

A Plea for Non-naturalism as Constructionism

Luciano Floridi¹ 

Abstract Contemporary science seems to be caught in a strange predicament. On the one hand, it holds a firm and reasonable commitment to a healthy naturalistic methodology, according to which explanations of natural phenomena should never overstep the limits of the natural itself. On the other hand, contemporary science is also inextricably and now inevitably dependent on ever more complex technologies, especially Information and Communication Technologies, which it exploits as well as fosters. Yet such technologies are increasingly “artificialising” or “denaturalising” the world, human experiences and interactions, as well as what qualifies as real. So the search for the ultimate explanation of the natural seems to rely upon, and promote, the development of the artificial, seen here as an instantiation of the non-natural. In this article, I would like to try and find a way out of this apparently strange predicament. I shall argue that the naturalisation of our knowledge of the world is either philosophically trivial (naturalism as *anti-supernaturalism* and *anti-preternaturalism*), or mistaken (naturalism as *anti-constructionism*). First, I shall distinguish between different kinds of naturalism. Second, I shall remind the reader that the kinds of naturalism that are justified today need to be protected and supported pragmatically, but they are no longer very interesting conceptually. We know how to win the argument. We just have to keep winning it. Whereas the kind of naturalism that is still interesting today is now in need of revision in order to remain acceptable. Such a kind of naturalism may be revised on the basis of a realistic philosophy of information, according to which knowing is a constructive activity, through which we do not merely represent the phenomena we investigate passively, but create more or less correct informational models (semantic artefacts) of them, proactively and interactively. I shall conclude that the natural is in itself artefactual (a semantic construction), and that the information revolution is disclosing a tension

✉ Luciano Floridi
luciano.floridi@oii.ox.ac.uk

¹ Oxford Internet Institute, University of Oxford, 1 St Giles, Oxford OX1 3JS, UK

not between the natural and the non-natural, but a deeper one between a user's and a producer's interpretation of knowledge. The outcome is a philosophical view of knowledge and science in the information age that may be called constructionist and a revival of philosophy as a classic, foundationalist enterprise.

Keywords Constructionism · Closure · Level of abstraction · Naturalism · Naturalisation · Open question argument · Reasonably reiterable query test

1 Introduction

This article is a plea for non-naturalism, not an argument for it, let alone a proof or demonstration of it. If you are a naturalist, I am afraid I believe you are mistaken, but I shall not try to convince you that you are, and that you should change your mind. Naturalism is not a bad position to hold. It is a dominating *Ur*-philosophical thesis—call it a Kuhnian paradigm, at least in contemporary, English-speaking philosophy departments¹—and a widespread faith (I shall specify later exactly in what), which ultimately can withstand any objection or counter-evidence. In this, I agree with Jack Ritchie (a naturalist himself):

There is no clever philosophical argument that can force you to make this naturalist move [that is, as Ritchie writes just before: “start with our well-developed science and build our philosophy from there”]. There is no argument to show that other approaches to philosophy will not succeed. Naturalism is a position that can claim at best only *motivations* [my italics] (Ritchie 2008), p. 196.

The dialogue between the best forms of naturalism and the best forms of non-naturalism remains open.² I think this very openness, an empirical fact that any naturalist should acknowledge, is also a very significant reflective reason in favour of non-naturalism, but more on this later in the article. For the moment, I would like to clarify that I am a non-naturalist because I find naturalism's *motivations* unsatisfactory. Serious naturalism is called science. And serious naturalisation is the reduction of some philosophical area of investigation to a scientific one. In both cases, a specific philosophical enterprise ends and a new scientific one begins, and that is the end of the story. Quine was right:

Naturalism does not repudiate epistemology, but *assimilates* [my italics] it to empirical psychology (Quine 1981), p. 72.

This holds true for all sorts of philosophical naturalisms, which do not repudiate their correspondent naturalised branches of philosophy. But make no mistake: *assimilation* is a matter of absorption, incorporation, conversion, and ultimately disappearance of the assimilated in favour of the assimilating. In Quine's version, naturalism is just a euphemism for the end of epistemology. The same seems to

¹ De Caro and Macarthur (2010), p. 3.

² See for example the attempt made by De Caro and Macarthur (2010) to develop a form of “liberal naturalism”.

apply to similar programs such as the one presented in Maddy (2007). And once assimilation is expanded to all areas of philosophy, from metaphysics to ethics, naturalism becomes another way of talking about the death of philosophy. Now, I do not disagree on the logic but, as it will become clearer below, I do differ on the premises. As I argued elsewhere,³ philosophy deals with ultimate questions that are intrinsically open to reasonable and informed disagreement. And these are not “assimilable” to scientific enquiries. True, whenever such questions are discovered to be “closed” in principle, then their answering is (and should indeed be) naturalised and moves out of the scope of philosophy. Western philosophy has been outsourcing “closed” questions since the Pre-Socratics. If one wants to know the nature of colours, for example, the best thing to do is to engage in some serious cognitive science or neuroscience, and to stop tinkering conceptually and pretending that the investigation has got anything to do with epistemology. The pretence generates only an internal “discourse” among equally minded philosophers, who are going to be ignored by science and society alike, what I have called “scholasticism”. The same holds true for the study of linguistic meaning, human psychological motivations, the nature of animal reasoning, or the relation between physical space (not spatiality) and physical time (not temporality). Of course, philosophy should be conversant and consistent with the best of our knowledge and understanding, and hence with science. And of course anyone is free to call any very abstract scientific investigation “philosophy”; most of the philosophy of quantum physics seems to be just very theoretical physics, for example. The same holds true in areas such as the philosophy of mathematics or the philosophy of biology. All this is perfectly fine. And if this is naturalism, then good philosophy has always been naturalist, and bad philosophy is merely anti-scientific nonsense. However, at least since Quine, naturalism means much more than just science-friendly philosophy or abstractly rarefied science. We saw that it means “assimilation”. And it is this “assimilation” than I find both unconvincing and irresponsible, as I shall clarify in the following pages.

Withstanding scientific assimilation is not the only motivation for my plea, which is also partly due to a strange predicament in which contemporary science seems to be caught today. On the one hand, science holds a firm and reasonable commitment to a healthy naturalistic methodology, according to which explanations of natural phenomena should never overstep the limits of the natural itself. Such a “closure” (more on this technical concept below) applies, as it should, also to social and human phenomena, from economics and sociology to neuroscience and psychology. On the other hand, contemporary science is also inextricably and now inevitably dependent on ever more complex technologies, especially Information and Communication Technologies, which it exploits as well as fosters. Yet such technologies are increasingly “artificialising” or “denaturalising” (no evaluation implied here) the world, human experiences and interactions, as well as what qualifies as real, and what may be epistemically accessible. One only needs to recall the data-based and software-driven nature of contemporary research in the physics of particles, in neuroscience, or in the social sciences. So the search for the ultimate

³ See (Floridi 2011) and (Floridi 2013a).

explanation of the natural seems to rely upon, and promote, the development of the artificial, seen here as an instantiation of the non-natural. I call it a “predicament” and not a “contradiction” because science and technology are allies, on the same (and right) side of the naturalist versus anti-naturalist divide. But I also qualify it as “strange” because the increasing and profound technologisation of science is creating a tension between what we try to explain, namely all sorts of realities, and how we explain it, through the highly artificial constructs and devices that frame and support our investigations. Naturalistic explanations are increasingly dependent on non-natural means to reach such explanations. In this article, I would like to try and find a way out of this apparently strange predicament, hence my plea. I shall argue that the naturalisation of our knowledge of reality is either philosophically trivial (naturalism as *anti-supernaturalism* and *anti-preternaturalism*), or mistaken (naturalism as *anti-constructionism*). I shall do so through the following steps. First, I shall distinguish between different kinds of naturalism. Second, I shall remind the reader that the kinds of naturalism that are justified today need to be protected and supported pragmatically, for example in terms of educational policies (think of the attempt to reject evolution in favour of creationism, or the whole debate about “post-truth”), but they are no longer very interesting conceptually. It is a battle won once and for all in theory, even if it still needs to be fought very hard in practice, again and again, to avoid the complacency of reason. Whereas the kind of naturalism that is still interesting today is now in need of revision in order to remain acceptable. Such a kind of naturalism may be revised on the basis of a realistic philosophy of information, according to which knowing is a constructive activity, through which we do not merely represent the phenomena we investigate passively, but create more or less correct informational models (semantic artefacts) of them, interactively.⁴ Note that this is a thesis about both philosophy and about science or, more precisely, it is a thesis about the nature of knowing, affecting both philosophy and science.⁵ When Quine defended his version of naturalism as scientific assimilation (see previous quote) he asserted that

Science itself tells us that our information about the world is limited to irritations of our sensory surfaces, and then the epistemological question is in turn a question within science: the question how we human animals have managed to arrive at science from such limited information (Quine 1981), p. 72.

Yet he was mistaken on several accounts. The “irritations of our sensory surfaces” may be limited, but this is an irrelevant distraction: by the same reasoning, one should then try to answer the question how Beethoven managed to arrive at his *Ode to Joy* from the seven-note diatonic musical scale, Leonardo to his *Mona Lisa* from the three colours in the RGB model, Orson Welles to his *Citizen Kane* from just black and white, and today any computer multimedia from just zeros and ones. Quine’s question is no longer a mystery as soon as one drops a *representationalist*

⁴ On this see (Floridi 2010a) and (Floridi 2014).

⁵ I am grateful to one of the anonymous reviewers for suggesting this clarification and some of the wording.

interpretation of knowledge—another unquestioned *Ur*-philosophical thesis in most English-speaking philosophy departments—and acknowledges that science tells us that our information about the world is also based very significantly on our abilities to combine and repurpose data and re-use them to help us make sense of the world. With an analogy, we are not the customers in reality’s restaurant, we are the chefs. Here is how today neuroscience puts it more precisely:

Increasingly, the brain reveals itself proactive in its interface with external reality. In the past, our conception of the brain changed from that of a mirror to that of an interpreter. Several current lines of research—in fields such as memory, motivation and attentional orienting—now begin to cast the brain as a predictor. The results of experience are integrated over various timescales in order to anticipate events relevant to the current task goals and motivational state of the individual and to tune the relevant perceptual and motor machinery accordingly. In particular, research on attentional orienting has shown how signals coding predictions about the location, identity or simple features of relevant events can influence several stages of neural processing. Recent evidence shows that these predictions are not restricted simply to the contents of events but also extend to their anticipated timing (Nobre et al. 2007), p. 465.

Kant was right: we construct the world that we experience as we experience it, we do not mirror it. So Quine’s previous quote is actually a *reductio ad absurdum*. This is why he was also wrong in his description of epistemological questions, which are not descriptive or scientific, but rather semantic and normative (e.g., “should this count as a case of knowledge?”), as Barry Stroud already objected very convincingly in the eighties.⁶

For all these reasons, I shall conclude that the natural is in itself artefactual (a semantic construction), and that the information revolution is disclosing a tension not between the natural and the non-natural, but a deeper one between a user’s and a producer’s interpretation of knowledge.⁷ The outcome is a philosophical view of knowledge and science in the information age that may be called constructionist, and a revival of a classic, foundationalist role for philosophy itself. Let me clarify this position.

2 The Nature of Naturalism

“Naturalisation”, broadly understood as a family of positions defending philosophical *naturalism* (more on the latter presently), used to be a popular strategy in philosophical analysis. Figures 1 and 2 show the Ngrams for “epistemology naturalized” and “naturalization”.⁸

⁶ See Stroud (1981) revised in Stroud (1984).

⁷ For a realist, non-correspondentist but “correctness-based” theory of truth see (Floridi 2010b), for a constructionist theory of perception see (Floridi 2014), and for an interpretation of the maker’s knowledge position see (Floridi forthcoming).

⁸ Note the American spelling, which I preferred in these two Ngrams because of Quine’s article and because it seems more influential than the British English, where the “z” is replaced by the “s”. The




Fig. 1 Ngram of “epistemology naturalized”; case-insensitive; year-range = 1800–2008; Corpus = Google Books in English; smoothing = 3




Fig. 2 Ngram of “naturalization”; case-insensitive; year-range = 1800–2008; Corpus = Google Books in English; smoothing = 3

Apparently it all started in the late sixties (Rysiew 2016), with Quine’s (1969) famous and influential article on “Epistemology Naturalized”. “Naturalization” then spread to other areas, such as the philosophy of mind, the philosophy of language, the philosophy of mathematics, and metaphysics. Today, its popularity seems to have peaked. The two Ngrams lend just a bit⁹ of statistical credibility to the anecdotal impression that the nineteen–eighties and nineties were the two decades when naturalisation became mainstream in philosophy, especially in English-speaking departments. Naturalising philosophical areas of investigation seemed less fashionable at the end of the millennium. Perhaps the debate has been exhausted, with all the main positions clarified, and only minor adjustments and personal positioning still adding to the literature. Or maybe other fashions have become predominant in philosophy, such as the recent trend concerning “grounding” analyses. Whatever the reasons, the decline in the popularity of “naturalisation” seems real. It certainly makes it reasonable to explore whether some version of non-naturalism may now be plausible even for

Footnote 8 continued

results are robust, as the reader may check by testing different variants at <https://books.google.com/ngrams>.

⁹ This “bit” should be taken with a pinch of salt because of course this is a very coarse analysis. For example “naturalization” also refers to the legal process by which a person acquires a new nationality.

supporters of “naturalisation” strategies. The heat of the debate is probably over, and a calm assessment of non-naturalism may be due.

“Naturalism”—as the goal, if not the outcome, of naturalisation—may mean different things to different philosophers,¹⁰ so a clarification of how I shall use this concept in this article is overly due. I shall start from the entry in *The Oxford Philosophy Dictionary*, not as an appeal to authority, nor as a naïve endorsement of the questionable view that philosophy is done by checking dictionaries, vocabularies, and Indo-European linguistic habits, but because it is particularly clear, and provides good evidence for what I take to be today’s philosophical *vulgata*¹¹:

Naturalism: Most generally, a sympathy with the view that ultimately nothing resists explanation by the methods characteristic of the natural sciences (Blackburn 2016).

Leaving aside the pointless reference to “sympathy”, I shall follow Blackburn and treat naturalism as a thesis about *explanation*, not about *reality*, i.e. about what there is or what entities exist. This priority of *epistemological* naturalism over its *ontological* twin has a Kantian origin: it is more reasonable to decide whether there are fairies in the world depending on what kind of explanation one has for their (alleged) occurrence, rather than evaluate what kind of explanation is acceptable depending on whether it can make sense of the existence of fairies. If this seems obvious and uncontroversial, the reader should consider the fact that it applies also to the existence of God.

Having clarified the kind of naturalism I shall discuss, I shall need three concepts to offer a refined definition of *epistemological naturalism*.

The first one is that of “explanation” itself. In other contexts, I have interpreted it in terms of “providing an account” (Floridi 2012a) but here it will be sufficient to treat it as a short cut to refer to an answer to a why question, in order to include both its “reason-giving” and its “clarity-making” nature.

The second concept is “closure”.¹² This is a simple but very powerful concept in mathematics, where a set is closed under an operation if carrying out that operation on members of that set always produces a member of that set. For example, the set of positive integers 1, 2, 3, ... is closed under addition: if one adds any two positive integers one always obtains a positive integer. Yet the same set is not closed (i.e., it is *open*) under subtraction: if one subtracts a larger positive integer, say 3, from a smaller positive integer, say 2, one moves out of the positive integers and enters into the set of negative integers, by obtaining in this case -1 .

The third and last concept is that of “natural phenomena”. I shall leave it at the intuitive level, but any identification in terms of biochemical and physical state or process, subject to causal order in the world, and at least in principle knowable empirically should be fine.

¹⁰ Ritchie (2008) provides a valuable introduction although it leaves out a major and most significant debate on normative non-naturalism.

¹¹ This is common. For example, the same approach is shared by Papineau (2015), even if in terms of “investigation”. It is roughly what (Ritchie 2008) defines as “methodological naturalism”.

¹² I have used this concept to discuss the nature of philosophical questions in (Floridi 2013a).

With the aid of these three concepts, let me suggest the following as a good definition of methodological naturalism (henceforth simply naturalism):

(*N**) Naturalism:_{def} the philosophical thesis according to which all natural phenomena are closed under explanation (in the sense that every explanation of a natural phenomenon will only have to refer to one or more natural phenomena).

By using NP to refer to the set of natural phenomena and *C* to indicate “closed under explanation”, one may reformulate *N** more briefly thus:

(*N*) Naturalism:_{def} the philosophical thesis according to which *C*(NP).

Usually, supporters of *N* have in mind a strong, assimilating version of it *à la* Quine, one that combines *N* with the following two assumptions:

- A.1. If a non-natural phenomenon is explainable in terms of natural phenomena, then it is to be re-assigned to the set of natural phenomena;
and
- A.2. There are no non-natural phenomena that cannot be explained in terms of natural phenomena.

Supporters of (A.1) but not (A.2) may remain agnostic about whether there may be any non-natural phenomenon that is not re-assignable to NP. I shall return to this position later in the article. For the moment, I am interested in naturalism as the philosophical thesis that combines *N* + (A.1) + (A.2), where (A.2) means that all non-natural phenomena can be explained naturalistically, at least in principle, and thus be re-assigned to the set of natural phenomena, and therefore also be forced to close under natural explanation. *N* + (A.1) + (A.2) is what I think philosophers assume when they talk about naturalism, but let’s label this position *strong naturalism* (*SN*) just to be fair to potentially less demanding positions. According to *SN*, all natural phenomena are closed under explanation and non-natural phenomena either do not occur at all—so there is no need for an explanation—or, if they seem to occur, their explanation entirely reduces them to natural ones, hence re-assigns them to NP, and forces them to close under natural explanation. In either case, the set of non-natural phenomena (NNP) is empty *de facto* or can be emptied *de jure*, in short:

(*SN*) Strong Naturalism:_{def} the philosophical thesis according to which *C*(NP)

$\wedge \exists x (x \in \text{NNP})$, where *x* may be interpreted as ranging on all phenomena.

One final clarification is in order before discussing *SN*. *SN* assumes that the explanatory *reduction* of non-natural phenomena to natural ones is achievable exhaustively and *non-reductively*. This latter condition may seem contradictory but it is crucial. Reduction is always possible, the question is whether a full reduction fails to include at least some important features of the *explanandum* that are both significant and in need of explanation. The idea of a non-reductive reduction can be made clear and precise by introducing a metaphorical comparison¹³ with “lossless

¹³ I am grateful to one of the anonymous referees for having called my attention to a potential confusion in a previous version of this article. The appeal to data-compression is potentially misleading if one understands explanation and reduction as concerned with phenomena and not (primarily) with data (Bogen and Woodward 1988). Since the primacy of data over phenomena is not the topic of this article,

information”. Data compression can be used to reduce the size of a file. When it is lossless, i.e. non-reductive, this means that the original data can be perfectly reconstructed from the compressed data (reversibility). When it is lossy, i.e. reductive, this means that the data compression allows only an approximate reconstruction of the original data. Important features have been irrecoverably lost. Using this distinction, one can talk of a lossless (i.e., non-reductive) explanatory reduction when the *explanans* leaves no residual information in the *explanandum* unexplained. The reduction is lossy (reductive) if some residual information remains unexplained. Some simple examples in the next section will illustrate the point more intuitively. What remains to be stressed here is that a good test to check whether the explanation provided is lossy is the fact that—after the natural explanation and the reduction of the non-natural to the natural phenomena has occurred—it is still perfectly legitimate to reiterate the same query: “yes, this is correct, but why...?”. In ordinary circumstances, assuming that Alice has fully and clearly understood Bob’s explanation, it would be unreasonable of Alice to reiterate her request for an explanation, because she could not hope to be better informed by doing so. To use Wittgenstein’s analogy, she would be buying another issue of the same newspaper to double-check the news. Not a smart move. However, if it is reasonable to reiterate the query, as a request for the original explanation, then clearly one is asking a question that has not been satisfied by the answer just received. Now this is what happens with lossy or reductive explanations. Alice may ask Bob why they are organising a party on Saturday, and all the empirical facts or observations and logico-mathematical proofs or calculations available will not make it unreasonable to keep asking the typical “yes, but...” question: “Yes, I understand, but why are we organising a party on Saturday?” because Alice may be asking a normative question. The careful reader will recognise in this test of “reasonably reiterable query” (henceforth the RRQ test) a family resemblance with the open question argument developed by E.G. Moore in his *Principia Ethica* to deal with the naturalistic fallacy.¹⁴ The affinity is strong, and not accidental, since in both cases it is a form of non-naturalism that is driving the reasoning.¹⁵

I shall say more about the RRQ test later, but for now we are ready to see how *SN* applies to all non-natural phenomena. There are two (compatible) ways in which the set of non-natural phenomena may be populated and therefore, conversely, emptied by *SN*. NNP may comprise phenomena that are non-natural not so much in themselves but because they require non-natural phenomena in order to be explained. I shall analyse these in terms of *supernatural* or *preternatural* phenomena. Or (inclusive or) NNP may comprise phenomena that are non-natural in themselves, quite independently of their non-natural explanations. I shall analyse

Footnote 13 continued

the reference to data-compression should be understood here only metaphorically and not as also something that can be used to compare phenomena as well.

¹⁴ For a recent presentation see Strandberg (2004).

¹⁵ For an application of the same mechanism to understand philosophical questions as open questions see (Florida 2013a). If the affinity turns out to be an identity this is fine, in this article I am not trying to claim some particular originality but rather wish to avoid any scholastic diatribe about whether the RRQ is or is not what Moore had in mind, and how it captures his own “open question” argument.

these in terms of normative and semantic phenomena. I shall illustrate the analysis with some simple examples and argue that the first form of non-naturalism is untenable, while the second remains defensible, and indeed preferable.

A note of warning before moving to the next section. It may seem out of place to argue here in favour of naturalism and naturalisation against the supernatural and the preternatural. The objection has some merit. It is true that the fact that billions of people believe in alternative medicine, astrology, bad luck and good fortune, magic, religious miracles, and other similar “explanations” is disappointing, even disheartening, but this does not mean that this article is the right context in which to defend the voice of reason. An op-ed would be a better place, perhaps. However, the objection is not final because the following section should help to achieve two theoretical goals: support the view that naturalism is a perfectly justified position in many cases, and show-case how the RRQ test applies in some cases that are less philosophically controversial. And this should clarify the alternative cases when neither condition is satisfied.

3 Two Indefensible Non-naturalisms: The *Supernatural* and The *Preternatural*

Consider first *supernaturalism*, defined here as the view that there are phenomena which are non-natural not in themselves but in terms of their explanation, which is supposed to be paranormal, i.e., appeal to non-natural phenomena. A great example is provided by the Cottingley Fairies. In 1917 two young girls, the cousins Frances Griffith and Elsie Wright, took a series of photos of fairies while playing in the garden of Elsie’s Cottingley village home, in England. Photographic experts initially declared them genuine and they soon became famous. It was only decades



Fig. 3 The first of the five photographs, taken by Elsie Wright in 1917, showing Frances Griffiths with the alleged fairies. *Source:* Wikipedia

later, in the 1980s, that Elsie and Frances admitted that the photographs were faked, made using cardboard cut-outs of fairies copied from a popular children’s book of the time. Figure 3 shows the first.

If something does not exist then you cannot take a photograph of it, so the Cottingley Fairies, which appeared in a series of five photographs, were thought to be real and their existence to require an explanation that had to appeal to phenomena beyond the realm of the natural, i.e., *supernatural*. Conan Doyle gave a spiritualistic explanation and made the fortune of the pictures. The reasoning at the time was that ontology—the view that the fairies are real—preceded epistemology (their intelligibility)—and therefore that one needed to find an explanation for their occurrence in ordinary circumstances. The only available one was non-natural. However, according to the first assumption (A.1) of strong naturalism, if a phenomenon is explainable naturalistically without any information loss (recall: exhaustively and non-reductively) then it is a natural phenomenon. Now queries such as “why were there fairies in the Cottingley Fairies photographs?” are fully answered by proving that there weren’t any in the first place, and by giving the details of the hoax. All these answers amount to a natural explanation, and there is no information loss, meaning that any reasonable requests of explanation raised by those pictures can be fully satisfied without ever overstepping the realm of natural phenomena. Using our RRQ test, we can check that asking again, after having received the full explanation, “yes, all this is correct, but why are there fairies in the Cottingley Fairies photographs?” would simply indicate that one had not understood



Fig. 4 Fairy trails from Freeverse (2009)

the explanation in the first place. This holds true for all forms of non-naturalism as supernaturalism (including magic and miraculous phenomena), which are untenable because they are fully reducible non-reductively to naturalism.

Consider next *preternaturalism*, defined here as the view that there are phenomena which are non-natural in terms of their *source* (that is, what gives rise to them), which is supposed to be outside the realm of natural phenomena. If we wish to remain in the context of fairies, a great example is provided by an old app, no longer available, called *Fairy Trails*. Fairy Trails was a beautiful augmented reality game for iPhone and iPod touch, created in 2009 by Freeverse, a company that no longer exists. One would shake the device to power the detector and then scan the surroundings for magical creatures (Fig. 4). The app used the camera and compass to merge the real world and game world. The user would tap the fairies on the screen to collect them in a jar. The more fairies one captured, the more power one had to detect the special, quicker fairies. Today, a vaguely comparable app works to catch Pokémon.

In this case too, of course, the fairies do not exist, yet not in the same sense as the Cottingley Fairies did not, because these fairies exist as artefactual phenomena, not as a hoax. They are equally naturalisable because their lossless explanation is entirely natural, in terms of the technology that generates them. To the request of explanation for the occurrence of fairies in one's own environment the answer would have pointed to the nature of the software involved. In this case too, any further question would betray a lack of understanding. The RRQ works well.

4 Two Defensible Non-naturalisms: the Normative and the Semantic

There are phenomena that some philosophers,¹⁶ including myself, consider non-natural in themselves, and not just because of the characteristics of their explanation or sources. These are *normative* and *semantic* phenomena. They are non-natural not because they work against the ordinary laws of nature or refer to non-natural domains, but because they do not pass the RRQ test. According to normative (also known as moral or ethical) and semantic non-naturalism, normative and semantic phenomena are not naturalisable because their explanation cannot be provided in a way that appeals exhaustively and non-reductively only to natural phenomena. In both cases, any naturalistic explanation is lossy, in the sense that it is perfectly reasonable to ask again for an explanation, correctly and informatively. Their failure to pass the RRQ test shows that the request for an explanation is really a request to change the Level of Abstraction (LoA, (Floridi 2008), (Floridi 2013b)) at which the answer is expected and needs to be given, from a naturalistic to a non-naturalistic one. Failure to understand such a request of change and the insistence to use a naturalistic LoA is equivalent to committing a category mistake in Ryle's sense, or a fallacy of *μετάβασις εἰς ἄλλο γένος* ("change/shift into another genus") to use

¹⁶ If two philosophers need to be chosen as their champions, these are probably E.G. Moore and Paul Grice, respectively. The reader of this article is likely to be acquainted with the work of both, so I shall not summarise it here, where I am more interested in addressing the two positions understood as theses about the nature of explanation.

Aristotle's terminology. This was already very clear to Husserl [see (Reynaert 2015)].

For example, suppose one finds pornographic photographs of semi-naked children dressed like fairies¹⁷ and declares them to be immoral. The question why there are such pictures is entirely naturalisable. But the question why they are immoral can be asked again and again, reasonably, correctly and in an informed manner, for as long as the answer is provided in terms that appeal only to natural phenomena, e.g., because that is what a society sanctions, or because some people find them repulsive. This is also the lesson we have learnt from Hume's "no ought from is" and from Moore's "naturalistic fallacy". Similarly, there is no naturalistic way to explain the presence of fairies in Shakespeare's *A Midsummer Night's Dream*, if not by missing the point, e.g. in terms of a historical account of beliefs in fairies in Shakespeare's time. The request for explanation can only be fully satisfied by referring to the LoA represented by the internal narrative, i.e. to the semantic construction to which the fairies belong. In a combined example, one may ask why Romeo and Juliet die. A naturalistic explanation, using a CSI-like LoA, could offer the following reasons: they do not die because they do not exist; or, they die in the tragedy because in the tragedy they are victims of a double accidental suicide due to lack of communication. A non-naturalistic explanation would have to refer to their love, their relations with their families, etc., in other words, to the internal meanings of the story. The more one asks for an explanation for normative and semantic phenomena ("why...?" sorts of questions) the clearer it becomes that they can be neither eliminated *de facto*, like the Cottingley Fairies, nor forced to close under naturalistic explanation *de jure*, like the fairies in Fairy Trails. The recalcitrant, residual element that remains unexplained is precisely the all-important element that requires an explanation in the first place. In the end, it is the contribution that the mind makes to the world, and it is up to the mind to explain it, not the world. This is the sort of mind's epistemic responsibility that I find convincing in Descartes, for example. Admittedly, as one of the anonymous reviewers rightly reminded me, one may object that the recalcitrant, residual elements still in need of explanation may be just the result of our own insipience (understood as the presence of a question without the corresponding relevant and correct answer), perhaps as just a (maybe even only temporary) failure to see that there is merely a false impression of an information deficit (by analogy with a scandal of deduction, (D'Agostino and Floridi 2009)). This objection may not be incorrect. But it may not be correct either. For this objection too is open to debate, reasonably, informatively, intrinsically, permanently. And so it could be used to undermine itself: once more we would have to metatheorise the debate on the tenability of the objection and conclude that there is an openness to the questioning that the questioning itself keeps open.

¹⁷ The example is based on a real story: "University tutor asked to photograph semi-naked children convicted of pornography" <http://www.telegraph.co.uk/news/uknews/2511121/University-tutor-asked-to-photograph-semi-naked-children-convicted-of-pornography.html>.

5 In Defence of Non-naturalism

We are amphibian beings, who live partly within natural environments as bodies and mostly within normative and semantic ones as minds. The normative and semantic environment is built by our minds, but it is no less real than the laws and social rules that regulate our societies, or the money we exchange. It is also the environment through which we know and semanticise (give meaning to and make sense of) the natural environment. Seeking to reduce the natural to the non-natural means misunderstanding both, as well as our own special role within each of them. The previous analysis suggests that a more fruitful way of approaching natural and non-natural phenomena is by interpreting *initially* both as being closed under explanation, but interpreting them independently of each other, in the same way as both the set of positive integers and the set of negative integers are closed under addition and yet neither can be reduced to the other or generate members of the other set through addition. This weaker form of naturalism, to be refined presently, allows the set of non-natural phenomena (NNP) to be explained non-naturally, with closure still applying to both sets respectively. As a *first* step, it makes more sense to acknowledge that the explanation of members of NP leads to other member of NP, and the explanation of members of NNP leads to other members of NNP. If this seems just a fancy philosophical distinction consider that a naturalistic defence of miracles can be based on it: miracles are explained by other miracles, and the non-miracles by other non-miracles. For example, one may argue that the explanation of Alice's recovery from cancer is entirely natural, no magic or miracles involved, and that the explanation of the fact that she should be grateful to those who supported her during her illness may be non-natural. Difficulties arise when the natural and the non-natural intersect, e.g. when someone argues that Bob's magnetic healing abilities exist, are explainable non-naturally, and in turn explain Alice's recovery from cancer. In order to prevent such forms of supernaturalism to enter surreptitiously in our definition of non-naturalism we need to add the assumption (A.1) formulated above. In this way, we obtain the following definition of non-naturalism:

(*NN*) Non-naturalism:_{def} the philosophical thesis according to which $C(\text{NP}) \wedge C(\text{NNP}) \wedge (\text{A.1})$.

It may now be tempting to see *SN* and *NN* as opposite and irreconcilable. One may see them as the roots of the two cultures dichotomy—with scientists supporting *SN* and humanists as well as legal scholars supporting *NN*, *positivism* versus *interpretivism*—or linked to the so-called “explanatory gap” existing between the external, objective, and naturalised (quantitative and physiological) understanding of mental states and their internal, subjective, and non-naturalised (qualitative and psychological) experience, hence between science and phenomenology. Yet all these kinds of dualism would be a mistake, because one more reflective step needs to be taken: where is the whole preceding analysis taking place? In the space of the natural or in the space of the non-natural? Answering this question means realising

that ultimately the natural is artefactual and there is no dichotomy, only articulation, as I shall argue in the next section.

6 The Artefactual Nature of the Natural

If the initial distinction between *SN* and *NN* is tenable, it seems clear that, metatheoretically, its very explanation cannot belong to the set of natural phenomena.¹⁸ A self-reflective understanding of a philosophical analysis of naturalism is partly semantic and partly normative. Even the debate about the tenability of naturalism, and the very view presented at the beginning of this article that “Naturalism is a position that can claim at best only motivations”, constitute that normative–semantic residue that escapes naturalisation. Questioning it means seeking to explain the difficulty conceptually, not through empirical data, experiments, or mathematical proofs. This leads to the conclusion that our interpretation of the naturalisation of phenomena, including cultural and philosophical ones, is itself a cultural phenomenon, and hence non-naturalistic. It is the non-natural that enables us to create categories such as “natural”, “naturalised”, and “naturalisation”. There is nothing anti-realist or relativistic in such an acknowledgement. It only means that such a construction would not be reducible to a natural process without information loss. We know, semanticise, and explain reality (the system) through the construction, expansion, and refinement of our semantic artefacts about it (the model), by taking as much advantage as possible of realities (data) as constraining affordances. So either the natural is only a very minimalist and uncontroversial way to refer to data in themselves (noumenal) as constraining affordances, and this is uninteresting, or the natural—as we conceive and perceive it in a rich and multifarious way—is a construction, a model of the noumenal (reality in itself). The naturalisation of the non-natural turns out to be an expression of the artefactual nature of the natural. Ironically, naturalisation is really a second-order semanticisation. This seems to me the same road taken by John McDowell, with the only, profound difference that we need to take the opposite direction: the non-natural is our first nature, and the natural is actually our second nature. And this means that what we need is a genealogy of the natural from the non-natural, not vice versa.¹⁹ We know there is no God’s eye external perspective from which to try to analyse and solve philosophical problems. The mistake is to think that the inevitable internal perspective we must adopt can be naturalistic, when naturalism is actually the philosophical outcome of the very process of re-embedding our perspective in our own non-natural predicament. We cannot stop half way from heaven, and call at Nature’s station, we must go all the way back to the initial point of departure, a fully human constructionism. Philosophy as the ultimate form of first-order semanticisation can only be foundational or be “assimilated”, *tertium non datur*.

¹⁸ For an early and more elaborated defense of this point see (Florida 1996).

¹⁹ See Bridges (2007) on Gaskin (2006) on McDowell (1996) and more recently (De Caro and Macarthur 2010).

All this should help explain why science is increasingly artefactual. It is because its existence and features depend on the mental. As we progress in building ever more complex explanations of what there is (whatever it is in itself), we develop models that inevitably rely on more and more artificial forms of understanding. This process is kept honest by the constraints provided by the data. Our methodological mistakes in the advancement of knowledge have not been due to a lack of naturalism for a long time now. For naturalism is rightly taken for granted in any decent scientific endeavour. They are due to our being enchanted by the affordances provided by the data—the sort of information that the data invite us to transform them into—and being forgetful or indeed wilfully negligent towards the constraints provided by the same data, the sort of information that the data intractably resist to become. On this I believe Lakatos was correct. And this seems to me a crucial and permanent lesson we must learn from Popper’s falsificationism.

Naturalism is a user’s philosophy: it endorses a primacy of ontology (the description of the system) over epistemology (the elaboration of the model of the system) and a representationalist/correspondentist interpretation of the knowledge of the system. It is almost irresistible in its intuitive and common-sensical features. Non-naturalism endorses a primacy of epistemology over ontology and a constructionist/correctness-based interpretation of knowledge. It is counter-intuitive and not in line with common-sense. Yet so are many truths that we have come to accept in science and mathematics. In a world that is increasingly dominated by technology, normative constructs and conventions, artificial products, engineered environments, design, maker’s knowledge and mental poiesis of all kinds, strong naturalism seems to be a mere long-term consequence of the modern scientific revolution, a hang-over from modern empiricism. It is what our mammalian brain cannot help assume. But it can no longer provide an appropriate way of approaching our responsibilities, opportunities, and limits as careful and respectful (recall the constraints) creators of our understanding of the world, ourselves within it, and of the ways in which we shape both. Constructivism is what our human mind is rather suggesting to accept. The conclusion is that it is non-naturalism, interpreted as a realist constructionism, that can claim the best motivations today. This is why elsewhere I have defended a foundationalist and non-naturalist interpretation of the philosophy of information (Floridi 2012b). To be human is to refuse to accept the natural as naturalistic, and to take full responsibility for such a refusal.

Acknowledgements I discussed several versions of this article on many occasions and I am deeply indebted to many colleagues and audience members for their feedback. In particular, the following meetings were especially useful for improving the final version: a presentation, discussed during the Fourth Philosophy of Information Workshop (University of Hertfordshire, 2012), and then the following invited talks and keynotes given at the Center for the Study of Interdisciplinarity, (University of North Texas, 2012), at the Philosophy of Education Research Seminar (Institute of Education, London 2012), at the Department of Philosophy of the University of Uppsala (2013), and at the Department of Philosophy of the University of Milan (2016). I am in debt to the organisers for giving me the opportunity to gather so much valuable feedback. I am grateful to David Sutcliffe for his editorial revisions. Finally, I am grateful to the two anonymous referees for their very helpful and constructive comments, which enabled me to improve the article significantly.

References

- Blackburn, S. (2016). *The oxford dictionary of philosophy* (3rd ed.). Oxford: Oxford University Press.
- Bogen, J., & Woodward, J. (1988). Saving the phenomena. *The Philosophical Review*, 97(3), 303–352.
- Bridges, J. (2007). Review of Richard Gaskin, “Experience and the world’s own language: A critique of John McDowell’s empiricism”. *Notre Dame Philosophical Reviews*, 2007(2).
- D’Agostino, M., & Floridi, L. (2009). The enduring scandal of deduction. *Synthese*, 167(2), 271–315.
- De Caro, M., & Macarthur, D. (2010). *Naturalism and normativity*. New York: Columbia University Press.
- Floridi, L. (1996). *Scepticism and the foundation of epistemology: A study in the metalogical fallacies*. Leiden: Brill.
- Floridi, L. (2008). The method of levels of abstraction. *Minds and Machines*, 18(3), 303–329.
- Floridi, L. (2010a). *The philosophy of information*. Oxford: Oxford University Press.
- Floridi, L. (2010b). Semantic information and the correctness theory of truth. *Erkenntnis*, 74(2), 147–175.
- Floridi, L. (2011). A defence of constructionism: Philosophy as conceptual engineering. *Metaphilosophy*, 42(3), 282–304.
- Floridi, L. (2012a). Semantic information and the network theory of account. *Synthese*, 184(3), 431–454.
- Floridi, L. (2012b). *The road to the philosophy of information*. Chapter in Luciano Floridi’s *Philosophy of Technology* (pp. 245–271). Netherlands: Springer.
- Floridi, L. (2013a). What is a philosophical question? *Metaphilosophy*, 44(3), 195–221.
- Floridi, L. (2013b). Distributed morality in an information society. *Science and Engineering Ethics*, 19(3), 727–743.
- Floridi, L. (2014). Perception and testimony as data providers. *Logique et Analyse*, 57(226), 151–179.
- Floridi, L. (forthcoming). ‘What the maker’s knowledge could be’, *Synthese*.
- Gaskin, R. (2006). *Experience and the world’s own language: A critique of John McDowell’s empiricism*. Oxford: Oxford University Press.
- Maddy, P. (2007). *Second philosophy: A naturalistic method*. Oxford: Oxford University Press.
- McDowell, J. (1996). *Mind and world: With a new introduction*. Cambridge: Harvard University Press.
- Nobre, A. C., Correa, A., & Coull, J. T. (2007). The hazards of time. *Current Opinion in Neurobiology*, 17, 465–470.
- Papineau, D. (2015). ‘Naturalism’, *Stanford encyclopedia of philosophy*, available online.
- Quine, W. V. (1969). ‘*Epistemology naturalized*’, *ontological relativity and other essays* (pp. 69–90). New York: Columbia University Press.
- Quine, W. V. (1981). *Theories and things*. Cambridge: Belknap Press of Harvard University Press.
- Reynaert, P. (2015). Does naturalism commit a category mistake? *Bulletin d’Analyse Phénoménologique*, 11, 3.
- Ritchie, J. (2008). *Understanding naturalism*. Stocksfield: Acumen.
- Rysiew, P. (2016). ‘Naturalism in Epistemology’, *Stanford Encyclopedia of Philosophy*, available online.
- Strandberg, C. (2004). In defence of the open question argument. *Journal of Ethics*, 8(2), 179–196.
- Stroud, B. (1981). The significance of naturalized epistemology. *Midwest Studies in Philosophy*, 6(1), 455–472.
- Stroud, B. (1984). *The significance of philosophical scepticism*. Oxford: Clarendon Press.