

OXFORD



The Logic *of* Information

LUCIANO FLORIDI

The Logic of Information

*A Theory of Philosophy as
Conceptual Design*

Luciano Floridi

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Luciano Floridi 2019

The moral rights of the author have been asserted

First Edition published in 2019

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data
Data available

Library of Congress Control Number: 2018953774

ISBN 978-0-19-883363-5

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

Contents

| | |
|-------------------------------------|------|
| <i>Preface</i> | ix |
| <i>Acknowledgements</i> | xvii |
| <i>List of Most Common Acronyms</i> | xix |
| <i>List of Figures</i> | xxi |

Part I. Philosophy's Open Questions

| | |
|---|----|
| 1. What is a Philosophical Question? | 3 |
| Summary | 3 |
| 1. Introduction: Russell's 'such ultimate questions' | 3 |
| 2. The Variety of Questions | 4 |
| 3. A Resource-oriented Approach to the Nature of Questions | 6 |
| 4. Three Kinds of Question | 7 |
| 5. Philosophical Questions as Open Questions | 8 |
| 6. First Objection: There are No Open Questions | 10 |
| 7. Second Objection: There are Too Many Open Questions | 13 |
| 8. Third Objection: Open Questions are Unanswerable | 15 |
| 9. Fourth Objection: Open Questions are Indiscriminate | 20 |
| Conclusion: Philosophy as Conceptual Design | 23 |
| 2. Philosophy as Conceptual Design | 27 |
| Summary | 27 |
| 1. Introduction: From the User's Knowledge to the Maker's Knowledge | 28 |
| 2. Plato's Wrong Step | 31 |
| 3. The Maker's Knowledge Tradition | 35 |
| 4. A Constructionist Methodology | 38 |
| 5. Minimalism | 39 |
| 6. The Method of Levels of Abstraction | 41 |
| 7. Constructionism | 47 |
| Conclusion: Against Degenerate Epistemology | 49 |
| 3. Constructionism as Non-naturalism | 53 |
| Summary | 53 |
| 1. Introduction: A Plea for Non-naturalism | 54 |
| 2. The Nature of Naturalism | 57 |
| 3. Two Indefensible Non-naturalisms: The <i>Supernatural</i> and the <i>Preternatural</i> | 62 |
| 4. Two Defensible Non-naturalisms: The Normative and the Semantic | 63 |
| 5. In Defence of Non-naturalism | 65 |
| Conclusion: The Artefactual Nature of the Natural | 67 |

Part II. Philosophy as Conceptual Design

| | |
|---|-----|
| 4. Perception and Testimony as Data Providers | 71 |
| Summary | 71 |
| 1. Introduction: The Relationship between Knowledge and Information | 71 |
| 2. A First Potential Difficulty | 73 |
| 3. Some Background | 73 |
| 4. Perception and the Phaedrus' Test (Plato) | 75 |
| 5. Testimony and the Parrot's Test (Descartes) | 76 |
| 6. Data Providers | 77 |
| 7. A Second Potential Difficulty | 79 |
| 8. More Background | 80 |
| 9. The Vice Analogy | 81 |
| 10. The Constructionist Interpretation of Perception and Testimony | 82 |
| 11. Informational Realism: Structures, Interactions, and Causality | 92 |
| Conclusion: The Beautiful Glitch | 97 |
| 5. Information Quality | 101 |
| Summary | 101 |
| 1. Big Data | 101 |
| 2. The Epistemological Problem with Big Data | 104 |
| 3. From Big Data to Small Patterns | 104 |
| 4. Information Quality | 106 |
| 5. The Epistemological Problem with Information Quality | 107 |
| 6. A Bi-categorical Approach to Information Quality | 110 |
| Conclusion: Back to Fit-for-Purpose | 112 |
| 6. Informational Scepticism and the Logically Possible | 113 |
| Summary | 113 |
| 1. Introduction: History and Analysis of Scepticism | 114 |
| 2. The Two Faces of Scepticism | 117 |
| 3. Non-naturalism and the Foundational Problem in German-speaking Philosophy | 119 |
| 4. Coherentism, Naturalism, and the Refutation of Scepticism in British Philosophy | 122 |
| 5. Pragmatist Epistemologies in American Philosophy | 124 |
| 6. Possible Worlds and Borel Numbers | 126 |
| 7. The Edit Distance as a Modal Metrics | 131 |
| 8. Informational Scepticism or the Sceptical Challenge Reconstructed | 135 |
| 9. The Redundancy of Radical Informational Scepticism | 137 |
| 10. The Usefulness of Moderate Informational Scepticism | 141 |
| 11. Objections and Replies | 142 |
| Conclusion: From Descartes to Peirce | 147 |
| 7. A Defence of Information Closure | 149 |
| Summary | 149 |
| 1. Introduction: The Modal Logic of Being Informed | 149 |
| 2. The Formulation of the Principle of Information Closure | 151 |

| | |
|--|-----|
| 3. The Sceptical Objection | 155 |
| 4. A Defence of the Principle | 156 |
| 5. Objection and Reply | 159 |
| Conclusion: Information Closure and the Logic of Being Informed | 160 |
| 8. Logical Fallacies as Bayesian Informational Shortcuts | 162 |
| Summary | 162 |
| 1. Introduction: A Greener Approach to Logic | 162 |
| 2. What are Logical Fallacies? | 163 |
| 3. Do Formal Logical Fallacies Provide Any Information? | 163 |
| 4. Formal Logical Fallacies and Their Explanations | 164 |
| 5. Bayes' Theorem | 165 |
| 6. Bayes' Theorem and the Fallacy of Affirming the Consequent | 166 |
| 7. Bayes' Theorem and the Fallacy of Denying the Antecedent | 167 |
| 8. Logical Formal Fallacies and Their Bayesian Interpretation | 168 |
| 9. Advantages of the Bayesian Interpretation of Formal Logical Fallacies | 169 |
| Conclusion: Rationality Regained | 170 |
| 9. Maker's Knowledge, between A Priori and A Posteriori | 171 |
| Summary | 171 |
| 1. Introduction: The Question about Maker's Knowledge | 172 |
| 2. Maker's Knowledge: Same Information | 174 |
| 3. Maker's Knowledge: Different Account | 177 |
| 4. Maker's Knowledge: <i>ab anteriori</i> | 180 |
| Conclusion: Some Consequences of the Analysis of the Maker's Knowledge | 185 |
| 10. The Logic of Design as a Conceptual Logic of Information | 188 |
| Summary | 188 |
| 1. Introduction: Two Modern Conceptual Logics of Information | 189 |
| 2. Design, Contradictions, and Dialetheism | 195 |
| 3. The Logic of Design as a Logic of Requirements | 197 |
| Conclusion: From <i>Mimesis</i> to <i>Poiesis</i> | 204 |
| Afterword—Rebooting Philosophy | 207 |
| Introduction | 207 |
| Scholasticism as the Philosophical Enemy of Open Questions | 207 |
| Philosophical Questions Worth Asking | 209 |
| A Philosophical Anthropology to Approach Philosophical Questions | 210 |
| How to Make Sense of the World and Design It Today | 211 |
| Conclusion: Creative Destruction in Philosophy | 212 |
| <i>References</i> | 215 |
| <i>Index</i> | 231 |

Preface

The title of this book contains two words that are semantically very loaded, “logic” and “information”. Before I say anything about its contents, I better explain how I use each of them.

Consider “logic” first. Philosophy is often concerned with the study of structural properties of systems—whether empirical, like political organizations, or conceptual, like scientific theories—and their dynamics.¹ Before the advent of modern mathematical logic, such structural studies might have been considered to be part of logic broadly understood.² Today, this would be very misleading. But so would be the use of labels such as “informal logic”³ and “philosophical logic”,⁴ already appropriated by other branches of philosophy. Thus, for lack of a better expression, I have chosen to refer to such structural studies as “conceptual logic”. A conceptual logic is what we have in mind when we ask questions about the abstract properties, principles, mechanisms, or dynamics characterizing or underpinning systems like international relations, a bureaucracy, a scientific theory, or the internal arrangement of elements in a device to perform a specific task. A conceptual logic focuses on formal features that are independent of specific implementations or idiosyncratic contingencies, on types rather than tokens, on invariants and their mutual relations, and on state transitions that are generalizable. This is the sense in which “logic” should be understood in the title of this book.

Next, take “information”. What remains unclear in the previous explanation is exactly what further qualifies as a conceptual logic of a system. Later in this book, I shall use a more precise vocabulary and talk about the “model” (semantic information about) of a “system” (the referent “abouted” by some semantic information). This means that the overall question I address across this book, and hence its subtitle, can be phrased more precisely by asking: what is the conceptual logic of information-modelling (i.e. generating a description of some structural properties of) a system? Note that, once unpacked, this question is not *ontological* but *epistemological*. It is not about the intrinsic, conceptual logic of a system in itself, but rather about the conceptual logic of the model (i.e. information) of a system, and how we design⁵ it. If this is

¹ See Alexander (1964, p. 8).

² For an enlightening reconstruction of the late modern developments of logic between Kant and Frege see Heis (2012).

³ See two excellent examples: Sinnott-Armstrong and Fogelin (2015), Walton and Brinton (2016).

⁴ See two excellent examples: Burgess (2009), Sainsbury (2001).

⁵ Parsons (2016) is a great source for a philosophical analysis of the debate on the concept of design. The book is also an essential starting point to frame the proposal made in this book. It dedicates quite some space to the debate on modernism. It is a must-read for anyone interested in understanding the conceptual nature of design together with the excellent Dunne and Raby (2013). These works, together with Barthes (1981) and Flusser (1999), have deeply influenced my views on the topic.

unclear, compare the conceptual logic of a watch's mechanism with the conceptual logic of the design of a watch. The focus is really about a conceptual logic of the design of a system, understood as a specific kind of conceptual logic of information about (the design of the model of) a system; and not any kind of information, but that particular kind we call *semantic* or *factual* information, that is, information about something. Thus, "information" in the title of this book is not the sort of quantitative, probabilistic, signal-based concept that one may find in an engineering textbook (what is sometimes called Shannon information)—although I shall use that too in the following chapters—but the well-formed, meaningful, and truthful kind that one may find in a tourist guide about Rome.

Focusing on the logic of a model rather than on the logic of a system captured by a model means that this book is closely linked with some of the epistemological ideas I supported in a previous book, entitled *The Philosophy of Information* (Floridi 2011c). However, this is not a book in the epistemological tradition. It is rather a book on the logic of design and hence of making, transforming, refining, and improving the objects of our knowledge. So, it is not a book about ontology either, at least not in the classic sense. For some philosophers, it may be obvious that the conceptual logic of semantic information is the deep, conceptual logic of the world. Or, to put it in the vocabulary just introduced, that the logic of a model of a system is the logic of the system. In this, Hegel, Marx, the first Wittgenstein, and the atomistic Russell appear to share the same view. I am not so sure. Actually, I really doubt it, but, most importantly, this is not an issue that needs to be resolved here. We can all refrain from stepping into any ontological discussion, and phrase anything that we need to discuss epistemologically, and therefore more safely in terms of commitment. If, as a bonus, some readers may believe that we pay one and get two, as it were—that is, if anyone is convinced that by developing a conceptual logic of the model we shall also have obtained a conceptual logic of the corresponding system—I shall not complain, but neither shall I join them.

Putting the two clarifications together, this book is a study in the *conceptual logic of semantic information*. The next thing I should be candid about is that it is a *constructionist* study.

The scientific-realist reader should not be alarmed. If anything, this is a pragmatist book about a world that needs to be modelled correctly in order to be made meaningful satisfactorily. I never seriously question for a moment the recalcitrant and often grating presence of outside realities. Actually, I have little patience for self-indulgent and irresponsible speculations about logically possible situations involving brains in a vat, dreaming butterflies, or Matrix-like pills. In the best scenarios, they were once devised to help discuss valuable philosophical questions. They have now become scholastic puzzles that engage only professional academics.

At the same time, the naïve realist may rightly be a bit suspicious. Because the constructionism I develop in this book is still of a Kantian kind (more on Kant below). The starting point defended in this book, as in my previous works, is that we do not have access to the world in itself, something I am not even sure one can explain in full.

We always access any reality through specific interfaces, what I shall define in Chapter 2 as *levels of abstraction*, borrowing a conceptual framework from computer science. The world provides the data, to be understood as constraining affordances, and we transform them into information, like semantic engines. But such transformation or repurposing (see Chapter 4) is not equivalent to portraying, or picturing, or photographing, or photocopying anything. It is more like cooking: the dish does not represent the ingredients, it uses them to make something else out of them, yet the reality of the dish and its properties hugely depend on the reality and the properties of the ingredients. Models are not representations understood as pictures, but interpretations understood as data elaborations, of systems. Thus, the whole book may also be read as an articulation and defence of the thesis that *knowledge is design*, and that philosophy is the ultimate form of *conceptual design*.

There is one more clarification I should make. Shifting from representing to interpreting the world may be controversial. We shall see that it raises problems in terms of quality of the outcome (Chapter 5) and scepticism about its truthfulness (Chapters 6 and 7). But it is not new and it can easily be mistaken for something that I criticize throughout the book: a passive understanding of knowledge as being always about something already present in the world. This Platonic view, to be found also in Descartes, is too reductive. Modelling is not just dealing with what there is; it is often designing what there could, or should be. Having a different word helps here: a *blueprint* is a model of something that does not yet exist but that we want to design, not of something that is already there and that we want to explain, for example. So, this book is a *constructionist* study in the *conceptual logic of semantic information* both as a *model* (*mimesis*) and as a *blueprint* (*poiesis*). We have reached the full description. And this can now be used to contextualize this book within the wider project for the foundation of the philosophy of information.

This is the third volume in a tetralogy that includes *The Philosophy of Information* (volume one (Floridi 2011c)) and *The Ethics of Information* (volume two (Floridi 2013b)). I began labelling the tetralogy *Principia Philosophiae Informationis* not as a sign of bottomless hubris—although it may well be—but as an internal pun among some colleagues. In a sort of rowing-like competition, I joked that it was time for Oxford to catch up with Cambridge on the 3–0 scoring about “*principia*”. Not a pun that many find funny, or even intelligible.

Within the *Principia* project, this book occupies a middle ground between the first and the second volume. However, as the reader should expect, the three volumes are all written as stand-alone, so this book too can be read without any knowledge of anything else I have ever published. Yet the three volumes are complementary. The essential message from volume one is quite straightforward. Semantic information is well-formed, meaningful, and truthful data; knowledge is relevant semantic information properly accounted for; humans are the only known semantic engines and conscious informational organisms who can design and understand semantic artefacts, and thus develop a growing knowledge of reality; and reality is the totality of

information (notice the crucial absence of ‘semantic’). Against that background, volume two investigates the foundations of the ethics of informational organisms (inforgs) like us, which flourish in informational environments (infosphere), and are responsible for their construction and well-being. In short, volume two is about the ethics of inforgs in the infosphere. Thus, in a classic Kantian move, we are shifting from theoretical to pragmatic philosophy. I already mentioned that this third volume, insofar as it focuses on the conceptual logic of semantic information as a *model*, is linked with the epistemological analysis provided in *The Philosophy of Information*. And insofar as it focuses on the conceptual logic of semantic information as a *blueprint*, it offers a bridge towards the normative analysis provided in *The Ethics of Information*. This discusses, among other things, what duties, rights, and responsibilities are associated with the poietic practices that characterize our existence, from making sense of the world to changing it according to what we consider normatively right or morally good. By working like a hinge between the two previous books, this third one prepares the basis for volume four, on *The Politics of Information*. There, the epistemological, conceptual, and normative constructionism supports the study of the design opportunities we have in understanding and shaping what I like to call “the human project” in our information societies.

The three volumes may be understood as seeking to invert four misconceptions, easily explainable by using the classic communication model introduced by Shannon: sender, message, receiver, channel. Epistemology focuses too much on the passive receiver/consumer of knowledge, when it should be concerned with the active sender/producer, because knowing is constructing. Ethics focuses too much on the sender/agent, when it should be concerned with the receiver/patient, because the keys to being good are care, respect, and tolerance. Metaphysics focuses too much on the relata, the sender/producer/agent/receiver/consumer/patient, which it conceives as entities, when it should be concerned with the message/relations, because dynamic structures constitute the structured. And logic focuses too much on channels of communication as justifying or grounding our conclusions, when it should be concerned with channels that enable us to extract (and transfer) information from a variety of sources reliably, because the logic of information design is a logic of relata and clusters of relations, rather than a logic of things as bearers of predicates. I will be utterly amazed if even one of these u-turns in our philosophical paradigms will be successful.

Let me now turn to a quick overview of the contents. The task of this third volume is still that of contributing to the development of a philosophy of our time for our time, as I have written more than once. As in the previous two volumes, it does so systematically (conceptual architecture is pursued as a valuable feature of philosophical thinking) rather than exhaustively, by pursuing two goals.

The first goal is meta-theoretical and is fulfilled by Part I. This comprises the first three chapters. There, I offer an interpretation of philosophical questions as open questions (Chapter 1), of philosophy as the conceptual design of such questions and of their answers (Chapter 2), and of constructionism as the best way of approaching such a way of doing philosophy (Chapter 3).

On the basis of this first part, the second part pursues not a meta-theoretical but a theoretical investigation. In Chapter 4, I investigate how a constructionist philosophy as conceptual design can interpret perception and testimony as data providers. This leads to the discussion of two further issues: the quality of the information that we are able to elaborate on the basis of such data (Chapter 5); and the truthfulness of the information constructed, namely the sceptical challenge faced by a constructionist interpretation of semantic information, and its related modelling processes (Chapter 6). In Chapters 7 and 8 I return to the features of a constructionist logic to analyse the value of information closure for the consistency of our models of the world, and the possibility of expanding such models even through fallacious reasoning. In Chapter 9, I finally define what the maker's knowledge could be, once we interpret ourselves as the designers responsible for our own understanding of the world. The book ends with a detailed analysis of what conceptual logic may underpin our designing activities, seen as depending on constraints and a relation of "a makes b a sufficient solution" (Chapter 10). The ten chapters are strictly related, so I have added internal references whenever it might be useful. They could be read in a slightly different order, as one of the anonymous referees pointed out. I agree. I offer here the structure that I find more helpful, but, for example, some readers may wish to go to the end of the story first, and learn about the logic of design as a conceptual logic of information before approaching any other chapter.

In terms of its philosophical roots, like volumes one and two, this too is a German book, written from a post-analytic continental-divide perspective, which I have the impression is increasingly fading away. The careful reader will easily place this work in the tradition linking pragmatism, especially Charles Sanders Peirce, with the philosophy of technology, especially Herbert Simon.⁶ Unlike volume one, and even more than volume two, this third volume is much less neo-Kantian than I expected it to be. And contrary to volume two, it is also less Platonic and Cartesian. In short, writing it has made me aware that I may be moving out of the shadows of my three philosophical heroes. Not a plan, but this is what happens when you follow your reasoning wherever it leads you. *Amici Plato, Cartesius et Kant, sed magis amica veritas*. In *The Ethics of Information*, I wrote that "some books write their authors". I now have the impression that only bad books are fully controlled by their authors. They are called airport novels and telenovelas. Let me illustrate this with an anecdote.

One day, when I was a boy, my mother found me distressed because I had just read about the death of Porthos, one of the four musketeers, in Dumas' *The Man in the Iron Mask* (1847–50, now (Dumas 2008)). There was not much she could do, but she tried to console me by pointing out that I was not alone. Dumas himself had been deeply upset by the death of his character. She recounted the story according to which a maiden had found Dumas crying because Porthos was dead. This seemed to me perfectly

⁶ The reader interested in exploring these connections may wish to consult Allo (2010) and Durante (2017).

reasonable. It was a very sad moment indeed. But what struck me at the time, and made me forget my sadness, was the alleged explanation that Dumas offered when asked why he, the omnipotent author, could not change the story of his own literary character. He said that Porthos *had* to die. That seemed to me to express, better than anything else I had experienced at the time, the force of a theorem, the strength of logical coherence, the inescapable, constraining grip of dynamic structures, which have features and interact in ways that are utterly independent of our wills and desires. You can choose the *if*, but not the *then*. Dumas was a great writer because he felt powerless. Porthos' death was the inexorable conclusion, given the development of the plot. That evening I thanked my mother for the lesson in logical thinking, and felt a bit less upset: necessity is somewhat soothing. Many years later, I went to check the episode again. It still saddens me. But I found a line that seems very appropriate to conclude this anecdote. "Parbleu!" said Porthos again, with laughter that he did not even attempt to restrain, 'when a thing is explained to me I understand it; begone, and give me the light.'

Regarding the style and structure of this book, as I wrote in the preface of volumes one and two, I remain painfully aware that this third volume too is not a page-turner, to put it mildly, despite my attempts to make it as interesting and reader-friendly as possible. I remain convinced that *esoteric* (in the technical sense) research in philosophy is the only way of developing new ideas. But *exoteric* philosophy has its crucial place. It is like the more accessible and relevant tip of the more obscure and yet necessary part of the iceberg under the surface of everyday life. The reader interested in a much lighter reading may wish to check *The Fourth Revolution: How the Infosphere is Reshaping Human Reality* (Floridi 2014a) or perhaps the even easier (Floridi 2010b).

This book will require not only graduate-level knowledge of philosophy, patience, and time but also an open mind. The last three are scarce resources. For in the last two decades or so of debates, I have been made fully aware—sometimes in ways less friendly than I wish—that some of the ideas I defend in the following pages are controversial. I have also seen how often mistakes are made by relying on “systemic attractors”: if a new idea looks a bit like an old idea we already have, then the old one is a magnet to which the new one is powerfully attracted, almost irresistibly. We end up thinking that “that new” is really just like “this old”, and if we do not like “this old” then we dislike “that new” as well. Bad philosophy indeed, but it takes mental strength and exercise to resist such a powerful shift. In the case of this book, I am sure some readers will be tempted to conclude that it is an anti-realist, anti-naturalist attempt to let the sceptic and the relativist have the final word about logic, knowledge, and science. So, let me say this as clearly as I can, quoting Margaret Thatcher: no, no, no. There are many ways of being anti-sceptic and anti-relativist, and being a realist constructionist is one of them. I actually believe that it is the best option we have today. The reader does not have to follow me that far. But no mistake should be made about the direction I am taking.

As in the previous two volumes, two features that I thought might help the reader to access the contents of this book more easily are the summaries and conclusions at the

beginning and at the end of each chapter, and some redundancy. Regarding the first feature, I know it is slightly unorthodox, but the solution, already adopted in volumes one and two, of starting each chapter with a ‘Previously in Chapter x . . .’, should enable the reader to browse the text, or fast-forward entire chapters of it, without losing the essential plot. Science-fiction fans, who recognize the reference to *Battlestar Galactica*, may consider this third volume as the equivalent of season three.

Regarding the second feature, while editing the final version of the book, I decided to leave in the chapters some repetitions and some rephrasing of recurrent themes, whenever I thought that the place where the original content had been introduced was too distant, either in terms of pages or in terms of theoretical context. If sometimes the reader experiences some *déjà vu* I hope it will be to the advantage of clarity, as a feature and not a bug.

A final word now on what the reader will not find in the following pages. This is not an introduction to formal logic from an informational perspective (a book that should be written) or to some modal logic of “ S is informed that p ”. Nor do I seek to provide an exhaustive investigation of all the issues that may go under the label of “logic of information”. In particular, I hope to work more extensively on testimony, epistemic luck, and causality in the future. They are a topic just touched upon in the following chapters. The book is not a textbook on information theory or on the sort of topics discussed in undergraduate courses entitled “Introduction to (library and) information science” either.⁷ The reader interested in such topics might wish to look at *The Routledge Handbook of Philosophy of Information* (Floridi 2016) or at *Information—A Very Short Introduction* (Floridi 2010b). This book is about the roots of some of the logical problems of our time, not their leaves. It is about design.

⁷ I discussed the relationship between the philosophy of information and library and information science in Floridi (2002a, 2004b).

Acknowledgements

After the experience of the first two volumes, I am bound to acknowledge here some of the things I already acknowledged there. The repetition does not diminish but reinforces my deep gratitude.

As before, I could not have worked on such a long-term project without dividing it into some feasible and much smaller tasks. Some of the earliest bits of this book date to 1999, although most of the work was done between 2011 and 2017. I am glad to see that so many different components still fitted together in the end. Of course, I hope this is the symptom of a well-planned project, but I am afraid that, the older I get, the more likely it is that it may actually be due to some mental resistance to change, also known as sclerosis. Old people do not change their minds easily.

As for the previous two volumes, many of the ideas and contents presented in the chapters were initially tested as conference presentations or (sometimes inclusive or) journal articles. The bibliographic details are in the list of references. This systematic and gradual way of working is laborious, but it seems to be fruitful and it may also be inevitable, given the innovative nature of the field. It does require a perseverance and commitment that I hope are not ill-exercised. I wished to assess the ideas presented in this volume as thoroughly as possible, and publishing the articles gave me the opportunity and privilege to enjoy a vast amount of feedback, from a very large number of excellent colleagues and anonymous referees. If I do not thank all of them here, this is not for lack of manners or mere reason of space, but because the appropriate acknowledgements can be found in the corresponding, published articles.

There are, however, some people I would like to mention explicitly because they played a significant role throughout the project and during the revisions of the final text. Kia first of all. She and I have been married for as long as I have been working on the whole *Principia* project. Without her, I would have never had the confidence to undertake such a task, and the spiritual energy to complete it. It requires a huge amount of serenity to invest so much time in philosophical thinking, and Kia is my muse. I am repeating something I already wrote in the previous volume: she keeps making our life blissful, and I am very grateful to her for the endless hours we spend talking about the topics of the *Principia*, for all her sharp suggestions, and for her lovely patience with an utterly obsessed, single-minded, overly focused husband, who must be unbearable to anyone else. I wish I could claim to have contributed to her research in neuroscience even half as much as she has to my philosophical development. Between the second and this third volume we moved into the house of our dreams. Mill Barn is the perfect place to think.

Patrick Allo, Selmer Bringsjord, Min Chen, Marcello D'Agostino, Massimo Durante, Giuseppe Primiero, Sebastian Sequoiah-Grayson, Ugo Pagallo, Mariarosaria Taddeo,

David Watson, and James Williams, as well as several members and visitors of the Digital Ethics Lab at Oxford during the past years, were very generous with their time and provided numerous opportunities for discussion and further reflection on the topics analysed in this book.

A very long list of meetings, conferences, and workshops (the details are in the published articles listed in the References section) provided stimulating and fruitful venues to elaborate and further test some of the ideas presented in this book. I have learnt a lot from the colleagues I met there, though I wish I had been able to learn more.

Peter Momtchiloff was pivotal for the realization of all three volumes, both because of his foresighted invitation to publish them with OUP, and because of his support and firm patience, when it seemed that I would never complete them.

The anonymous reviewers appointed by OUP were incredibly helpful with their comments and suggestions. I was impressed especially by one of them for her or his careful reading and very insightful recommendations. I wish I could thank her or him in person.

Jessica Antonio and Louise Parker, my personal assistants at different times, provided the exceptional support and impeccable managerial skills without which I could not have completed this project.

I would also like to thank all the people who made the task of working on this book more difficult than it should have been. Sometimes friction is essential to make progress.

Finally, this book was made possible by the very generous and friendly support of the Universidad Carlos III de Madrid, where I spent a fruitful and quiet period of research as the 2015 Chair of Excellence, thanks also to the support of the Santander programme. The colleagues at Carlos III, and especially Ilia Galan Diez, provided the best intellectual and social environment to develop some of the ideas presented in the following pages. I am immensely grateful for the privilege. The book is dedicated to them.

List of Most Common Acronyms

| | |
|------|--|
| A | agent, Alice |
| B | agent, Bob |
| C | agent, Carol |
| ICTs | Information and Communication Technologies |
| LoA | Level of Abstraction |
| PI | Philosophy of Information |

The use of Alice as synonymous with Agent is not merely an Oxford-related reference. The reader acquainted with the literature on quantum information will recognize in Alice, Bob, Carol, and Dave the more memorable and concrete placeholder names for otherwise abstract agents or particles; see http://en.wikipedia.org/wiki/Alice_and_Bob, retrieved January 2018.

List of Figures

| | |
|---|-----|
| 1. The scheme of a theory | 46 |
| 2. The SLMS (System-LoA-Model-Structure) scheme with ontological commitment | 47 |
| 3. Ngram of ‘epistemology naturalized’; case-insensitive; year-range = 1800–2008; Corpus = Google Books in English; smoothing = 3 | 58 |
| 4. Ngram of ‘naturalization’; case-insensitive; year-range = 1800–2008; Corpus = Google Books in English; smoothing = 3 | 58 |
| 5. Fairy Trails from Freeverse (2009) | 64 |
| 6. Shannon’s model of communication not applicable | 83 |
| 7. The agent-system interactions elicit data as signals at a bodily LoA | 84 |
| 8. Data source and referent | 85 |
| 9. The data processing theorem | 87 |
| 10. Non-natural meanings as data hacking | 89 |
| 11. The construction of perceptual information | 91 |
| 12. Perception, testimony, and language | 93 |
| 13. The growth of big data | 102 |
| 14. Global information (data) created vs. memory (storage) available | 105 |
| 15. Examples of IQ categories and dimensions | 108 |
| 16. Data become information within a context (Z), at a LoA (Y), chosen for a purpose (X) | 109 |
| 17. Example of bi-categorical IQ analysis | 111 |
| 18. Graph of a bi-categorical IQ analysis | 112 |
| 19. Example of a state of a system S consisting of two cards | 130 |
| 20. System and model share the same Borel number | 131 |
| 21. Example of a ’s misinformation about the state of system S consisting of two cards (see Figure 19) | 134 |
| 22. An example of Hamming distance between two Borel numbers | 134 |
| 23. Example of informational scepticism | 136 |
| 24. Some implications of informational scepticism | 137 |
| 25. An example of the redundancy of the sceptical challenge | 140 |
| 26. Valid inference schemes and formal logical fallacies | 165 |

| | |
|---|-----|
| 27. A simple application of Bayes' theorem | 166 |
| 28. Valid inferences and formal logical fallacies | 168 |
| 29. Alice has moved the King's pawn two steps (e2–e4) | 174 |
| 30. Mapping the four conceptual distinctions | 181 |
| 31. Commutative diagram for the maker's and the witness's knowledge | 184 |
| 32. Transcendental logic | 189 |
| 33. Dialectical logic | 192 |
| 34. The logic of design | 194 |
| 35. Elements of system design | 200 |

PART I

Philosophy's Open Questions

Advanced research develops at an ever-faster pace, so it is becoming increasingly hard to anticipate its future. This holds true in any field of investigation, including philosophy, where the scope and varied nature of the discipline make forecasting its development perhaps an even more challenging task. At the same time, there seems to be a growing demand for philosophical understanding, in a world in which conceptual problems, ethical issues, social and political tensions, economic imbalances, religious and cultural clashes, international conflicts, and environmental challenges have acquired global dimensions and visibly affect the lives of billions of people simultaneously. The world needs more philosophical insights and more ideas on how to cope with a complex future. It needs conceptual theorizing and philosophical answers. Therefore, it is simply irresponsible to dismiss philosophy as a useless waste of time, with cheap jokes, populist rhetoric, and unscientific remarks, as Stephen Hawking (more on him in Chapter 1), Neil deGrasse Tyson, or Bill Nye have done. True, philosophy has its ups and downs, as I shall argue in Chapter 1, but so have other disciplines, and it is essential that we should try to identify the most promising lines of philosophical research, and anticipate trends in which it might be worth investing our intellectual efforts and dwindling resources. What kind of philosophy are we developing? And what kind of philosophy should we be fostering? These are difficult questions. They also represent a classic *leit-motif* in the history of philosophy, because no serious philosophy worth pursuing shies away from the reflection on its own status. It comes with the job, as they say. They are also questions that should be asked preferably to (I did not write by) experts.¹ The main

¹ They have been; see for example Schlette (1968), Leaman (1998), Bencivenga (2006), Dietrich and Weber (2011), and Floridi (2013a). In the latter case, I sought to collect views from philosophers who are also editors in chief of influential philosophical journals and philosophy editors of important publishing houses. They may contribute particularly helpful perspectives on the future of philosophical research because they enjoy a privileged role with respect to its current state. Not only are they likely to have a better sense of what is the most promising work in progress, and what kind of fruitful patterns might be emerging, as the journals' long backlogs of forthcoming articles and publishers' long-term editorial plans clearly suggest. They are also the ones who may have an overall perspective about what is not published, and for what reasons. As in the case of negative results, this too is valuable information. According to a widely cited piece of research, (Fanelli 2012), every year there is a 6 per cent decrease in the publication number of scientific papers with negative results. It would be interesting to compare this situation to philosophical journals.

2 PHILOSOPHY'S OPEN QUESTIONS

questions to address are inevitably indexical: what now, what next, what for. The strategy is always, to quote Schlick:²

to ask the question, 'What will be the future of philosophy?' entirely from the point of view of the philosopher. However, to answer the question we shall have to use the method of the historian because we shall not be able to say what the future of philosophy will be except in so far as our conclusions are derived from our knowledge of its past and its present.

Inevitably, readers of this book will have their own answers. In this first part of the book, I present my own. I hope they are convincing, but even if they are not, they should be helpful in understanding how I approach the philosophical questions discussed in the second part.

For in the best scenarios, when new research topics begin to emerge, it is conceivable that attempts at studying them might initially be unsuccessful, and yet, the presence of such failures may still betray the importance of the topic itself. Finally, editorial strategies and choices also contribute to shaping the future of the discipline, so editors seem to be well positioned to discuss the future of philosophical research.

² Schlick (1932), republished in Schlick (1979).

1

What is a Philosophical Question?

Summary

There are many ways of understanding the nature of philosophical questions. One may consider their morphology, semantics, relevance, or scope. This chapter introduces a different approach, based on the kind of informational resources required to answer them. The result is a definition of philosophical questions as questions whose answers are in principle *open* to informed, rational, and honest disagreement, *ultimate* but *not absolute*, *closed* under further questioning, possibly *constrained* by empirical and logico-mathematical resources, but requiring *noetic* resources to be answered. The chapter concludes with a discussion of some of the consequences of this definition for a conception of philosophy as the study (or ‘science’) of open questions, which uses conceptual design to analyse and answer them. That is the topic of Chapter 2.

1. Introduction: Russell’s ‘such ultimate questions’

In 1912, Russell published *The Problems of Philosophy* (Russell 1912). A hundred years later, it is still a classic introduction to philosophy, accessible to beginners and insightful for experts. As the reader probably knows, it is a tiny book. In such a short space, Russell explicitly chooses to privilege epistemological problems over metaphysical ones. As he writes in the preface:

In the following pages, I have confined myself in the main to those problems of philosophy in regard to which I thought it possible to say something positive and constructive, since merely negative criticism seemed out of place. For this reason, theory of knowledge occupies a larger space than metaphysics in the present volume, and some topics much discussed by philosophers are treated very briefly, if at all. (Russell 1912, p. 1)

Such an epistemological orientation makes the book especially suitable for an informational interpretation, as the analysis of the problem of sense data shows, for example. Yet, in this chapter, I do not wish to engage with Russell’s selection of philosophical questions, with his own formulation of them, or even with the answers he advocates. Instead, I intend to concentrate on only one paragraph. It is the first of the book, and

is based on a simple but ground-breaking lesson we have learned from studies in computational complexity since the 1960s: a lesson that had its roots in the work of Alan Turing (Fortnow and Homer 2003), who was born the same year Russell published his book. It seems that 1912 was some kind of accidental threshold.

3. A Resource-oriented Approach to the Nature of Questions

To understand the approach to the nature of philosophical questions that I intend to advocate, we need to take a step back.

Turing gave us a clear analysis of what an algorithm is. This is crucial in order to shift our perspective on the nature of computational problems, because having a standard way of formulating algorithms means having a universal parameter to calculate the complexity of the problems that such algorithms are supposed to solve. Thus, a Turing Machine works a bit like the standard metre in Paris. The result, a few decades after Turing's ground-breaking research, is that, in computational theory, one investigates the nature of problems by looking at their complexity, understood as the degree of difficulty of solving them, by studying the resources required to solve them.⁷ This means that one does not focus on the specific *morphology* of the problems—because this is where having a universal model, such as a Turing machine, helps us not to be side-tracked—on their *semantic features*—because we are interested in whole classes of problems independently of their specific content and topic—or on their *scope* and *relevance*—because we are interested in their complexity, independently of their timeless applicability. Rather, one investigates the complexity of computational problems by studying the quantity and quality of *resources* it would take to solve them. This usually means considering the amount of *time* (number of steps) or *space* (quantity of memory) required. In this way, it becomes possible to study classes of problems that share the same degree of complexity, and to structure such classes in hierarchies of different degrees of complexity.

The importance of such a shift in perspective may be easily conveyed by means of a simple analogy. When studying animals, one may concentrate on their morphology (e.g. quadruped), or reproductive method (e.g. oviparous), or indeed many other features, such as their habitat (e.g. palustre, or swamp animal). Or one may consider what it takes it to keep an animal alive. According to this diet-oriented approach, a class of animals could be catalogued as frugivore, or fruit eater, for example. This is more or less what happens in computational theory: classes of computational problems are organized in terms of their complexity understood as the kind of resources they (would) consume in order to be solved.

⁷ An excellent introduction is provided by Sipser (2012). I have provided a simple and philosophically oriented outline of computational complexity in Floridi (1999).

The previous analogy should help to clarify the resource-oriented approach adopted below: we can study questions by considering the resources required by their answers. It should also dispel any concern about the strategy I shall follow. It is not my intention to try the impossible, that is, to classify and study philosophical problems on the basis of their computational complexity. On the contrary, later in the chapter it will become obvious that this should be considered an outdated, Leibnizian dream. In philosophy, there are no experiments and no calculations, for reasons that will soon be obvious. What I wish to borrow, from computational complexity theory, is only the simple, yet very powerful, insight that the nature of problems may be fruitfully studied by focusing on what it may take in principle to solve them, rather than on their form, meaning, reference, scope, and relevance.

Equipped with this new conceptual tool and a different approach, let us now return to our initial attempt to clarify what a philosophical question is.

4. Three Kinds of Question

Questions may be of many kinds;⁸ think of Socrates' rhetorical questions, for example, or his testing questions. And therefore they may serve a variety of purposes; for example, implicitly asserting that p , or explicitly checking whether Alice knows that p . However, it seems clear that, in this context (recall Russell's book), we are discussing questions as genuine requests for information⁹ or *queries*: we are looking for an answer we do not have but would like to obtain.¹⁰ It follows that, from a resource-oriented approach, and simplifying a bit,¹¹ we may distinguish between different kinds of questions as *queries* (henceforth I shall drop this qualification), depending on what kind of information it would take to answer them.

Obviously, some questions require empirical information to be answered, as when Alice asks Bob who is coming to their party on Saturday, and whether any guest is vegetarian. Equally obviously, some other questions require logico-mathematical information to be answered,¹² as when Alice asks Bob how many people might be coming

⁸ Analyses abound. In connection with the nature of philosophical questions, Cohen (1929) already provided a clear analysis of the other senses that I disregard in this chapter, and further bibliographical references.

⁹ "The question has usually been described as a request for information" (Cohen 1929); see also Llewelyn (1964) for an analysis of questions as ways to elicit information—what I have called queries.

¹⁰ The proposal in Cohen (1929)—critically discussed by Llewelyn (1964), but see also Carnap, C. I. Lewis, Reichenbach, Ryle, Hamblin, and Presley—to construe questions as propositional functions is perfectly compatible with the theses defended in this chapter, but its adoption would merely burden the text with an unnecessary complication.

¹¹ I am simplifying here because, of course, hybrid cases are very common, combining features of more than one class (hence my use of the inclusive "or" in the text), but this refinement only makes things more obscure without helping to make the point any more convincing, so I shall rely on the reader's intelligence.

¹² Whether logico-mathematical resources may be defined as logico-mathematical *information* is a difficult issue because the informational nature of logic and hence mathematics is problematic; see Hintikka's "scandal of deduction" (Hintikka 1973). I shall not discuss this issue here. For a full discussion and a constructive proposal see D'Agostino and Floridi (2009).

if each guest brings a spouse, and whether there is one chair per person. However, still other questions may be answered by neither empirical nor logico-mathematical information (or a combination of the two), as when Alice asks Bob whether they should have a party on Saturday at all, or why it may be preferable to anticipate it to Friday. In this case, no matter how much empirical and logico-mathematical information Bob may use to formulate his answer, he may still fail to address Alice's question convincingly. This is repeatedly stressed in many contexts by philosophers of almost any school. To quote Uygur again:

the question 'Is consciousness a mechanism or organism?' requires primarily a thorough investigation of the respective concepts 'mechanism' and 'organism.' The ground for treating philosophical questions is thus not various world-facts—e.g. machines or living beings—but different discourses embodying these concepts. (Uygur 1964, p. 67)

It gets worse. For Alice might marshal an equal amount of empirical or (inclusive or, I shall not repeat this clarification unless it is necessary) logico-mathematical information to formulate an answer different from Bob's, even if she and Bob may still be well informed, reasonable, sensible to each other's perspectives, and honestly willing to compromise. For this reason, I shall call this third kind of questions *open* to informed, rational, and honest disagreement. The reader will have guessed that philosophical questions are open questions, in the following sense.

5. Philosophical Questions as Open Questions

Empirical and logico-mathematical questions are in principle closed, whereas philosophical questions are in principle open. By this I do not mean to suggest that people necessarily agree on how to answer empirical and logico-mathematical questions (I wish). Nor am I suggesting that the relevant empirical or logico-mathematical information is always available as a resource to formulate a correct answer to any empirical and logico-mathematical questions (more wishful thinking). What I am suggesting is that empirical and logico-mathematical questions are such that, once we have the necessary and sufficient resources to formulate a correct answer, any further disagreement on that formulated answer may speak volumes about the parties involved but says nothing about the answer itself. This is how I interpret Kant (and I may be forcing this interpretation on him) when he writes:

there are sciences the very nature of which requires that every question arising within their domain should be completely answerable in terms of what is known, inasmuch as the answer must issue from the same sources from which the question proceeds. In these sciences it is not permissible to plead unavoidable ignorance; the solution can be demanded.

(Kant 2007, A 476/B 504)

If Alice and Bob disagree on who or how many people did come to their party, one of them is un- or misinformed, or irrational, or confused, or perhaps just stubborn, or a

*image
not
available*

answers) may also lead backwards, to previous questions: ‘Did you buy six bottles of wine?’ may lead to ‘Did you buy any wine?’ which in turn may lead to ‘Did you go shopping?’ When this happens, one may say that Q_1 presupposes Q_2 , which presupposes Q_3 , and so forth. More formally, erotetic implication—the logical relation interpreting the linguistic relations of ‘posing’ and ‘presupposing’ or of ‘arising from’, in the vocabulary of Wiśniewski (1994)—may be as formally stringent as that of ordinary implication. In the long run, it gives rise to a network, sufficiently loose to admit different topologies but not, for that reason, less robust. It follows that it is a mistake to think that ultimate questions are questions that come first or last in time, or at the beginning or the end of a chain of questions. Given the network of questions that mutually pose or presuppose each other, ultimate questions are those questions whose answers are most influential in terms of a cascade of further questions and answers to other related questions within that network. With an analogy, one may say that they are like key pieces in a jigsaw puzzle: once they are placed, it is much easier to locate the other pieces. For example, how one answers the question about the existence of God, or what is a morally good life worth living, has immensely more influence—in terms of how one answers other questions within one’s own network—than how one answers the questions about the choice between black or brown shoes, or whether the party should be organized at all. Kant spends so much time discussing the question about causality because that is one of the cornerstones of our understanding of the world. Answer that, and a whole wave of consequences follows, leading to other questions. Ultimate questions are therefore attractors, in terms of systems theory: (answers to) less important questions within the system of questions tend to evolve towards (answers to) them over time. In terms of visual metaphors, it is better to imagine philosophical questions at the main nodes of our set of questions, big roundabouts in life’s roads.

In light of the second objection, we can further refine our definition by saying that philosophical questions are in principle *open* and *ultimate* questions, *closed under further questioning*. This is even better, but unfortunately the new definition still fails to deal with a third difficulty: scepticism about the available resources. What if philosophical questions, as defined above, cannot be answered because of an inescapable lack of the right kind of information? This is the third objection.

8. Third Objection: Open Questions are Unanswerable

Philosophical questions may be philosophical doubts (see the quotation from Russell in section 1). When they are, they still preserve the features of being open to reasonable disagreement and closed under further doubting. So, one can follow Descartes, rely on closure, and deploy doubt to reach certainty about the doubting process itself: radical doubts generate more doubts, but at least this much is certain. The problem, even in a Cartesian-friendly scenario, is not so much that certainty is unreachable in this way but that a purely formal, that is, content-empty, status is all that one can reach safely.

image

not

available