A Defence of Constructionism: Philosophy as Conceptual Engineering

Luciano Floridi[1, 2, 3]

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

Abstract

The article offers a broad account and a defence of *constructionism*, both as a metaphilosophical approach and as a philosophical methodology, with some references to the philosophical tradition that has inspired it, the so-called "maker's knowledge" tradition. The main thesis defended is that Plato's "user's knowledge tradition" should be complemented, if not replaced, by a constructionist approach to philosophical problems in general and to knowledge in particular. To put it simply, an epistemic agent knows something when that agent is able to build (reproduce, simulate, model, construct etc.) that something and plug the obtained information in the correct network of relations that account for it. Or in even more intuitive terms, an agent qualifies as an epistemic agent not when she is a passive user of some information, but when she is a critical producer of it. Her epistemic expertise increases in relation to the scope and depth of the questions that is able to ask and answer on a particular topic. The maker's knowledge is knowledge of the ontology of the semantic artefact and this is a fundamental epistemological lesson we can learn from poietic disciplines such as computer science and economics. So constructionism shifts the focus away from the *mimetic, passive and declarative knowledge that* something is the case, in order to concentrate more on the *poietic, interactive and practical knowledge of* something being the case, that is, of semantic artefacts. Once applied to the interpretation of philosophy itself, constructionism suggests adding conceptual engineering to conceptual analysis as a fundamental method.

Keywords

Constructionism, constructivism, epistemology, levels of abstraction, minimalism, philosophy of information.

Introduction: from the User's Knowledge to the Maker's Knowledge Approach

There are critical crossroads, in the history of philosophy, when a small conceptual step in one direction has immense consequences for the kind of further investigations it encourages, and the potential alternatives it obliterates. Plato took one of those influential steps when he endorsed the distinction between *episteme* and *techne*, and insisted on grounding our understanding of human knowledge on a user-oriented approach, favouring a passive and mimetic reception of semantic information. That crucial step ended up influencing twenty-five centuries of epistemological work.

As we shall see, it would be incorrect to suggest that nobody ever questioned "the user's knowledge" approach, what I shall call the *Platonic dogma*. One may even argue that the historical Plato himself never quite endorsed the Platonic dogma in the way I shall describe below. Perhaps, but it is the Plato that we find in the history books that I have in mind. And along the same line, it is fair to state that, in the same books, the Aristotelian tradition *partially* (on this qualification see below) reinforced the Platonic dogma – recall the piece of wax passively taking on the impress of a signet-ring – while Christian philosophy never radically challenged it, attributing to God, like Plato, the only creative intellect in the universe. Alternative views, clustered under the "maker's knowledge" approach, amounted to a minority report at most; witness contemporary epistemology textbooks, which hardly even mention it. Nowadays, despite some attempts to promote a more nuanced, if not an opposing view about the interactive, constructive and creative (henceforth *poietic*) origins and nature of our knowledge of the world, we still find our academic culture and its philosophers accepting, by default, the priority of knowledge-that over knowledge-of and knowledge-how, of theory over practice, of thinking over doing, of representations as copies and reproductions over representations as models and constructs.

The difficulty lies partly in the fact that the Platonic dogma has its strongest ally in the unquestionable success of our commonsensical dealings with everyday reality – what might be called the "What You See Is What You Get" experience of the world – partly in the fact that, together, the Platonic dogma and

the commonsensical naïve realism shape the very framework within which any further reflection on the nature of knowledge is critically developed.

Given the deep entrenchment and far-reaching roots of the Platonic dogma, pessimism with regards to its overturning may seem fully justified. And yet, some room for hope is offered by the increasing discrepancy between Plato's epistemological orthodoxy and many of our actual epistemic practices. The scientific revolution was made possible by the abstraction and mathematisation of phenomena, and the pervasive technologisation of science: *episteme* and *techne* may not have entered into philosophical marriage yet, but they have been bedfellows for centuries now. The information societies in which we live are neo-manufacturing systems in which the expert and intelligent handling of data and information is the primary value-adding occupation of the majority of the working population (Floridi 2010). Any child who learnt by doing, any person aware of the fact that understanding requires much more than passive observation, any student trained in a lab or in a field, any engineer who ever designed an artefact, any scientist who ever ran an experiment or devised a simulation, any user who ever felt the need to know more about a technology than just how to enjoy it, any academic who has realised that teaching is a great way of learning, and, in general, anyone involved in the business of information creation, refinement, transmission and acquisition must have perceived, at some point, that our ever richer insights into the nature of reality have their foundation in our practical and creative interactions with it. We do things with information, to paraphrase Austin (Austin 1962). So it is because we *know how* to do this and that creatively, interactively and collectively that we can rightly report to *know that* such and such is the case passively and individually. Propositional knowledge is the glorious conclusion of the informational process, not its humble beginning.

The intellectual divide between epistemic practices and epistemological theorization is widening, not least because the pressure coming from the immense amount of information we are exponentially generating is pushing us towards developing new techniques and technologies to handle them, and new epistemic ways to exploit them. We should be very suspicious of the increasing pedagogical insignificance of our epistemology based on the user's knowledge

approach.[1] It seems clear that, if epistemology wishes to be more than just an intellectual game, utterly divorced from the actual dynamics of human knowledge, and hence irrelevant to its understanding and furtherance, it must retrace Plato's steps, and take a different direction. The risk is to foster an epistemological culture of passive information receivers and consumers rather than of critical and proactive information producers. The time has come to be epistemologically heretical, to abandon a passive, mimetic, user-oriented perspective as to how we generate our knowledge of the world and join forces with some of the less orthodox thinkers in our philosophical tradition, in favour of a maker-oriented approach. Knowledge is not about getting the message from the world; it is first and foremost about negotiating the right sort of communication with it.

It would be easy to be radical, yet such temptation would also be naïve. For it would mean failing to appreciate the true difficulty of our task. Merely moving in the direction opposite to the one taken by Plato and classic epistemology after him would be difficult but simple. This is the unsuccessful strategy adopted by Romantic idealism and post-modern *constructivism* (with a *v*, not an *on*). Both have ended up as friendly fire for the sort of *constructionism* I suggest we should develop. For if the choice is mistakenly reduced to be between denying

i)   either any poietic value to our epistemic activities,

ii)   or any existence, or at least ontological independence, to the external world,

then you only need to miss the train once, or spill some coffee in the morning, to regain some sanity and make up your mind about who is right and who must be wrong. Various forms of constructivism are nothing more than *haute couture*, made possible precisely by that excessive reliance on the propositionalization of our knowledge of the world that they then seek to criticise. They are innocuous parasites that live in a symbiotic relation with the culture they criticise but on which they entirely depend for their survival. They are the best enemy that a

---

[1] The philosophical constructionism defended in this paper goes hand in hand with constructionist learning theory (Harel and Papert 1991), according to which learners understand the world by building mental models, and that this practice is facilitated by the actual handling of physical objects (see also Piaget's experiential learning theory).

conservative, user-oriented epistemology may wish to have, for they pose no credible challenge, since their position is even less tenable than the Platonic dogma they seek to overthrow.

So our difficulty is complex, because it consists in being radically moderate: we need to identify and follow the middle course, represented by the *design* of the world. This hardly thrills young minds, smacks of compromise to older ones and, worst of all, cannot escape the constant risk of being confused with either Scylla or Charybdis, *discovery* or *invention*. During an arm wrestling event, you neither expect nor look forward to a balanced draw between the two wrestlers. That is not the exciting outcome of the competition. Likewise, in philosophy we seek clear-cut alternatives and resolutions. In epistemology, we wish either knowledge or reality to pin the other's arm onto the surface of a conclusive answer over primacy once and for all. We can hardly hold firm the view that *constructionism* is neither *realism* nor *constructivism*, because knowledge neither *describes* nor *prescribes* how the world is, but *inscribes* it. The powerful and intuitive polarization between realism and idealism, discovery and invention, naturalism and anti-naturalism, represents the intellectual temptation that makes our difficult task so delicate. For we need to ensure that what we appreciate as the highest and most cherished form of knowledge – our increasingly successful capacity to capture the world in a network of propositional contents and relevant accounts – maintains its deserved high status, keeps flourishing and remains open to further progress, while replacing the foundations that make it so reliable. We wish to be friendlier to truth, without becoming enemies to Plato. With a political analogy, our revolution will not consist in dethroning classic epistemology as the philosophical Queen but in transforming her kingdom into a constitutional monarchy. Or think of the difficult but simple task of removing, piece by piece, a fragile dinner service from a table, compared to the difficult and very complex task of pulling a tablecloth from under the same dinner service at once. At the end of the process, if we are successful, what will have changed is what the items are placed on, not their positions.

A change as radical and yet complex and delicate as the one outlined above is not achievable by an individual, let alone by a single article. So, in the

following pages, I intend only to contribute to what I hope may be a larger movement, in which I enlist Kant's transcendental epistemology and Peirce's pragmatism as being among its best expressions. I shall do so by gathering supporters and defending a constructionist philosophy of information. Here is the line of reasoning in synthesis.

In section one, I shall describe briefly the Platonic step, why it was not a necessary step, and one that could have been recovered. The Platonic dogma left us with some problems, and in section two I shall argue that the maker's knowledge tradition, from Bacon to the philosophy of information, may provide a viable alternative to their solution. The rest of the article is devoted to outlining a constructionist philosophy (section three), by discussing minimalist principles (section four), the method of levels of abstraction (section five), and constructionism itself (section six). The article ends with a brief conclusion (section seven).

1. Plato's Wrong Step

Who really knows an artefact? Plato addresses this question several times in his dialogues. His answers are *largely* (more on this qualification later) consistent, despite being provided at different stages of his intellectual development.

In the *Cratylus*, a dialogue traditionally attributed to the so-called transitional or middle period, Plato argues (390b-d) that the *user* of an artefact (in this case a weaver using a shuttle) knows it better than its *maker* (the carpenter who makes the shuttle). The argument is at best controversial. Allegedly, it is the user who knows whether the proper form (the blueprint) is correctly given to the artefact. Today, we would say that it is up to the user to determine the (or at least up to the maker to guess the user's) functional requirements and the usability features of an artefact first and then assess how far the actual product satisfies both. Yet, why the maker is supposed to lack (and indeed the user possess) such knowledge of the form remains unexplained, even in the best scenario in which Plato might be talking only of the functional requirements and usability features. Equally puzzling is how the maker could actually produce a good artefact in the first place without a solid grasp of its blueprint.

A later and more famous statement of the "user's knowledge" dogma fails to be more convincing. The first half of the tenth book of the *Republic* is dedicated to the issue. At the beginning, we find the famous argument of the three couches (597b): the ideal couch is produced by God, the physical couch is made by the carpenter, and then the "fake" one results from the painter's imitation. The imitator gets bashed, as expected, but the interesting point here is to note that God has intrinsic knowledge of the couch because he is the source of its blueprint, the designer of the perfect exemplar. Intriguingly, no user is mentioned. Later, however, the tripartite distinction shifts into the familiar pattern we have already encountered in the *Cratylus*. In 601c-d, the new example involves reins and bit; the painter fails to know their real nature, since he merely imitates them, and the makers, the cobbler and the smith, have no real knowledge of them but only justified beliefs about their nature, a lower kind of knowledge. The genuine knower is the horseman, who understands how to use them and is truly acquainted with their nature. Plato uses the three-way distinction between *imitating*, *making*, and *using*, in order to argue now (all quotations are from (Plato 1989):

> That there are some three arts concerned with everything, the user's art, the maker's, and the imitator's. [...] [and] the user of everything is the one who knows most of it by experience, and that he reports to the maker the good or bad effects in use of the thing he uses. (601-d).

The user now provides the bench-mark in the artefact's evaluation, in view of functional requirements and usability features. The maker is merely reduced to a first-class imitator, the imitator of God's creations. Plato has forgotten the real maker, God the Ur-maker, who is the master of a fourth, poietic art, not mentioned. The user becomes the judge, who evaluates how closely the human maker is able to reproduce the blueprint provided by the divine maker.

Historically, Plato might have been motivated by cultural and social biases to degrade the maker's knowledge to mere true belief. Perhaps the artisan must be seen as a mere living tool, a slave, in the hands of the user ("the work of the carpenter is to make a rudder, and the pilot has to direct him, if the rudder is to be well made", *Cratylus*, 390d).

Plato might have been unaware of such biases: for he seems to write as if it were utterly obvious and uncontroversial that an artisan could never possibly qualify as a knower with respect to its artefacts. Philosophically, it is important to note that, in both dialogues, the primary targets of Plato's arguments are not the makers (the legislators in the *Cratylus*, and the artisans in the *Republic*), but the imitators (the sophists and the artists, respectively). In both cases, the dismissal of the maker's knowledge as second-rate appears to be collateral damage in a battle fought against more dangerous and worthier foes. That this is so becomes clearer once we look at another transitional or middle period dialogue, the *Euthydemus*.

In the *Euthydemus*, Plato defends a more complex interpretation of the relationship between maker's and user's knowledge, not entirely coherent in itself, as we shall see presently, and less consistent with the conclusions reached in the other two dialogues, but much more interesting. Initially, the basic argument looks similar. However, perhaps because the context is not immediately represented by the controversy against potentially competing sources of knowledge such as rhetoric (the sophists) and traditional culture (poets and artists), Plato seems willing to endorse for a moment a more balanced view. Here are the important passages in question:

> [288e-289a] […] even if we knew how to turn all the stones into gold, the knowledge would be worth nothing to us; for unless we know *also* how to use the gold, we saw there was no benefit in it. […] In the same way, it seems there is no benefit in any other knowledge, of business or of physic for example, or anything else which *knows how to make something but not how to use what it makes* (ποιεῖν τι ἐπίσταται, χρῆσθαι δὲ μὴ ᾧ ἂν ποιήσῃ). […] [289b] Even if there is a knowledge how to make men immortal, without the knowledge how to use immortality even this seems to bring no benefit. […] Then […] we need such a knowledge (ἐπιστήμης) as *combines* both how to make something (ποιεῖν) and how to use what is made.

The Greek word translated as *combines* is συμπέπτωκεν, literally "falls along with". The conciliatory tone could not be more explicit: user's and maker's knowledge are complementary types of know-how, which must be joined

together in order to reach full and useful *episteme*. This is the crucial crossroad where Plato could have taken a different path. He does not, because he immediately recalls the specific makers he wishes to attack, the speechmakers, and more generally the sophists. And like someone who, running away from danger, decides to take not the path that leads in the right direction, which he acknowledges to be the closest, but rather the path which is diametrically opposite to where the danger emanates, out of fear rather than reason, so Plato cannot resist the temptation to revise the relationship between the two complementary kinds of knowledge into an opposition, as soon as the old foes reappear. He simply overreacts and, quite astonishingly, in the following passage, which comes immediately after the one quoted above, he completely undermines what he just defended:

> [289b] Then we must be *nothing at all* like those who make harps, for instance; we do *not want to become masters of knowledge like that*; [289c] for there the making art is one thing, the using art quite another, and each art deals separately with the harp. For the arts of harp making and harp playing are very different.

One would have expected an appeal to multi-disciplinarity, to use a buzzword of our times. Plato could have taken the direction of a more complex education of the philosopher both as a maker and as a user, as a knowledge produmer (producer and consumer). Instead, he prefers to contradict the complementary thesis he just supported and establish a dichotomy between the user's superior and the maker's inferior knowledge of an artefact. In this way, he can build a stark contrast between the good and the bad guys, the philosophers and the sophists. The step is taken; maker's and user's knowledge, *techne* and *episteme*, and hence practical and theoretical knowledge start moving apart; the dogma has entered our epistemology. It immediately shows its power: a few lines after the passage quoted above, Plato compares hunters to makers, who do not know well enough what they have hunted and must hand it over to the cooks, the users; he tells us that generals, like hunters, capture cities and armies, but they must then hand them over to politicians, the users who know better, and finally, he writes

[290c] Geometers and astronomers and calculators – for these are sort of hunters too, since they are not mere makers of diagrams, but they try to find out the real meanings – so because they do not know how to use them, but only how to hunt, they hand over their discoveries, I take it, to the dialecticians to use up, at least all of them hand over who are not quite without sense.

Something seems to have gone badly awry. Some sort of ideology of power and control, together with the philosophical agenda to cast the highest level of discredit on those wordsmiths and makers of discourses that are the poets and the sophists, push Plato in a direction that is practically untenable and philosophically implausible. Who, if not a philosopher, could believe and argue that the maker of an artefact knows it less well than its user? Try that next time your car breaks down and needs to be repaired. Plato was right in stressing the importance of both kinds of knowledge, and he was wrong (and for the wrong reasons) to argue that the user's knowledge should be preferred to the maker's.

The last quotation from the *Euthydemus* introduces a further complication. Geometers, astronomers and calculators are makers of a special kind of artefact. Today, they would be called knowledge workers, who produce and manage information. This is something that we, as epistemic agents, create, share and transmit among ourselves and across generations, and accumulate and refine through time. It is difficult to follow Plato when he holds that users know their iPhone better than Apple, at least without adding several qualifications. And it is impossible to agree with him when Wikipedia users are said to know the information they access better than those who actually generated it in the first place. This is a *reductio*. Unless what we are supposed to have in mind is the equation producer = reproducer; that is, unless we hold the view, as Plato seems to, that the ultimate knowledge of things is something we can at most access (reminiscence), but that we do not build. If there is an ultimate, ideal Wikipedia, which producers of the entries of the actual, human Wikipedia are only gradually approximating, then those producers do not have better knowledge of the contents they are conveying, no more than a policeman knows a crime scene better than the criminal. What counts is the exposure to the truth and how you handle it (user's knowledge) not how you construct it in the first place (maker's

knowledge). Geometers, astronomers and calculators are not co-workers in the scientific enterprise, but mere communication channels through which pre-existing information reaches the dialecticians. Research is reduced to uncovering and reporting. This is disappointing not least because it prevents understanding the contraposition between dialectic (διαλέγεσθαι) and eristic (ἐρίζειν) in terms of two procedures to engineer semantic artefacts, to be judged on the quality of their products. If sophists and ignoramuses are mere imitators, like painters, and if philosophers or scientists are not the makers with whom we should compare them, then one may argue that none of them is closer to knowing the genuine nature of the semantic artefacts in question. Charlatans of all sorts, in all ages and societies, can feed irresponsibly human cultures with irrationalism, obscurity, dogmatism and relativism, at least by claiming to be no worse than their opponents. Such charlatans are disgraceful not because they sell us conceptual replica – as Plato argues – but because they sell us conceptual lemons. Since in Plato it seems that

i) all that matters is how you use some information and assess it against the actual, genuine referent,

then the user's knowledge dogma further requires

ii) a form of external realism about information (the ideal Wikipedia out there) and

iii) some theory explaining our potentially noiseless access to it (recollection is perfectly suited).

Very coherently, Plato defended all three positions. Yet each of them seems to be highly questionable and has been subject to endless criticisms. Is there an alternative to a Platonic philosophy of information? How can information, as a conceptual/semantic artefact, be (better) known by its makers?

2. The Maker's Knowledge Tradition

The alternative view is a constructionist philosophy of information. This is based on the maker's knowledge tradition, which in turn has its roots in some strands of the Aristotelian-Scholastic philosophy. If genuine knowledge is knowledge of the intrinsic nature of the object known (knowledge of the ontology of the known), and if there is no innate acquisition of such blueprint, then knowing a

phenomenon, an artefact or, in our case, gaining information and being able to account for it, means being able to produce it and reproduce it, to assemble and disassemble it, to build and dismantle it, to improve it and answers questions about it, and all this for the right reasons. For knowledge to be possible, *mimesis* must be replaced by *poiesis*. A similar thesis can be found in the Aristotelian-Scholastic tradition, which holds that knowing something means having a full understanding of the causes that bring about the known. From a Christian but non-Greek perspective, God the Ur-Maker is the only artisan that creates *ex nihilo*. This is how God the maker may be seen to be omniscient – recall how Plato himself had to acknowledge the superior nature of the Ur-maker's knowledge in the *Republic* – and how humanity, by partaking in God's nature, may have some epistemic access to the created universe. Nowadays, as I have argued elsewhere ((Floridi forthcoming; Floridi 2011)), we may be better off by analysing knowledge in terms of holding the information that such and such is the case and being able to provide the correct account of why this is so. The user's knowledge approach can easily be atemporal and single-agent: it does not really matter how many people are looking at the wall, or how many get out of the cave. The maker's knowledge approach, on the contrary, requires trial and error and time-consuming processes of accumulation, so knowledge becomes a collaborative enterprise of growth and refinements in a multi-agent system (humanity) across generations. It is based on man-hours. As Aristotle writes in *Metaphysics* (II,[α], 1, 993b, 1-5):

> [...] each thinker makes some statement about the natural world, and as an individual contributes little or nothing to the inquiry; but a combination of all conjectures results in something considerable.

Of course, the Aristotelian-Scholastic tradition does not hold a constructionist epistemology yet, for the view of knowing as a passive process of discovery and acknowledgement of the state of the world in itself is still predominant, witness the very history of the word "information" ((Floridi 1994)). But when Francis Bacon writes, in the *Novum Organum*, that "Vere scire, esse per causas scire" (To know truly is to know through causes) he is more than just paraphrasing an Aristotelian maxim: he is pointing us in a new, anti-Platonic direction, according to which we, as epistemic agents, can only know what we make as Ur-makers.

This is a major transformation in perspective brought about by Bacon and later on by the scientific revolution. *Constructionism* holds that knowledge is acquired through the creation of the right sort of *semantic artefacts*, information modelling, in other words. We are the builders of the infosphere we inhabit, Bacon's "intellectual globe". Thus, both the philosophy of language and the philosophy of mind should be at least equally concerned with creative rather than reproductive capacities. The obsessive emphasis on mimetic representations and propositional knowledge is misplaced. We do not and cannot gain knowledge by passively recording reality in declarative sentences, as if we were baskets ready to be filled; instead we must handle it interactively.

A constructionist philosophy might be cause for concern. Even from a historical point of view, the maker's knowledge tradition has often been suspected of sceptical sympathies. Kant himself was accused of being a sceptic. So Hintikka was right in arguing that

> [I]n many respects pessimistic conclusions [from the *verum ipsum factum* thesis] should have been as close at hand in Bacon's time as the optimistic ones. […] It is a sobering thought that [the maker's knowledge approach] the leading idea of the intellectual background of modern applied science and scientific technology can be traced back to sceptical and theological principles calculated to extol the superiority of the Divine practical reason over the human one. ((Hintikka 1974), pp. 85-6).

Yet such concern overlooks the fundamental task of modern constructionism, which is that of soldering together the Platonic dichotomy between human and divine making and hence between the two different epistemological "arts", by reinterpreting the dualism between human and divine knowledge in ontological terms, between *noumena* and *phenomena*, between the reality of data in themselves and the world of information as we know it.

This is not scepticism, and it is not relativism either. It would be difficult to overstress how realistically oriented such constructivism is. Since the beginning, it has been based on a new understanding of technological production and control (e.g., in Bacon) and perceptual cognition (e.g., in Locke). Recall the "false friends" mentioned in the previous section. Bacon was certainly no idealist. Locke has been interpreted –rightly in my view – as a supporter of a

maker's knowledge epistemology ((Tully 1980); (Jacovides 2002)). And I am certainly not advocating a *constructivist* epistemology. But just in case some readers were still suspicious about the empirical pedigree of the maker's knowledge approach, and hence of the corresponding *constructionist* philosophy of information, here is how Hobbes aptly sums up the view that the only true knowledge is knowledge possessed by the maker of his creation:

> Of arts some are demonstrable, others indemonstrable; and demonstrable are those construction of the subject whereof is in the power of the artist himself, who, in his demonstration, does no more but deduce the consequences of his own operation. The reason whereof is this, that the science of every subject is derived from a precognition of the causes, generation, and construction of the same; and consequently where the causes are known, there is place for demonstration, but not where the causes are to seek for. Geometry therefore is demonstrable, for the lines and figures from which we reason are drawn and described by ourselves; and civil philosophy is demonstrable, because we make the commonwealth ourselves. But because of natural bodies we know not the construction, but seek it from the effects, there lies no demonstration of what the causes be we seek for, but only of what they may be.[2]

Hobbes still lacks a view of science as a poietic activity, and of perception as a way of interpreting data, but he is on the right track: in the mathematical as well as in the human sciences we can acquire full knowledge because we can come to know the causes of our objects of investigation. After all, you know a theorem if you know how to prove it, not if you can merely state its result, and experiments do not imitate the world, they shape it in controllable ways, as Galileo understood.

So far the new theoretical outlook. We have a new direction towards which *constructionism* might develop. The question is how. One way to answer this question would be by embarking on a thorough historical exploration.[3] Unearthing the various intellectual stages through which the maker's knowledge tradition developed, as one of the important epistemological undercurrents of

---

[2] (Hobbes 1656), Epistle Dedicatory; EW, VIII, 183-4.
[3] On the history of the knowledge maker's tradition see (Child 1953) and (Pérez-Ramos 1988). I read (Rorty 2009) as a criticism of the user's tradition.

modern philosophy, into a contemporary constructionist movement, would be both fascinating and enlightening. Yet this is not what I intend to do in the rest of this paper. Another alternative would be to show how empirical knowledge itself develops through the constructions of our information about the world. This is how I interpret Kant's epistemological project, and in (Floridi 2011) I have tried to show how Kant's transcendentalism might be combined with a constructionist philosophy of information. Here, I wish to keep the dialogue with Plato open, and hence work at the same level of reflection: if the user's knowledge approach is joined by that of a maker's knowledge, as a viable way of making sense of human knowledge, can it be applied to itself, and hence be used to make sense of a constructionist methodology? In other words, could the maker's knowledge approach be coherently applied to the very theorization of its value? Can we make sense of the possibility that it might be constructionism all the way up, and hence include philosophy as well? The short answer is yes, by interpreting philosophy in general and the philosophy of information in particular as conceptual engineering. The long answer will keep us busy for the rest of this article.

3. Philosophy as Conceptual Engineering

Too much ink has been spilt on philosophy as conceptual analysis. The alternative view, that philosophy is at least as much, if not actually more, engaged with creating, refining and fitting together our conceptual artefacts in order to answer open questions, i.e. questions that are not answerable in principle empirically or mathematically,[4] has received too little attention.[5] So, much work lies ahead, and what I shall highlight here is no more than a contribution to such a large project. I shall divide the overview into three parts: MINIMALISM, the METHOD OF LEVELS OF ABSTRACTION (LoAs) and CONSTRUCTIONISM. Together, they flesh out a constructionist methodology, i.e., a philosophy that takes seriously the view that the maker's knowledge is the right approach from which to interpret all expressions of human knowledge, from our empirical interactions with the world to the self-reflective interpretation of our

---

[4] I discuss the nature of open questions and of erotetic foundationalism in (Floridi 2011).
[5] For an interesting exception, of course coming from two non-analytic philosophers, see (Deleuze and Guattari 1994).

epistemology. This means that I take MINIMALISM, the METHOD OF LEVELS OF ABSTRACTION and CONSTRUCTIONISM as providing the core of a constructionist philosophy of information. However, I am also aware that other philosophers of information may wish to endorse only MINIMALISM and the METHOD OF LEVELS OF ABSTRACTION, but not a full constructionist philosophy. I suspect Dretske might be one of them.[6] For this reason, I shall restrict the use of that label to refer only to the third part, instead of the whole. If you are a constructionist, you are likely to be better served in your philosophy by a minimalist approach and the method of LoAs. But you might wish to help yourself with the latter two tools and stop short of accepting a maker's knowledge perspective.

4. Minimalism

MINIMALISM is an erotetic principle. The view is simple: philosophical questions often pose multi-faceted problems, which can form a problem space, which, in turn, following Descartes, can be decomposed and dealt with by means of a divide-and-conquer approach (Raftopoulos 2003). The outcome is a set of more approachable sub-problems, interconnected in a sort of Quinean web of dependencies (Quine 1951). When dealing with a philosophical question, the starting problem often presupposes other open problems. So the strength of the answer depends on the strength of the corresponding assumptions. A minimalist starting problem relies as little as possible on other open problems, thereby strengthening the final answer to the philosophical question. Very often, one may improve the tractability of a problem space by choosing a model with which to study it. MINIMALISM outlines three criteria to orientate this choice: *controllability*, *implementability* and *predictability*.

A model is controllable when its features can be modified purposefully. Given this flexibility, the model can be used as a case study to test different solutions for the problem space.

The second criterion recommends that models be implementable, usually through the description of conceptual mechanisms (e.g., thought experiments, analogies, logic constructs, ideal models, counterexamples, etc.), sometimes

---

[6] I have learnt much from Dretske in general, and I agree with him in several ways, but I believe that, because of (Dretske 1994), he would not welcome the sort of synthetic (as opposed to naturalised) constructionism I defend in this article.

through virtual simulations (also known as experiments *in silico* (Grim 2003)), and seldom through physical realizations. The more transparent a model becomes, that is, the more its mechanism and dynamics move from making it a black to making it a white box, the better. Metaphorically, the maker of the model is a Platonic "demiurge", fully cognisant of its components and of its state transition rules. The model can therefore be used as a conceptual laboratory to test specific constraints on the problem space.

The third criterion follows from the previous two: ideally, the chosen model should be such that its behaviour should be predictable, at least in principle. The demiurge can predict the model's behaviour in that she can infer the correct consequences from her explanations of the model. The model outcomes become then the benchmarks of the tested solutions.

Given the previous three criteria, three properties further characterise MINIMALISM as I am advocating it here. First, MINIMALISM is *relational*. Problems and models are never absolutely minimalist, but always connected with the problem space posed by the philosophical question. Such relational nature percolates through the approach, making it easier to understand that, often, a conceptual difficulty lies in the innate temptation to see only Boolean solutions. I have already referred to this tendency above, when mentioning the difficult task of steering a middle course between extremes that work as very powerful attractors. For example, as soon as a particular issue is framed in terms of internalism vs. externalism, we lose the possibility of opting for a third, *liminalist* solution. Imagine being forced to define the concept of "food" either in externalist terms, only as physical nutrients in the world, or in internalist terms, only as dietary requirements for a specific biological species. The debate could easily be endless, since food is a relational concept. Recall my early warning: relationalism or, as I prefer to call it, *liminalism*, is most emphatically *not* to be confused with *relativism*. It is simply false to say that something is food independently of the nature of the feeder, but it is also obviously false to say that, therefore, it just depends on the eater, and hence that anything might qualify as food. Relativists should be made to eat their hats. Semantic information is another useful example, for it is neither only in the environment nor only in the mind, but arises from the interactions of specific agents within their

environment. The relational nature of MINIMALISM further clarifies the need to *triangulate* our basic concepts: it is often easier to start with a Boolean dichotomy, but it is equally often essential to treat the dichotomy as providing, as in geometry, the two known points at either end of a distinction which can help to determine the third point that provides the right perspective on the issue under examination. The reader might perceive rightly in this traces of a Hegelian logic.

Second, MINIMALISM provides a way to choose critically the most fruitful starting problem for the analysis of a problem space, thus enhancing the strength of the next step forward in the process of answering the philosophical question. According to a minimalist approach, the tractability of a philosophical problem is a function of the three criteria outlined above. They allow the use of dynamic models to test possible solutions and to derive properties of the problem space.

Finally, MINIMALISM is a matter of inferential relations between a problem and its space, but it is not a way to privilege simple or elementary problems. Minimalist problems may be difficult or complex. MINIMALISM is an economic method related to, but not to be confused with, Ockham's razor. The two methods are of course compatible and indeed complementary. However, whilst Ockham's razor increases clarity and elegance and avoids inconsistencies and ambiguities by eliminating redundant, explicative or ontological elements in a theory, MINIMALISM offers criteria for choosing problems and models relative to a given specific question. Moreover, Ockham's principle of parsimony is absolute and is applied to any theoretical element, while MINIMALISM's main maxims of strength and tractability are always relative to a given problem space. Thus, a quick and dirty way of defining MINIMALISM is by describing it as the erotetic version of Ockham's razor.

The definition of MINIMALISM is based on two main assumptions. One concerns the existence of a problem space. MINIMALISM does not give an account of the decomposition process of the problem space to which it applies. The other is that a dynamic model, whether conceptual, virtual or physical, is useful in finding the answer to the investigated question. MINIMALISM does not explain why and how this approach works. In the next sections, two other methods are presented to ground both assumptions: the METHOD OF LEVELS OF ABSTRACTION and

CONSTRUCTIONISM. The METHOD OF LEVELS OF ABSTRACTION is used to describe the observables that compose the problem space and how they are related. MINIMALISM is always relative to a given Level of Abstraction as the structure of the problem space depends on the LoA assumed by the investigator. CONSTRUCTIONISM clarifies how to devise the model and how to use it in order to investigate the set of minimalist problems that will lead to the required answer.

5. The Method of Levels of Abstraction

The METHOD OF LEVELS OF ABSTRACTION (henceforth METHOD OF ABSTRACTION) has been formalised in (Floridi 2008). The terminology has been influenced by an area of Computer Science, called *Formal Methods*, in which discrete mathematics is used to specify and analyse the behaviour of information systems (Zeigler 1976). Despite that heritage, the idea is not at all technical and for the purposes of this paper no mathematics is required, for only the basic idea will be outlined.

Let us begin with an everyday example. Suppose we join Anne (A), Ben (B) and Carole (C) in the middle of a conversation. Anne is a collector and potential buyer; Ben tinkers in his spare time; and Carole is an economist. We do not know the object of their conversation, but we are able to hear this much:

A. **A**nne observes that it has an anti-theft device installed, is kept garaged when not in use and has had only a single owner;

B. **B**en observes that its engine is not the original one, that its body has been recently re-painted but that all leather parts are very worn;

C. **C**arole observes that the old engine consumed too much, that it has a stable market value but that its spare parts are expensive.

The participants view the object under discussion according to their own interests, which teleologically orient their perspectives or, more precisely, their own *levels of abstraction* (LoA). They may be talking about a car, or a motorcycle or even a plane, since any of these three referents would satisfy the descriptions provided by A, B and C above. Whatever the referent is, it provides the source of information and is called the *system*. Each LoA makes possible an analysis of the system, the result of which is called a *model* of the system (Figure 1). For example, one might say that Anne's LoA matches that of an owner, Ben's that of a

mechanic and Carole's that of an insurer. Evidently a system may be described at a range of LoAs and so can have a range of models.

We are now ready for a slightly more formal definition of a LoA. A LoA is a finite but non-empty set of *observables*, which are expected to be the building blocks in a theory characterised by their very choice. Since the systems investigated may be entirely abstract or fictional, the term "observable" should *not* be confused here with "empirically perceivable". An *observable* is just an *interpreted typed variable*, that is, a typed variable together with a statement of what feature of the system under consideration it stands for.

From the previous example and definition, it follows that a LoA is comparable to an *interface*. An interface is an intra-system, which transforms the outputs of system S into the inputs of system T and vice versa, producing a change in data types. LoAs are comparable to interfaces because:

1. they are a network of observables;
2. the observables are related by behaviours that moderate the LoA and can be expressed in terms of transition rules;
3. they are conceptually positioned between data and the agents' information spaces;
4. they are the place where (diverse) independent systems meet, act on or communicate with each other.

An interface (technically, a *gradient of abstractions*) consists of a collection of LoAs. An interface is used in analysing some system from varying points of view or at varying LoAs. In the example, Anne's LoA or interface might consist of *observables* for security, method of storage and owner history; Ben's might consist of observables for engine condition, external body condition and internal condition; and Carole's might consist of observables for running cost, market value and maintenance cost. The *gradient of abstraction* might consist, for the purposes of the discussion, of the set of all three LoAs.

The METHOD OF ABSTRACTION allows the analysis of systems by means of models developed at specific gradients for specific purposes. In the example, the LoAs happen to be disjoint but in general they need not be. A particularly important case is that in which one LoA includes another. Suppose, for example, that Delia (D) joins the discussion and analyses the system using a LoA that

includes those of Anne and Carole, plus some other observables. Let's say that Delia's LoA matches that of a buyer. Then Delia's LoA is said to be more concrete, or finely grained or lower, than Anne's and Carole's, which are said to be more abstract, or more coarsely grained or higher; for Anne's or Carole's LoA abstract some observables which are still "visible" at Delia's LoA. Basically, not only has Delia all the information about the system that Anne and Carole might have, she also has a certain amount of information that is unavailable to either of them.

It is important to stress that LoAs can be nested, disjoined or overlapping and need not be hierarchically related, or ordered in some scale of priority, or support some syntactic compositionality (the molecular is made by more atomic components).

We can now use the METHOD OF ABSTRACTION and the concept of LoA to make explicit the ontological commitment of a theory, in the following way. A theory comprises at least a LoA and a model. The LoA allows the theory to analyse the system under analysis and to elaborate a model that identifies some properties of the system at the given LoA (see Figure 1).
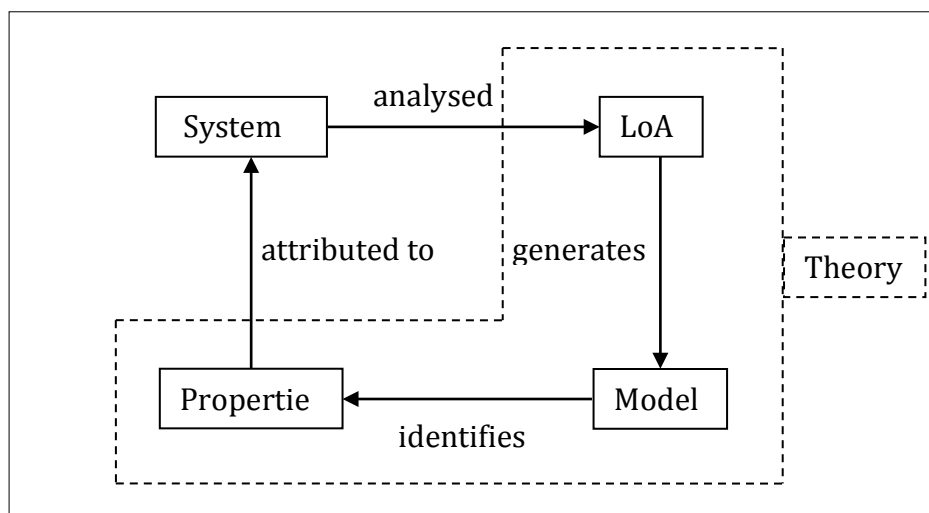
Figure 1 The scheme of a theory

The ontological commitment of a theory can be clearly understood by distinguishing between a *committing* and a *committed* component, within the scheme.

A theory commits itself ontologically by opting for a specific LoA. Compare this to the case in which one has chosen a specific kind of car (say a Volkswagen Polo) but has not bought one yet. On the other hand, a theory is ontologically committed in full by its model, which is therefore the bearer of the specific commitment. The analogy here is with the specific car one has actually bought (that red, four-wheeled, etc. specific object in the car park that one owns). To summarise, by adopting a LoA a theory commits itself to the existence of certain types of objects, the types constituting the LoA (by deciding to buy a Polo Volkswagen one shows one's commitment to the existence of that kind of car), while by adopting the ensuing models the theory commits itself to the corresponding tokens (by buying that particular vehicle, which is a physical token of the type Polo Volkswagen, one commits oneself to that token, e.g. one has to insure it). Figure 2 summarises this distinction.
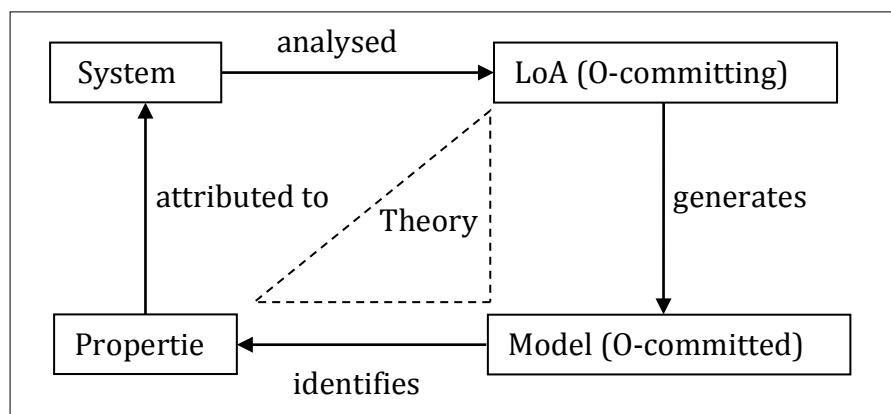
Figure 2 The SLMS scheme with ontological commitment

By making explicit the ontological commitment of a theory, it is clear that the METHOD OF ABSTRACTION plays an absolutely crucial role in any conceptual process, whether analytic or synthetic.

6. Constructionism

Providing the guidelines for choosing a problem and supplying a method for observing and analysing it are two of the fundamental steps in the information-theoretical, constructionist treatment of philosophical problems. In order to be complete and sound, the general methodology must also give an account of how

23

the problem must be investigated once it has been fully delineated. I shall refer to this method as CONSTRUCTIONISM.

CONSTRUCTIONISM is at the core of the maker's knowledge approach and of the kind of philosophy of information that I advocate in (Floridi 2011). As for MINIMALISM and the METHOD OF ABSTRACTION, CONSTRUCTIONISM too finds its roots in both the philosophical tradition and in computational theories and practices. We have seen that a black box is a system whose internal structure, rules and composition remain opaque to the observer. A white box is a system about which the observer knows everything, because the observer is actually its maker. This perspective, well known in Computer Science and Artificial Intelligence, lays in the wake of the so-called maker's knowledge tradition, according to which:

1. one can only know what one makes, and therefore
2. one cannot know the genuine nature of reality in itself.

Like Vico[7] and Hobbes, philosophers who hold (2) argue that, since any attempt to know the intrinsic nature of reality will inevitably fail, it is better to concentrate on those sciences whose subject is created by us, such as politics and social sciences. Philosophers who hold (1) argue that it is possible to improve our knowledge of reality through the development, application, and improvement (of our knowledge) of the techniques by which reality is investigated. We saw that this tradition finds its champion in Francis Bacon's philosophy of technology and it is related to Kantism. Following Bacon, technological mediations – which should be seen as including language as well, as Plato rightly perceived, see above the discussion about the dialectic method – become a crucial subject of philosophical enquiry because they are both a human product and the means through which the world is conceptualised, appropriated and investigated. Likewise, when Kant stresses the importance of understanding the conditions of possibility of our knowledge, he is working within the constructionist tradition. One can investigate scientifically the phenomena one experiences only insofar as one is epistemically responsible for them. Combining

---

[7] Vico famously argued that *verum ipsum factum* or *verum et factum convertuntur* (what is true and what is made are interchangeable). (Zagorin 1984) provides an insightful discussion of Vico's epistemology in the more general context of the maker's knowledge tradition. For a critical assessment see (Gaukroger 1986).

the Baconian and the Kantian lesson, Friedrich Dessauer argued convincingly that technology establishes a positive, poietic contact with noumena.[8]

The constructionist method consists of the following six principles:

1. The Principle of Knowledge: only what is constructible can be known. Anything that cannot be constructed at least conceptually could be subject, at most, to working hypotheses.

2. The Principle of Constructability: working hypotheses are investigated by conceptual models (including simulations) based on them ((Humphreys 2004)).

3. The Principle of Controllability: models must be controllable.

4. The Principle of Confirmation: any confirmation or refutation of the hypothesis concerns the model, not the modelled system.

5. The Principle of Economy: the fewer the conceptual resources used in the conceptual models the better. In any case, the resources used must be fewer than the results accomplished.

Constructionism suggests that, given a theory, one implements and tests it in a model. Because one constructs the model, one can also control it. As Newell and Simon remarked

> neither machines nor programs are black boxes; they are artefacts that have been designed, both hardware and software, and we can open them up and look inside (Newell and Simon 1976), p. 113.

Suppose that a robot one has built behaves like the ant one observed. The Principle of Confirmation prevents her from generalizing the working hypotheses, as if they were the real cause (or internal structure) of the modelled. It is obviously possible to provide an endless number of models with different internal structures whilst still obtaining the same behaviour. From this, the Principle of Context-dependency is derived:

6. The Principle of Context-dependency: isomorphism between the simulated and simulation is only local, not global.

In the previous example, the robot accounts for the behaviour of the ant only under the constraints specified by the model. If the constraints change, so does the evaluation of the hypotheses.

---

[8] (Dessauer 1958). I owe this insight to (Mitcham 1994).

CONSTRUCTIONISM is in plain contrast to any mimetic approach in epistemology and the philosophy of information, according to which reality is approached through some reproductive or representational mechanism. Ideas, mental images, corresponding pictures, concepts and so forth are supposed to be mere copies or portraits of some otherwise mysterious reality in itself. From the constructionist point of view, on the contrary, knowledge is a modelling process, which shapes and edits reality to make it intelligible. It therefore rejects more "mimetic" theories such as Plato's, Aristotle's or Descartes', in favour of a more Kantian approach. The Principle of Economy refers to the careful management of resources. On the one hand, in defining knowledge processes, mimetic theories are ontologically very wasteful, for they use a large amount of resources. Assuming that there is a reality and that it works in some particular way means making a heavy ontological commitment. On the other hand, CONSTRUCTIONISM does not state anything about reality in itself. This more modest and cautious commitment makes errors less likely. As in the case of MINIMALISM, the constructionist Principle of Economy differs from Ockham's razor too. While the latter is a post-production revision tool, for it provides a criterion for choosing among theories already produced, the former is a pre-design planning norm, for it requires the designer to be fully aware of the initial assumptions before she undertakes the investigation process and it binds the construction of any explanatory model to the conceptual resources available. Because in philosophy we always deal with intrinsically open questions, most of the conceptual costs supported in reaching plausible answers are hidden in the starting assumptions. The more powerful the latter are, the easier it will be to reach the wanted conclusions, yet the more costly the initial investment will be. In philosophy, as in life, free lunches are not available, and what we do not pay at the end, we pay upfront. The real difficulty, in a conceptual cost-benefit analysis, is balancing resources employed and results obtained. Cost-free conclusions are a chimera.

7. Conclusion

We have come to the end of this long journey. In the previous pages, I have argued that a constructionist approach in philosophy, in line with the maker's knowledge tradition, is both viable and fruitful. Constructionism is a very

common and unproblematic approach in all those sciences that have a poietic attitude towards their subject, which they tend not only to study but also to build, like computer science and economics. They could teach some important methodological lessons to philosophy, once the latter is seen as the highest expression of conceptual engineering. Plato was right in defining someone who knows (the dialectician) as the person "who knows how to ask and answer questions" (*Cratylus* (390c). He was wrong in identifying such knower as the information user rather than its maker. I am embarrassingly aware that much more work lies ahead, not only in order to develop and refine the constructionist approach in more detail, but also, and perhaps more importantly, in order to apply it to philosophical problems, both old and new.[9] This should not be discouraging. The world in which we live seems in great need of all the possible help we can give it, and a constructionist philosophy capable of devising the required concepts that will enhance our understanding may definitely lend a hand, if we can manage to develop it.

---

[9] An important example is provided by the question what kind of understanding a constructionist philosophy of information both supports and requires, see (Kuorikoski 2009).

References

Austin, J. L. 1962. *How to do things with words*, *The William James lectures*. Oxford: Clarendon Press.

Child, A. 1953. Making and knowing in Hobbes, Vico, and Dewey. *University of California Publications in Philosophy* 16 (13):271-310.

Deleuze, Gilles, and Fe lix Guattari. 1994. *What is philosophy?* London: Verso.

Dessauer, Friedrich. 1958. *Naturwissenschaftliches Erkennen*. Frankfurt am Main: Josef Knecht.

Dretske, Fred. 1994. If You Can't Make One, You Don't Know How It Works. *Midwest Studies in Philosophy* 19 (1):468-482.

Floridi, L. 1994. 'Objective Knowledge': the Disappearance and Revaluation of 'Knowledges' from John Sergeant to Karl Popper. *Nouvelles de la République des Lettres* 1:97-122.

Floridi, Luciano. 2010. *Information - A Very Short Introduction*. Oxford: Oxford University Press.

Floridi, Luciano. 2011. *The Philosophy of Information*. Oxford: Oxford University Press.

Floridi, Luciano. 2008. The method of levels of abstraction. *Minds and Machines* 18 (3):303-329.

Floridi, Luciano. Forthcoming. Semantic Information and The Network Theory of Account. *Synthese*.

Floridi, Luciano, Mariarosaria Taddeo, and Matteo Turilli. 2009. Turing's Imitation Game: Still a Challenge for Any Machine and Some Judges. *Minds and Machines* 19 (1):145-150.

Gaukroger, Stephen. 1986. Vico and the Maker's Knowledge Principle. *History of Philosophy Quarterly* 3 (1):29-44.

Greco, Gian Maria, Gianluca Paronitti, Matteo Turilli, and Luciano Floridi. 2005. How to Do Philosophy Informationally. *Lecture Notes in Computer Science* Volume 3782/2005:623-634.

Grim, Patrick. 2003. Computational modeling as a philosophical methodology. In *The Blackwell Guide to the Philosophy of Computing and Information*, edited by L. Floridi. Oxford, New York: Blackwell.

Harel, Idit, and Seymour Papert. 1991. *Constructionism*. Norwood, N.J.: Ablex Pub. Corp.

Hintikka, Jaako. 1974. Practical versus Thoeretical Reason - An Ambiguous Legacy. In *Practical Reason*, edited by S. Körner. New Haven: Yale University Press.

Hobbes, Thomas. 1656. *Elements of Philosophy, the first section, concerning body. Written in Latine ... and now translated into English. To which are added Six lessons to the professors of mathematicks of the Institution of Sr Henry Savile, in the University of Oxford. [With plates.]*: 2 pt. Printed by R. & W. Leybourn, for Andrew Crooke: London.

Humphreys, Paul. 2004. *Extending ourselves : computational science, empiricism, and scientific method*. Oxford ; New York: Oxford University Press.

Jacovides, Michael. 2002. The Epistemology under Locke's Corpuscularianism. *Archiv für Geschichte der Philosophie* 84 (2):161-189.

Kuorikoski, Jaakko. 2009. Simulation and the Sense of Understanding. In *conference on "Models and Simulations 3"*. Charlottesville, Virginia.

Mitcham, Carl. 1994. *Thinking through technology : the path between engineering and philosophy*. Chicago ; London: University of Chicago Press.

Newell, A., and H. A. Simon. 1976. Computer science as empirical enquiry: Symbols and search. *Communications of the ACM* 19 (3):113-126.

Pérez-Ramos, Antonio. 1988. *Francis Bacon's idea of science and the maker's knowledge tradition*. Oxford: Clarendon.

Plato. 1989. *The collected dialogues of Plato: including the letters, edited by Edith Hamilton and Huntington Cairns, with introduction and prefatory notes*. 14th printing ed. Princeton: Princeton University Press.

Quine, W. V. O. 1951. Two dogmas of empiricism. *The Philosophical Review* 60 (1):20-43.

Raftopoulos, A. 2003. Cartesian analysis and synthesis. *Studies in History and Philosophy of Science* Part A, 34 (2):265-308.

Rorty, Richard. 2009. *Philosophy and the mirror of nature*. 30th anniversary ed. with a new introduction by Michael Williams and a new afterword by David Bromwich. ed. Princeton, N.J.: Princeton University Press.

Tully, James. 1980. *A discourse on property : John Locke and his adversaries*. Cambridge: Cambridge University Press.

Zagorin, Perez. 1984. Vico's Theory of Knowledge: A Critique. *The Philosophical Quarterly* 34 (134):15-30.

Zeigler, Bernard P. 1976. *Theory of modelling and simulation*. New York ; London: Wiley-Interscience.