic information and the distal access to
 6 its target system.

7 It seems that a correctness theory of truth places Informational
 8 Structural Realism (ISR), a bit closer to Bueno's Structural Empiricism,
 9 yet Bueno is concerned that I might actually be a sceptic. Allow me to put
 10 myself in excellent company: this was, and still is sometimes, the same
 11 accusation moved against Kant. He who denies epistemic access to the
 12 ultimate nature of things in themselves must be a friend of Sextus
 13 Empiricus. Absolutely not. Scepticism, when properly understood, is a
 14 family of philosophical arguments in favour of the impossibility of
 15 establishing, with total certainty, whether we have reached the truth
 16 about some particular matter. Translated into informational terms, it is
 17 an attack against the possibility of determining whether some semantic
 18 content c about a target system s might actually be a case of semantic
 19 information i about s . Is $c_s = i_s$? The sceptic does not argue in favour of a
 20 negative answer but seeks to show that one can never tell. ISR, on the
 21 contrary, is in favour of the possibility of answering an endless number of
 22 occurrences of such type question, although from a fallibilist position of
 23 course, since we might be, and have been, wrong in the past (see Floridi
 24 1996 and forthcoming). So the real divide is not between my sceptical and
 25 Bueno's anti-sceptical position but a constructionist understanding of
 26 knowledge, which is essential to grasp ISR but which Bueno disregards.
 27 This is not surprising, since Bueno seems to favour some representation-
 28 alist theory of information/knowledge, which I do not endorse. To put it
 29 very simply, I support a maker's knowledge approach (Bacon, Hobbes,
 30 Vico, Kant, Cassirer) and hold that gaining information and hence
 31 knowledge about the world is a matter of data *processing*, where
 32 "processing" is taken very seriously. As with cooking, the end result of
 33 our cognitive (including scientific) elaborations is absolutely realistic,
 34 since without ingredients and proper baking there is no cake; but the
 35 outcome does not represent or portray, or x-morph, or take a picture of
 36 the ingredients or of the baking. Knowledge delivers conceptual artefacts,
 37 which are as real and objective as the cake you are eating. This anti-
 38 representationalism should not be confused with any version of anti-
 39 realism.

40 A final remark before closing. So far, I have been talking as if ISR
 41 concerned information only understood semantically, epistemically, cog-
 42 nitively, or methodologically. It does not. ISR defends primarily an
 43 *ontological* thesis, namely, an informational understanding of reality. This
 44 seems a significant difference from Bueno's structural empiricism. It is
 45 well captured by Steven French (forthcoming):

1 In effect what Floridi's approach allows us to do is separate out the
 2 commitments of ESR [Epistemic Structural Realism] from both Worrall's
 3 agnosticism and Poincaré's espousal of 'hidden natures'. At the first order LoA
 4 [level of abstraction], ESR—as the name suggests—offers us an epistemic form
 5 of realism to the effect that what we know in science are the relevant structures
 6 (and if we are to follow this line of analysis, we should perhaps drop or at least
 7 modify the afore-mentioned slogan). Beyond that, Poincaré, Worrall *et al.*
 8 should remain quiet; there should be no talk of natures, hidden or otherwise,
 9 no adoption of forms of agnosticism, but rather a 'quietist' attitude to any
 10 further commitments. To make those is to proceed to the next level, as it were,
 11 and here the appropriately metaphysically minimal attitude is that offered by
 12 OSR [Ontic Structural Realism], which reduces the amount of humility we
 13 have to swallow by reconceptualising the underlying (putative) objects them-
 14 selves in structural terms.

15 In a popularization of his views on the ultimate nature of reality, Frank
 16 Wilczek (2008) presents ordinary matter as a secondary manifestation of
 17 what he calls *the Grid*, namely (what we perceive as mere) empty space
 18 (but that is actually) a highly structured entity. Wilczek was awarded the
 19 Nobel Prize in physics in 2004 for his contribution to the discovery of
 20 asymptotic freedom in the theory of the strong interaction. I shall not
 21 pretend to understand the sophisticated physics that led him to his views
 22 about matter. Metaphysically, however, I am very sympathetic, because it
 23 seems to me that Wilczek's Grid is the physicist's counterpart of what I
 24 have defined as the *infosphere*: a non-materialist view of the world as the
 25 totality of informational structures dynamically interacting with each
 26 other. This is the ontology I defend in ISR.

27 The final essay on which I shall comment looks at my work on the ethics
 28 of the infosphere. Volkman's main contention is clearly stated at the
 29 beginning of his contribution: information ethics (IE) is too foundational,
 30 impartial, and universal to "do full justice to the rich data of ethical
 31 experience." The use of "data" might have been a Freudian slip, but the
 32 point is unmistakable: IE is partly useless, partly pernicious. Reading the
 33 article, one has the impression of a Manichean dichotomy between two
 34 moral discourses: the good one, which is warm, human, careful about the
 35 richness of life and the complexity of our difficult choices, bottom-up,
 36 with roots in real cases and full of *phronesis*; and the bad one, which is
 37 cold, objectifying, abstract, unable to capture the nuances of everyday
 38 experience, top-down, detached and algorithmically calculating. Virtue
 39 ethics (VE) versus IE, in case you had any doubts. If only things were so
 40 simple. The outcome of such a Manichean view is that Volkman's essay
 41 contains many insightful remarks, but very few of them concern IE.
 42 Anyone interested in an informed and reasonable discussion of IE might
 43 prefer reading the essay by Terry Bynum that provides the Epilogue to
 44 this special issue, or the excellent essay by Charles Ess (2008).

© 2010 The Author

Journal compilation © 2010 Metaphilosophy LLC and Blackwell Publishing Ltd

1 The problem with Volkman's approach is that it seeks to build a
 2 conflict of views at the cost of unfairness and lack of objectivity, when a
 3 more constructive and fruitful dialogue between IE and VE could have
 4 identified many convergences, as well as potentially bridgeable disagree-
 5 ments and complementary divisions of interest, thus acknowledging what
 6 each theory might be better positioned to provide. Personally, I have
 7 often argued that the distance between IE and VE is small, since both call
 8 our attention to the need to develop morally good constructions of agents
 9 and their societies and of the natural and artificial environments. *Poiesis*
 10 is a fundamental activity that requires careful ethical investigations, and it
 11 is fair to claim that it is IE that has stressed its crucial ethical importance
 12 in the current debate (cf. the concept of *homo poieticus*). The increasing
 13 interest in the ethics of design is proof of such timely focus. But this is
 14 how far Volkman is prepared to be friendly towards IE. Having grasped
 15 this point, he adopts the Manichean dichotomy illustrated above and
 16 tries, unsuccessfully though strenuously, to transform differences in
 17 focus, emphasis, and scope into a deep and irrecoverable fracture.

18 Interest in the essay as a critical discussion of IE starts waning once it
 19 charges IE with the patently impossible pretence of "incorrectly suppos[ing]
 20 that there are judgments regarding the being and flourishing of
 21 information entities that are not bound to the perspective of some agent,
 22 and that these judgments can enter into human decisions about what to
 23 do and who to be." This straw man, the cold and dry view from the sky
 24 that Volkman is keen to slap onto IE, is nowhere to be found in my or
 25 indeed other colleagues' work on IE. Not least because, as Volkman
 26 acknowledges, IE firmly holds that ethical investigations must be devel-
 27 oped by adopting and specifying the levels of abstraction (LoA) at which
 28 they are conducted, and therefore the context and purposes for which a
 29 LoA is privileged. If the reader is put off by "levels of abstractions," as
 30 something that sounds too close to a cold logical formalism, let me
 31 suggest replacing them here with warmer "human views." What IE argues
 32 is that our intrinsic animal biases, our egocentric drives, and our
 33 anthropocentric inclinations can be withstood, mitigated, and rectified,
 34 through reflection, education, social pressure, and a progressive improve-
 35 ment in our understanding of our roles in the universe. We start as selfish
 36 egoists interested only in ourselves, Hobbes is right, but we can and must
 37 hope to become unselfish and altruistic stewards of the world. Failure
 38 along the way is inevitable, especially at the individual level, but whatever
 39 small degree of success is achieved, it should be most welcome. We can
 40 become better agents by progressively balancing the demands of the
 41 shouting "me, always me, only me!" with the demands of the other, both
 42 biological and artificial. This is why, for example, IE is regularly
 43 compared to Buddhist ethics. Of course, we must "start where we are,"
 44 as Volkman repeatedly recommends. Yet this is trivial. There is no other
 45 place where we could start. The interesting question is whether staying

1 where we accidentally find ourselves thrown by natural evolution is good
 2 enough. IE argues that it is not (Hongladarom 2008). The alternative is an
 3 ego-colonialism that is unappealing. We read in the essay article that
 4 “although I cannot succeed in my life by becoming someone else, it is
 5 equally true that my own success depends on extending my self by
 6 including others in my very constitution.” It is on this well-meant
 7 inclusiveness from within, rather than respectful acceptance from with-
 8 out, that some of the worst deeds have been justified. Especially now-
 9 adays, it seems irresponsibly self-indulgent to enjoy the reassuring
 10 scenarios in which there are only friends and loving agents in ethics,
 11 while the rest is politics. What happens when **the world** neither wishes nor
 12 consents to be included in our “very constitution” but asks respectfully
 13 to remain other from us? How can we deal with conflicts between polarised
 14 agents, all bent on “starting where they are” and unwilling to step out of
 15 their egocentric predicaments? IE rejects the option “Go out into the
 16 highways and hedges and force them to enter that my house may be
 17 filled” (Luke 14:23). Augustine was keen on that passage, which provided
 18 textual justification for the Crusades.

19 Unfortunately, having made the crucial false step of misunderstanding
 20 IE for a cold, objectifying, abstract approach to human morality, the
 21 essay stumbles on several other points. **The** contemporary shift of the
 22 ethical discourse, from being entirely agent-centred to being progressively
 23 (and at least equally if not) more patient-centred is disregarded at a cost,
 24 although it represents a crucial novelty in such areas as medical ethics,
 25 bioethics, environmental ethics, or indeed information ethics. Accusa-
 26 tions of *historicism* (or, alternatively, *anachronism*, if the historical
 27 development fails to support a theory) leave the conceptual debate
 28 untouched. The list of other missed opportunities to debate IE in its
 29 real nature rather than as a caricature is too long not to become tedious.
 30 For example, pluralism is intrinsic to IE, which also defends the crucial
 31 importance of the *overridable* nature of the respect to be paid to
 32 informational entities, a feature that explicitly makes IE both willing
 33 and able “to discriminate between the information entities that merit
 34 respect and admiration and those that have not earned this status.” Or
 35 take Volkman’s misrepresentation of the boy in the junkyard example. I
 36 provided it as an excessively simplified thought experiment to illustrate
 37 pros and cons of different ethical theories. Volkman uses it as a target to
 38 which he addresses rhetorical questions: “Is the boy really just getting
 39 mindless **kicks, or is** he rehearsing his shot? How much time are we
 40 talking about? What are the alternatives open to him? What brought him
 41 here, and where is he going? There are myriad coherent stories in which
 42 it would be perfectly O.K. to smash things.” But the rhetorical game of
 43 adding “richness” to an intentionally streamlined example is trivial, and
 44 anyone can play it: “Nobody grants that breaking windscreens necessarily
 45 leads to a bad character, life is too short to care and, moreover, a boy who

1 has never broken a car windscreen might not become a better person after
 2 all, but a repressed maniac, who knows? Where did David practice before
 3 killing Goliath? Besides, the context is clearly described as ludic, and one
 4 needs to be a real wet blanket to reproach a boy who is enjoying himself
 5 enormously, and causing no apparent harm, just because there is a chance
 6 that his playful behaviour may perhaps, one day, slightly contribute to
 7 the possible development of a moral attitude that is not praiseworthy”
 8 (Floridi 1999, 54).

9 I fully subscribe to the view that “if impartialism and universalism turn
 10 out to be undesirable in themselves, at least when carried beyond their
 11 appropriate domains, then much of ethics since the Enlightenment has
 12 been a mistake, with IE as the most recent and most glaring example.” It
 13 is exactly because I believe that much of ethics since the Enlightenment
 14 has been a success and that IE is the most recent development of such a
 15 worthy tradition that I wholeheartedly hope that ethics will maintain a
 16 reasonable defence of both impartiality and universality. A fair and
 17 tolerant society depends on them, and we are getting more global by the
 18 day. We need to find a way to dialogue impartially and universally. What
 19 one might argue is that the impartial and universal application of
 20 morality needs to be consistent with the diversity of the agents and
 21 patients involved, and the variety of their predicaments. This is not a
 22 point made by Volkman, but I doubt anyone would disagree about it.

23 In conclusion, the essay represents a missed opportunity. Since it opens
 24 with a famous and beautiful quotation from Emerson, allow me to close
 25 my few remarks with a classic one by Shakespeare:

HORATIO: O day and night, but this is wondrous strange!

HAMLET: And therefore as a stranger give it welcome.

There are more things in heaven and earth, Horatio,

Than are dreamt of in your philosophy. (*Hamlet*, act 1, scene 5)

33 Conclusion

34 Information has been a subject of philosophical interest for a very long
 35 time. In a way, one could read the whole history of philosophy as
 36 containing a thin electric-blue line that runs from the pre-Socratic
 37 philosophers to us. Obvious developments in our technology, society
 38 and culture have brought to light such continuous, uninterrupted thread,
 39 which I have characterised in my work as the philosophy of information
 40 (PI). PI has opened up a very rich area of conceptual investigations. Now,
 41 the development of new philosophical ideas seems to be more akin to
 42 economic innovation than we usually acknowledge. For when Schumpeter
 43 (1943) adapted the idea of “creative destruction,” in order to
 44 interpret economic innovation, he might as well have been talking about
 45 intellectual development. This is the way I understand the metaphor of

1 the digital phoenix used by Bynum and Moor (1998) (see the next essay).
 2 This special issue shows how much creative destruction has been caused
 3 by PI. I hope it is only the beginning.

4
 5 *Department of Philosophy*
 6 *University of Hertfordshire*
 7 *de Havilland Campus*
 8 *Hatfield, Herts AL10 9AB*
 9 *United Kingdom*
 10 *l.floridi@herts.ac.uk*

11 **Acknowledgments**

12
 13 I am most grateful to all the contributors in this special issue for their
 14 time, insightful comments and tough challenges; to Patrick Allo for his
 15 patient and excellent work as guest editor, and his insightful Introduction;
 16 and to Otto Bohlmann (managing editor) and Armen T. Marsoobian
 17 (editor in chief) of *Metaphilosophy* for having organised and hosted this
 18 special issue and allowed us to work together.
 19

20 **References**

- 21
 22
 23 Allo, Patrick. Forthcoming. "The Logic of 'Being Informed' Revisited
 24 and Revised." *Philosophical Studies*.
 25 Barwise, Jon, and Jerry Seligman. 1997. *Information Flow: The Logic of*
 26 *Distributed Systems*. Cambridge: Cambridge University Press.
 27 Bynum, Terrell Ward, and James Moor, eds. 1998. *The Digital Phoenix:*
 28 *How Computers Are Changing Philosophy*. Oxford: Blackwell.
 29 Ess, Charles. 2008. "Luciano Floridi's Philosophy of Information and
 30 Information Ethics: Critical Reflections and the State of the Art."
 31 *Ethics and Information Technology* 10, no. 2:89–96.
 32 Floridi, Luciano. 1996. *Scepticism and the Foundation of Epistemology: A*
 33 *Study in the Metalogical Fallacies*. Leiden: Brill.
 34 ———. 1999. "Information Ethics: On the Philosophical Foundations of
 35 Computer Ethics." *Ethics and Information Technology* 1, no. 1:33–52.
 36 ———. 2002. "What Is the Philosophy of Information?" *Metaphilosophy*
 37 33, nos. 1–2:123–145. Reprinted in James H. Moor and Terrell Ward
 38 Bynum, eds., *CyberPhilosophy: The Intersection of Philosophy and*
 39 *Computing* (Oxford: Blackwell, 2003).
 40 ———. 2004. "Outline of a Theory of Strongly Semantic Information."
 41 *Minds and Machines* 14, no. 2:197–222.
 42 ———. 2005a. "Consciousness, Agents and the Knowledge Game." *Minds*
 43 *and Machines* 15, nos. 3–4:415–44.
 44 ———. 2005b. "Is Semantic Information Meaningful Data?" *Philosophy*
 45 *and Phenomenological Research* 70, no. 2:351–70.

- 1 ———. 2006. "The Logic of Being Informed." *Logique et Analyse* 49, no.
2 196:433–60.
- 3 ———. 2007. "In Defence of the Veridical Nature of Semantic Informa-
4 tion." *European Journal of Analytic Philosophy* 3, no. 1:1–18.
- 5 ———. 2008a. "Information Ethics: A Reappraisal." *Ethics and Informa-
6 tion Technology* 10, nos. 2–3:189–204 (special issue entitled "Luciano
7 Floridi's Philosophy of Information and Information Ethics: Critical
8 Reflections and the State of the Art," edited by Charles Ess).
- 9 ———. 2008b. "Understanding Epistemic Relevance." *Erkenntnis* 69, no.
10 1:69–92.
- 11 ———. 2010a. *Information: A Very Short Introduction*. Oxford: Oxford
12 University Press.
- 13 ———. 2010b. *The Philosophy of Information*. Oxford: Oxford University
14 Press.
- 15 ———. Forthcoming. "Information, Possible Worlds, and the Cooptation
16 of Scepticism." *Synthese*.
- 17 ———. Submitted. "Semantic Information and the Correctness Theory of
18 Truth.
- 19 Floridi, Luciano., ed. 2003. *The Blackwell Guide to the Philosophy of
20 Computing and Information*. Oxford: Blackwell.
- 21 French, Steven. Forthcoming. "The Interdependence of Structure, Ob-
22 jects and Dependence." *Synthese*.
- 23 Hongladarom, Soraj. 2008. "Floridi and Spinoza on Global Information
24 Ethics." *Ethics and Information Technology* 10, no. 2:175–87.
- 25 Schumpeter, J. A. 1943. *Capitalism, Socialism, and Democracy*. London:
26 G. Allen and Unwin.
- 27 Wilczek, Frank. 2008. *The Lightness of Being: Mass, Ether, and the
28 Unification of Forces*. New York: Basic Books.

