



THE THINGS BEFORE US

On What it Is to Be an Object

E.J. Booij



**The Things Before Us:
On What it Is to Be an Object**

Elbert Jan Booij

The Things Before Us:
On What it Is to Be an Object

ILLC Dissertation Series DS-2021-02



INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

For further information about ILLC-publications, please contact

Institute for Logic, Language and Computation
Universiteit van Amsterdam
Science Park 107
1098 XG Amsterdam
phone: +31-20-525 6051
e-mail: illc@uva.nl
homepage: <http://www.illc.uva.nl/>

Copyright © 2020 by Elbert Jan Booij

Cover design by Ezra Booij, Hella Snijder, and Elbert Booij.
Printed and bound by Proefschriftenprinten.nl, Ede.

ISBN: 978-90-831117-7-3

The Things Before Us:
On What it Is to Be an Object

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. K.I.J. Maex
ten overstaan van een door het College voor Promoties ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op woensdag 27 januari 2021, te 13.00 uur

door Elbert Jan Booiij
geboren te Baarn

Promotiecommissie

Promotores:

prof. dr. ing. R.A.M. van Rooij	Universiteit van Amsterdam
prof. dr. F. Berto	Universiteit van Amsterdam

Overige leden:

prof. dr. A. Betti	Universiteit van Amsterdam
prof. dr. D. Cohnitz	Universiteit Utrecht
prof. dr. A.C. Bottani	University of Bergamo
prof. dr. M.J.B. Stokhof	Universiteit van Amsterdam
dr. P.J.E. Dekker	Universiteit van Amsterdam

Faculteit der Geesteswetenschappen

to Mirjam

to Ezra and Abel

to my parents

Contents

Acknowledgments	ix
Introduction	1
1 The Draft Ontology	7
1.1 Things	7
1.2 On Where to Start	10
1.3 Layers and Roots	13
1.4 Principles	16
1.5 The Draft Ontology	20
1.6 On Where to Go	23
2 Objects	27
2.1 The Logic of Objects	27
2.2 Inferential Reasoning	32
2.3 Perception, Conception, Conversation	35
2.4 Being	38
2.5 Plurality	46
3 Properties, Procedures, and Images	53
3.1 A Conflict Within the Draft Ontology	53
3.2 A Provisional Diagnosis	58
3.3 Procedures	61
3.4 Images	64
3.5 Grades of Sparseness	66
4 Mereological Fusion as a Case Study	73
4.1 Parts and Wholes	73
4.2 Pleading Innocence	75
4.3 The Challenge of Nihilism	76

4.4	Nihilese and Mereolese	78
4.5	Complex Procedures	81
4.6	Vicinitive Relations	84
4.7	Emergence	88
5	Abstraction	93
5.1	Abstract Objects	93
5.2	Triviality	96
5.3	Kinds	97
5.4	The Primacy of Generality	99
5.5	Abstraction vs Fusion	101
5.6	Attributes	105
5.7	In Search of Individuality	107
5.8	Deeper Abstraction	111
5.9	Perfection	113
6	Being and Non-Being	115
6.1	Identity	115
6.2	Roots	121
6.3	Existence	125
6.4	Imagination	130
6.5	Modality	133
7	Discussion	137
7.1	On Balance	137
7.2	An Explanatory Limitation	138
7.3	Objectivity	140
7.4	Conceptual Content	142
7.5	Indeterminacy	146
7.6	Causality	151
	Bibliography	154
	Samenvatting	167
	Summary	169

Acknowledgments

Ideas are always developed in interaction with others, and also the interest in philosophical matters itself is an attitude that strongly benefits from encouragement by those around you. I am grateful to my parents for always having provided this. My mother (by engaging in many of them) taught me how to have a discussion, and also that arguments need not always have a conclusion. My father taught me not to lose sight of the fact that ideas, however beautiful or inspiring, still deserve a critical eye with respect to truth.

My wife Mirjam was the one who convinced me to turn the old passion into a more serious project: to take courses at the University of Amsterdam and, bit by bit, to realize the aim of writing a PhD thesis. I am very thankful for her support and that of my sons Ezra and Abel. My family have always taken my philosophical efforts highly seriously and never stopped granting me the time and space to pursue them. At the ILLC there are many people I am indebted to, but Katrin Shulz's advice to do the MSc Logic proved particularly fruitful.

Special gratitude, of course, goes to my supervisor Robert van Rooij. All through the process that led up to this thesis, Robert turned out to be a stable factor, always ready to give advice, always enthusiastic, always critical. He was the one with whom I took my first real courses in logic, and after he supervised my master's thesis I was very glad to have him as my PhD supervisor. Many thanks for all those comments, suggestions, supportive remarks, and misgivings as well; I would not have got very far without them. Thanks also for spotting all the typos!

What proved to be one of Robert's best ideas was to ask Franz Berto to be my second supervisor. Viewing the subject matter from a slightly different angle—more metaphysical, more historical—Franz provided a complementary perspective that I greatly appreciated. His insistence on writing in ongoing dialogue with the literature was a welcome correction to my slightly solitary tendencies. This is also the place to mention his gentle—and apt!—criticism regarding my inclination to use what he referred to as 'colourful language'. I hope to have made

some progress. . .

Our tripartite meetings were always enjoyable and instructive. Both Robert and Franz to an admirable extent exhibit the attitude of ‘agreeing to disagree’: matters of opinion never interfered with the quality and depth of their considerations, remarks, or advices for further reading.

Finally I would like to thank Hella Snijder and Ezra for their contributions to the cover design.

Amsterdam
December, 2020.

Elbert Booij

Introduction

When it comes to making progress in solving the problems it has committed itself to, philosophy has a poor reputation, with no exception for metaphysics. This state of affairs does not generally drive the workers in these fields to despair. Often it is faced with a sense of acceptance, somewhat like the acceptance of one's mortality: uninvited but inevitable. More charitably one could hold that philosophers are only unusually patient, knowing that even the most revolutionary intellectual achievements in human history once started with writing in the sand of Mediterranean courtyards. Had not the noble art of medicine been around for millennia as a mere idea, little more than ineffectual magic, before it became a science? But however true that may be, there is something about the state of ontology in particular that is testing this mood to the point of being unbearable. Although science has provided vast amounts of insight with respect to what lies behind the properties of common objects, we do not seem to be able to shed light on what a property itself is, or what it means for an object to have one. And what is an object anyway?

In the attempts to gain some foothold here, most researchers in the field make extensive use of *logic*. This is a sensible choice: logic and mathematics together comprise the general language of structure, of regularities, rules, laws. In it, all the great results of physics are written and many of those achieved in other scientific disciplines. Given the startling propensity of nature to comply with descriptions in this language, it is to be expected that the same will be true with respect to the fundamental notions of ontology, especially since not a single remark about nature can be made *without* making use of these notions. Saying that there are electrons is, after all, using existential quantification. Saying that they are negatively charged is using the subject-predicate construction.

And indeed, up to a point this expectation is met: *first-order logic* is a formalism that fits our ontological intuitions remarkably well. Surely our world *is* a world of objects, characterizable by means of predicates; moreover none of the basic principles of first-order logic (such as individuals, predication, and exis-

tence) could be left out in any description of it. But first-order logic only seems to capture the outlines; in some sense it stops there. Clearly, it is possible to improve on the basic architecture in a multitude of ways, and few natural language expressions finally escape a fair enough paraphrase in one of the resulting logics. Yet, certain undeniable results notwithstanding, it seems fair to say that none of these extended logics gets anywhere near the combination of conceptual ingenuity and computational power of the original. Despite all the effort devoted on the matter, there appears to be much ontology that formal logic cannot do much more than account for in rather *ad hoc* ways.

Could disappointment about these limitations have made us inattentive to the unique merits of first-order logic? The sheer *indispensability* of its core concepts is something which has not, to my mind, received the interest it deserves. First-order logic is appreciated first and foremost for its qualities as a *formal system*: it is hailed for the ability to provide a safe embedding for mathematics, an invaluable basis for programming languages, and a Platonic universe in its own right, about which many deep facts can be proved. Not, however, for its *ontological* adequacy. And of course the match is not a perfect one: first-order logic is straightforwardly about individuals that have properties and entertain relations, but not about time or modality. Nor is the formalism naturally poised to represent part-whole relations or facts about abstract objects (the colour blue) and their exemplifications (a blue wall). Nevertheless, *what* first-order logic describes, it describes perfectly—all it can be blamed for is that there is much it fails to describe.

But maybe even this impression of insufficiency is largely due to flawed presumptions. As for the connection between objects and properties, a very common strategy is to understand properties as being *immanent* to the objects that have them. The property, so to speak, is where the object is: ‘out there’. This is the line of thought that typically prompts one to reject disjunctive, negative, and certain types of relational (‘mere Cambridge’) properties and to extend the options for property-building by considering properties of mereological wholes as some sort of composition out of their parts. In the course of time I have become increasingly convinced that all this works to some extent, but without in the least disclosing new ways to look at *other* property-issues, such as the one-over-many problem (there are many red things, but one colour red). And the extent to which it does work is hardly satisfactory.

That our intuitive stock of properties is *far* more extended than whatever can be defined by simple logical means I take to be an important clue for where to locate the failure of this general approach. Property descriptions of this origin do not even begin to cover the tiniest part of what is encountered in real life, and this is not because all ‘real’ properties are unanalysable. Being a bachelor often consists in being male and unmarried. For a wall to be blue is certainly for (some dominant subset of) its *parts* to be blue. To be living is more than to be reproductive or to breathe, but there is definitely a connection with these

conditions. However, rather than logical (and, say, mereological) recombinations only, it seems that we are intuitively willing to accept as a property *every* condition that can be determined by running a certain fixed recipe, an algorithm, a procedure on other conditions. Ontological misgivings notwithstanding, to be Swedish-or-Russian, or non-Irish, or north of Cambridge, or a basket with a prime number of eggs—all these are, as properties, fully legitimate. But if this is so, then it seems that such properties are to be understood as based *within our cognitive system*. They are not themselves ‘in the head’—the procedure that determines whether the eggs add up to a prime number is performed on the eggs, and they are elsewhere. The eggs do not need a human brain to exist, and there *is* a prime number of them. Yet, a property like this can be experienced by us only thanks to a certain computational ability of our mind. The property is therefore not immanent in any usual sense of the word. And pondering things over, we must conclude that there are a great many properties for which this is the case.

The question is how many. One might proceed by reasoning in the following way: look, *real* properties are immanent. There is an ultimate layer of ways things really are, and on top of that we humans consider a wealth of conditions that we compute out of them. We call them ‘properties’ but they are not; that is just a trick of the tongue. One *might* reason so, but then, which are those real properties supposed to be? The wall’s being blue—but what if our colour-vision were different? The electron’s being negatively charged? But negative charge is only a (very successful) theoretical posit explaining the way the electron moves. Its position then, or its velocity? But, according to quantum physics, exactly those are determinable to a limited degree only! If physical objects have real properties at all, then there is not a single one that we have direct access to. Whoever insists on this fundamental layer of immanent properties is in for a challenging quest, since none of the attributes that we ascribe to common things remotely qualifies. And that is what makes it hard to maintain that non-immanent properties are in some way invalid as such. They are—allowing for a few speculative exceptions—the only properties that we know of!

So the properties we know and live by are non-immanent. But then, what of the objects that have them? Our world is a world of sound and shape, distance and duration, prime numbers and patterns, and if we take these seriously as properties (and relations), then it would be quite outrageous to hold that their bearers are somewhere else completely. Following this line of thought to its conclusion would imply that the objects *themselves* are in part based in our cognitive system.

Let us pause here for a while, for it is an important junction we have reached. If we take this road, we skip the one with the signpost *realism*, and this decision should not be taken lightly. Realism is the view that lies behind the immanence of properties just mentioned. I shall use this term to refer to any view that considers the existence of at least some objects completely independent from any human observer. Traditionally the position diametrically opposite realism is *idealism*, by

which all objects are essentially ‘in the mind’. Plain, unsubtle idealism is a non-starter: were there no objects in the age of the dinosaurs—were there dinosaurs at all? Few would buy this view wholesale. But plain realism, without any qualification, is not very attractive either. The world is given to us through our senses, which have their own peculiar, species relative way of presenting it.

There is, or so I shall argue, a middle ground that, after Immanuel Kant, I shall call *transcendentalism*. Transcendentalism is the view that, broadly speaking, does not deny the existence of a world independent of our mind, but acknowledges that the world as we experience it is strongly shaped by the fundamentals of our cognition. The brand of transcendentalism that I shall argue for has it that this is true in particular for the concept of *object* itself. This notion (plus, inevitably, all that comes with it) is a *human* way of cutting the pie of reality. What everyone knows to be true for colour and smell, and Kant said was also true for space and time, equally applies to the concept of object itself.

A Very Deep Thought—time for the reader to be on guard all the more, as depth is a dubious predictor of truth. To turn transcendentalism as intended here into a solid theory it must be possible to make it specific enough to provide explanations with a reasonable claim on being superior to those realism can offer. This book is written in the conviction that there are good arguments to make this case. This appears to be so, especially thanks to the centrality of the said notion of *cognitive system*, plus the availability of resources unheard-of in Kant’s days. Today we know something about how our brains actually work, and we know much more about computation and the structure of algorithms. With the help of insights like these, it should be possible to build and test theories about how the world appears to us transcendentially.

As might be expected, moving from realism to transcendentalism involves more than only the reconsideration of what an object is, what a property is, and what it means for the first to have the latter. Various more specific questions present themselves. Of all these, it is especially the following issues that are up for discussion in the chapters below:

- When do objects compose a whole? How does the latter relate to its parts?
- What are abstract objects? How do they relate to concrete ones?
- What are fictional objects? How can there be such things?
- How do mental representations relate to objects?

What these issues have in common is that they all touch on classic metaphysical themes and have proved quite persistent under realistic presumptions. I believe that, when seen through a transcendentalist lens, they will turn out to be far less intractable. These questions are essentially probing the concept of object, a concept which, despite being the most general of concepts, is not as bare as it looks. Indeed, to formulate answers to the questions above requires little more

than just spelling out the different aspects of this one concept, or so I contend, guided by the elements of first-order logic. A tiny bit of that job will be done below.

Chapter 1

The Draft Ontology

The world is everything that is, and every thing is a thing—an object. Starting out from this idea, the concept of object is adopted as the central theme of the book. Following Immanuel Kant in important respects, I advocate a transcendentalist (as opposed to realist) approach to metaphysics, according to which objects belong to our human perspective on the world, rather than to the world as it is in itself. The mind-dependence of objects is defended with the help of the concept of ‘rootedness’, and it is argued that first-order logic should be seen as an expression of the main principles guiding human cognition. A short survey is done of what I call ‘the draft ontology’, our intuitive picture of the world, that these principles are intended to describe. The chapter ends with a defence of the study of the human cognitive system as an essential element of the methodology to arrive at a more systematic ontology.

1.1 Things

Some say reality is made of things. Others, such as Wittgenstein in his *Tractatus*, say it is made of facts (a view that, in a different form, returns with Armstrong’s (1993) states of affairs). Still others, like Ladyman and Ross (2009), say that it is made of patterns.

There might be more options yet, but before deciding, let us reflect for a moment and realize how remarkable it is that we humans, from our armchairs, feel that we have the ability to give a sensible verdict on the matter in the first place. Reality is a big thing to grasp! Suppose I am among those who want to argue for ‘things’, how do I go about? Let there be things on my table presently, things that, since I am not hallucinating, I take to be part of reality. There are more things around me, and all over the planet. And even beyond: stars and galaxies as far the eye—if necessary aided by telescopes and satellites—can see.

All *things*, all providing quite impressive evidence to sustain my point. But could there be nothing else still further away? And short of that, can I be sure that, in between all things far and near, there are no other forms of reality? Are black holes things? Is dark matter? Even when we stay within eye's reach: what of space itself? What of gravity or any other force? What of time? How am I to find out?

The latter of these questions is of an *epistemological* nature, and different from (if closely related to) the *ontological* questions preceding it. In the course of developing the approach to be defended in this book, one of the things I shall argue for is that the epistemological questions are the crucial ones. Addressing the ontological question directly, however, is the more common pathway. It is well-agreed that there are aspects of reality that elude a description in terms of 'just' things (Russell, 2009). There are no unicorns, but a summing up of no matter how many things that do exist will not accommodate this fact, so it has to be added (for the sake of efficiency it would seem that we need some fact on top of all things to the effect that *this is all*, but never mind the details). This would provide a reason to prefer a fact-based ontology. But this solution has its own drawbacks: whoever states that all the world is facts, will have to accept things as well, since the latter are contained within the former. It is a fact that there are many things in this room, hence, *there are many things*. Things cannot be gotten rid of.

Or can they? Some authors go so far as even to reject the reality of things. In their work, aptly titled *Every Thing Must Go*, Ladyman and Ross (2009) argue that the world is not made of objects, but of *patterns*, a conception related to Dennett's (1991a) 'real patterns', and inspired by the work of Worrall (1989), who in turn defers to Poincaré (1905). All that matters, and all that is real, is *structure*. The authors base their counterintuitive thesis on what modern physics has revealed about the ultimate constituents of matter, in particular the incompatibility of those findings with the ways we normally think about 'things'.

Considerations like these doubtlessly make sense; however, with respect to the problem before us, they will not, to my mind, help us to an answer. This is not because there *is* no answer. There is one, and it is embarrassingly simple—so simple that it is bound to sound like a bad joke, a wordplay. Yet I think that unpacking the idea behind this answer bit by bit will yield a very different way of looking at reality, at our cognitive relation with reality, and, consequently, at the field of metaphysics and the way it relates to science. Here is the answer: reality is made of things, because *everything is a thing*. Things are or course things, but facts are also things, and so are patterns, whereas not everything is a fact (the Eiffel Tower is not a fact), nor is everything a pattern (the Eiffel Tower is not a pattern either). How can we know? Well, facts and patterns can be recognized by the behaviour of our references to them in language and thought. If Johnny is fat, then we can refer to this fact by saying things like: 'Paul told me that Johnny is fat,' but it makes no sense to say 'Paul told me that the Eiffel

Tower.’ And patterns are repeatable, but the Eiffel Tower is not (despite the large amount of copies, ranging from key-chains to actual buildings, in which many of the patterns *in* the Eiffel Tower are repeated; but none is the Eiffel Tower itself). Things, therefore, are the most general ontological type.

How can we find out if something is a thing? Things have their own characteristic ways in which they are referred to. They can be spoken about by fixing a reference, which we seem to do in three ways: by referring directly, by quantifying, and by declaring some thing to be *the same as* something else (which is then, at second glance, *not* something else). Things which are not the same, are different, whereby the possibility of a *plurality* of objects is given. But the most striking feature of things is that they can be subject to *predication*: they can be said to have certain properties, or have relations to other things. With these conditions in mind it is easy to show that e.g. facts and patterns are things as well. With respect to facts, much of the evidence for their being things can be found in a sentence like:

All the facts I heard at the meeting are quite serious, but most of them were the same I had learnt about earlier, such as that Jerry had been lying.

Patterns behave like things in an even more straightforward way (cf. also Berto and Tagliabue, 2014).

Be all this as it may, however, most of these diagnostics are provided by natural language. They are dependent on human conventions of referring to reality. To what extent can such criteria be taken seriously as a way of probing reality itself? The origin of the philosophical tradition of using language as the main metaphysical measuring device lies well before the ‘official’ Linguistic Turn, but the same is true for the appreciation of language as potentially misleading:

Of course, what confuses us is the uniform appearance of words when we hear them spoken or meet them in script and print. For their application is not presented to us so clearly. Especially when we are doing philosophy! (Wittgenstein, 2010)

‘The uniform appearance of words’: it is hard to think of a case to which this quote applies more than that of ‘thing’! Could anyone be blamed for thinking that ‘thing’ is just an extremely syncretic piece of vocabulary, multifunctional to the extent that nothing deep should be associated with it? Nevertheless, what I shall argue for in this thesis is the exact opposite of this: that reality is made of things is fundamental.

From now on I shall speak of ‘objects’ instead of ‘things’. This is a pragmatic, not a semantic shift: when I say ‘object’ I shall be speaking about my target of investigation, saving ‘thing’ for more general contexts. So this is the fundamental idea:

Reality is made of objects.

To see why, however, requires some expanding upon.

1.2 On Where to Start

Metaphysics is a strange enterprise. What is there? And what is it like? If answers to such questions come from a scientist, the general consensus is that the knowledge in question should be obtained by formulating theories and testing them by making observations. Never mind the epistemological subtleties; this is more or less the routine. These theories are worded in natural language plus whatever the formal sciences, logic and mathematics, bring in. So consider the theory: *All DNA-based organisms on earth share strongly similar genetic codes.* To test a hypothesis like this is not particularly easy, but as a theory it makes perfect sense. The question is this: Is it true that nowhere on the planet non-man-made self-replicating DNA-based structures can be found whose way of translating DNA into protein building-blocks is more than marginally different from that of known organisms? Since it will never be possible to search every location on earth or in the oceans, the answer will remain tentative. Also, there is room for clarification: how many codons must translate differently to disprove ‘strong similarity’? It is quite common that a need to revise the original wording of the theory arises due to new, unexpected findings.

There are, however, questions that the biologist in charge will *never* ask. Clearly, a DNA-codon is an object, however, issues like:

- What is an object?
- Do our senses inform us about objects?
- What does it mean for objects to be distinct?
- What does it mean for objects to share a property?
- How do abstract objects relate to their exemplifications?
- Can we rely on conclusions reached by induction?

(and so on) are not on the table. This is not because such questions are meaningless. In fact, the biologist must have answers to *every one* of them to get anywhere at all. If she could not trust her senses, recognize DNA-codons, tell them apart, notice similarities, categorize, and extrapolate, she would be lost. Not that she should be able to give explicit wording to the answer; what counts is the ability to apply the concept of object in a sensible way: e.g. appreciate that what is seen through a microscope are objects just as well as what is seen with the naked eye, and to distinguish glimmers and after-images from real objects.

Without all these conditions fulfilled, no answer with respect to the hypothesis about the DNA-code could get one inch off the ground; yet they are not the sort of questions that science counts as its business to answer. The answers to them are the silent axioms that scientists—as well as anyone else—may consider their birthright.

But what is given to some, is hard-fought by others. The tragedy of metaphysics is that the competence to answer the above questions is mostly taken to be constitutive of being a sensible human being. Nothing spectacular! Yet, even though the efforts to shed light on ontological issues are unlikely to change the course of science, only a minor departure from the mindset that leads us through daily life—by no means alien to science itself—is sufficient to show that they are real questions, open to investigation and argument. Even so, which are the foundations on which to build? Workers in the empirical sciences have their ontological convictions plus mathematics and logic, in recent decades comfortably packed in lots of software. Workers in the formal sciences have the axioms of arithmetic, plus sets and all that comes with them. What, if anything, can the metaphysician take for granted? Explanations must start somewhere. Where to start?

The point of departure that I would suggest is arrived at by first considering the epistemological question: ‘*How do we know?*’—in particular, how do we know that reality is made of objects? In the first section I gave some clues, but they looked like being about human capacities, rather than objects themselves. Here is similar question—not an unfamiliar one. How do we know that everything we will ever see will turn out to be coloured? This is, of course, because that is the way visible objects are presented to us. It is, in short, the way our brains-plus-senses, our cognitive system, works: our ‘mysterious’ omniscience in this respect turns completely trivial once we realize this. So how is it that we know that everything we will ever see, smell, bump into, hope for, put to work, think or talk about, is an object? What if the answer is the same? Does it make sense?

Thoughts about the subjectivity of reality—that what we experience is not the real world, but only some sort of (impoverished) representation of it, are *at least* as old as philosophy itself. Plato’s Cave is already a pretty well worked-out version of it, and one need not be surprised if ideas of this sort have occurred to the first speaking communities of the human race. The subjectivity of sense experience, the fallibility of it; the distinction between appearance and reality (which seems to have relevance even for the smarter among the animals¹), the possibility of false reports and false belief: all of these things give reason to question the objectivity of every way in which the world presents itself to us, and from there it is only a small step to question them all at once. There must be *some* wisdom in this line of thought, but to turn it into a constructive way

¹There is some evidence for intentional deception by primates, (e.g. De Waal, 1992; Hare et al., 2006)

of improving our metaphysical grasp of the world is a challenging task. It is all too easy to turn subjectivist ideas into a low-cost trick of explaining sort of everything. What is asked for is a way of approaching these issues that eschews embracing ‘big ideas’ because of their size only, and has a proper sense of detail.

Immanuel Kant must be seen as by far the greatest explorer of this territory, and my reason for preferring his work over that of many others is not *primarily* his doctrines *per se*, but first and foremost his attitude towards the subject and towards philosophy in general. Kant is not a man of extremes, and yet his conclusions are radical. In his striving to reconcile conflicting doctrines (realism and idealism, empiricism and rationalism) he arrives at a fundamental distinction between the world as it appears to us on the one hand, and as it is in itself, on the other. Our knowledge of objects cannot be but transcendental:

I apply the term *transcendental* to all knowledge which is not so much occupied with objects as with the mode of our knowledge of objects, so far as this mode of knowledge is possible *a priori*. Kant (1993, B25)

In working out this fundamental idea, Kant developed a system of great detail and sophistication, and even though few of his conceptions are still widely accepted in their original form, his general approach, as witnessed by the fact that his works have continued to be studied today, still strikes a chord. It is the idea that human cognition is not a passive recipient of reality as it is in itself, but contributes in shaping it, as it appears to us. There is something like a ‘template’ present in us beforehand (*‘a priori’*), for objects and their properties to fit in. That this is so for colour, indeed for sense perception in general, is relatively easy to see (the world is different for the colour-blind), but Kant’s important contribution was that the subjective component of what objects are goes considerably deeper, including (notably) their spatio-temporal constitution.

As Kant was careful to emphasize, this does not imply plain (Berkeleyan) idealism: the world does not cease to exist when we shut our eyes; not even when all human understanding comes to an end. It is not that our ideas of objects fail to refer to reality, only that this reference is not passive. It adds something to the world we live in. This point is of special importance, which becomes clear once we realize the deeply troublesome nature of a distinction like that between ‘primary’ and ‘secondary’ properties (as propagated by Locke). Whoever buys this distinction full force, buys herself into believing that the sofa in the corner is not really green; that ice is not really cold, and that when we lift a spoon, what we feel between our fingers is some sort of tactile go-between. But that is hard to accept as a report about the world we actually inhabit. In that world, such mind-dependent features (of which mind-dependence we are tolerably well aware!) seamlessly belong to the things that we *also* believe to exist autonomously, in ways *inaccessible* to our senses. So what I propose is this: just like colour is the way our vision shapes the world, the arrangement of reality in objects—over which

we quantify and to which we apply predicates—is the way our cognition shapes the world.

If this should be correct, then the task of the metaphysician becomes very different from that of the scientist or the mathematician. It becomes, to a large extent at least, a *reflexive* enterprise: for fundamental notions like that of object, our understanding of reality must include a grasp of this understanding itself as well. For, although we cannot know reality beyond the limits of our own understanding, to study and reflect on the latter seems to be within our power.

1.3 Layers and Roots

Before continuing on this path, however, let us spend a moment to see if the idea just presented is worth the purchase. The position is quite a radical one, after all. Not many people will have troubles with the idea that *some* of the objects in our thoughts and experience are mind-dependent, but to see this as the norm is another matter.

Of the category of generally accepted mind-dependent objects, *the average beach resort* is a good example. Such an ‘object’ may seem more like a mental tool for organizational purposes: talking about the average beach resort is an indirect way of speaking about specific beach resorts, which themselves are real. Or so it seems reasonable to hold—there is an ongoing dispute about whether it is possible, linguistically, to replace utterances about the former with utterances about the latter, but even if that should turn out hard (which it does! viz. e.g. Wetzels, 2009), there might be other ways to support the claim. Any way around, there is much attraction in the idea of a solid base of *real objects*, that exist unqualifiedly, upon which supervenes a jungle of metaphors, abstractions, summary representations, useful fictions, *façons de parler*, and all that jazz. In informal conversation it may be confusing, but in a context of full metaphysical strictness it seems fine to say, of things like the average beach resort, that they just do not exist. And the basis of things that *do* exist, it will commonly be held, is the totality of *physical objects*, existing in space and time. There is, so to speak, a layer of fundamental reality, superimposed over which there is a second layer of a mind-dependent nature.

For doctrines of this sort I shall use the general term *two-layer supervenience*—but not with the intention to endorse them! What I want to argue for instead is that there is a distinctive sense in which *all* objects are in layer two. Yet, and this explains my affinity with Kant, this is not intended as some form of anti-realism, for with the demise of layer one the ontological weight shifts to layer two as well. This could be elucidated in the following way: it is quite obvious that to say, at the dinner table, that the average beach resort is ‘unreal’ is not a helpful contribution to a discussion about the meaning of coastal tourism for the national economy or its role in polluting the seas. At least in such a context, the average beach resort is as real as any specific beach resort, even though

what underlies the ‘reality’ of the former is the reality of the latter. However, the scare quotes around the first occurrence of ‘reality’ in this sentence and the lacking of them around the second suggest that the latter reality (in contrast to the former) is built on unquestionably solid ground, since it belongs to particular beach resorts. Now, this seems less certain.

Let me introduce the term *rootedness* for the type of relation between the average beach resort and the particular beach resorts for which it seems to go proxy: the former is rooted in the latter. Some would prefer to speak of *grounding* for cases like this, and maybe this is correct, but this term has a broader use than would suit my present purposes. ‘Rootedness’ has the advantage that it does not yet have any theoretical load hung on it. What I intend to denote with the term is a relation that is quite widespread, and will turn out to be one of the central concepts in this book. It is not a mere linguistic relation: it does not depend on, e.g. the possibility of replacing utterances containing the one expression with utterances containing only the other (even though linguistic data can be used to support claims about rootedness). To clarify *what* it is, I shall at this point confine myself to the rather impressionistic scheme of giving examples. Not until Chapter 6, after having tested the idea in several contexts, an attempt at a sound definition will be made.

Armed with this piece of conceptual equipment, let us return to the example. No doubt the average beach resort is indeed rooted in individual beach resorts, but its roots do not end there. That reality is not two-layered does not mean that it is *not* layered; for let us consider Casa Marina, one of those resorts to which *the average beach resort* owes its relevance. Could its ontological credentials be called superior to the latter? Surely its existence boils down to that of the apartments, the hired chairs, a vaguely demarcated amount of sand until where the beach ends, and some of the air above the place. It may consist of two parts on a considerable distance. When we speak of ‘Casa Marina’, our words refer to the resort we see today, but also to the one that was here six years ago, with the same name but a different owner, after which period everyone said it had become ‘a completely different place’. The enterprise furthermore owes its being—in some respects its essence—to the legal provisions that make it a beach resort. And so on and so forth, and nowhere we stumble upon anything whose objecthood is primitively given. Any attempt to gain clarity about these issues will force us to consider, of the things in which Casa Marina is rooted, in what they are rooted in turn. And pretty much all of these things will place us before similar questions.

Tracing down the roots further and further, where do we end? We may well skip some intermediate layers and proceed directly to the particle level. Since the time of the Greek atomists it has been thought that some sort of basis of all physical objects could be positioned here: an immense lot of objects, very small but otherwise not unlike visible ones, of a limited number of basic types, able to entertain a couple of basic relations with one another. A lucid wording of this idea we thank to David Lewis (1987), who spoke of *Humean supervenience*:

It is the doctrine that all there is to the world is a vast mosaic of local matters of particular fact, just one little thing and then another. (...) We have geometry: a system of external relations of spatiotemporal distance between points. Maybe points of spacetime itself, maybe point-sized bits of matter or aether or fields, maybe both. And at those points we have local qualities: perfectly natural intrinsic properties which need nothing bigger than a point at which to be instantiated. For short: we have an arrangement of qualities. And that is all. There is no difference without difference in the arrangement of qualities. All else supervenes on that. (p. ix)

With the dawn of modern physics this model first looked like being vindicated; however, when quantum mechanics made its entrance it became clear that the smaller the particles get, the less they look and behave like physical objects of the kind we are familiar with². They seem to lose definiteness of position, of movement; even of identity. Also, the autonomous causal interplay that we associate with ‘normal’ physical objects appears to be absent: particles seem to ‘care’ whether they have been observed or not, and they have their behaviour coordinated with their ‘siblings’ miles away. Now, all this could still be interpreted as strange, indeed, but no reason to doubt the unqualified existence of, say electrons and protons. But electrons do not differ fundamentally from particles like photons and force-bearing particles, whose kinship with ‘typical’ objects is even less convincing. Both light and forces have been described, in theories preceding quantum mechanics and in general relativity, *without* making use of particles, and it is quite conceivable that future theories have still other ways to approach these phenomena.

Observations like these have led theorists to conclude that the essence of what science, basic physics in particular, reveals is *structure*, not the bearers of structure. This is what Worrall (1989) says about difference and continuity between Fresnel’s ‘luminiferous aether’ and Maxwell’s electromagnetic field:

There *was* an important element of continuity in the shift from Fresnel to Maxwell—and this was much more than a simple question of carrying over the successful empirical content into the new theory. At the same time it was rather less than a carrying over of the full theoretical content or full theoretical mechanisms (even in “approximate” form). And what was carried over can be captured without making the very far-fetched assumption of Hardin and Rosenberg that Fresnel’s theory was “really” about the electromagnetic field all along. There was continuity or accumulation in the shift, but the continuity is one of *form* or *structure*, not of content. (p. 117)

²The basic idea of Humean supervenience can be improved upon to create more harmony between it and modern physics (Lewis, 1994)

If knowledge about structure is indeed what is being accumulated in science, whereas notions about content (action at a distance, aether, fields, waves, particles) are just passing conceptions of what embodies it, then particles seem poor candidates to play the role of the objective basis of everything. This conclusion is reinforced by the fact that nothing that might turn up below the level of currently known particles is likely to have much better prospects in this respect than photons or gluons.

What speaks in favour of *structural realism* is that the structure that shows itself in nature, as it is approximated by the theories of fundamental physics, at least looks like being the most *objective* layer of all—objective in the sense that it is fairly invariant across different ways of interpreting ‘the phenomena’. As such, it has by far the best claim on being real. The reason why structure on this level is *not* a basis in the sense of two-layer supervenience is that structure cannot exist independently: it needs a substratum in which to manifest itself, *even* though we put our trust in the first more than in the second! Mathematics after all is a language in which to describe the world. It is not the world itself.

But then again, what is? The structure in question is rooted in the ‘content’ that Worrall speaks of. But this content is given to us in the shape of what is called ‘the phenomena’—that is, by *perceptual information*. Hence, following the roots all the way down something quite remarkable happens: just before they could be expected to end, they bend off the road to objectivity, to terminate in the wholly *subjective* realm of sensory perception. Realizing this may evoke some sympathy towards the project, as it was pursued in the early days of the twentieth century, to turn *sense-data* into an objective basis of the physical world (e.g. Russell, 1914). There is much to be said about this idea, but little that has not been said before. I shall not pursue this line of thought here. In conclusion, acknowledging that there is at least *something* perfectly right about structural realism, I take scepticism about some fundamental level of objects in general to be vindicated. There is no basic layer.

1.4 Principles

Kant’s epistemology is constructive, but against a pessimistic backdrop:

What may be the nature of objects considered as things in themselves and without reference to the receptivity of our sensibility is quite unknown to us. We know nothing more than our own mode of perceiving them, which is peculiar to us, and which, though not of necessity pertaining to every being, does so to human beings. With this alone we have to do. (Kant, 1993, A42/B59-60)

For Kant, what organizes this transcendental knowledge is the *Categories*. They are the first principles guiding the way our understanding works. Of course they

are not premises that we have chosen by an act of decision: they are hard-wired in our understanding, thus imposing their order on the world we live in.

Any transcendentalist approach to ontology will have to acknowledge some set of principles playing the role that Kant's system of the Categories plays in his. Whether the latter fills the bill seems doubtful: despite the vast amount of work that has been devoted to analysing and understanding Kant's logic, little of it has made it beyond the resulting literature. Today the 'mother of all logics' is not the Categories, but *first-order logic*; not because it is the strongest logical system (which it is not), but because in it the most fundamental of the logical notions are packed together in a seamless unity. Clearly, there are fundamental concepts in addition to these, but expressing them in new and more powerful logics virtually always proceeds by re-using the principles already present in the first-order formalism. Examples of this can be found in, e.g. second-order logic and plural quantification, where the conceptual shift essentially consists in treating predicates or pluralities of individuals, respectively, as individuals. A similar move can be discerned in the use of possible world semantics for modal logic, where there is just an extra layer of quantification, viz. over worlds. Although these (and many other) extended logics allow for much richer domains of interpretation, and although they are computationally vastly more complex than first-order logic, semantically they are still about *things, what properties they have and how they are related*; and in that sense (but only in that sense) do not go beyond what was already present in first-order logic.

The collection of notions that together underlie first-order logic, I therefore take to be the most important of the ordering principles that govern the world of objects as we find it to be. They form the core of what I shall call the *principles of the understanding*, whose role in our knowledge of the world is much like what Kant had in mind for the Categories, or so I propose. The idea is that these principles are hard-wired in the human cognitive system. We can reflect on them, but not revise them. It is virtually impossible ever to be sure that we have an exhaustive list of the principles of the understanding, so the following summing-up has no claim on completeness:

- object
- predication

Objects in first-order logic are normally referred to as 'individuals'. It is objects that are subject to predication: they can have properties and relations—henceforth 'predication' as well as 'property' will always include relations as well. Predications come in rich variety. The primitive phenomenon of predication should be distinguished from the issue of property (or relation) embodying *objects*, like *whiteness*, even though both issues are closely connected.

- identity

In first-order logic identity (sameness) is a special predicate—so special³, however, that it seems reasonable for it to be acknowledged as a separate principle.

The next principle of first-order logic would have to be *quantification*. For reasons to become clear in further chapters, however, I do not think that quantification corresponds to a single principle of the understanding. There appear to be (at least) two principles that together comprise the logical procedure of quantification. One of them is:

- existence

the other, generality. But, as I see it, generality is better associated with other principles of the understanding (see below).

Of course first-order logic includes propositional logic, thus we have all the propositional operations:

- conjunction
- disjunction
- negation
- implication
- (...)

The operations of propositional logic are certainly principles of the understanding, but since they do not immediately apply to the general notion of object, their discussion will not have a prominent place in this book. Finally, among the principles expressed by first-order (or any) logic we have:

- consequence
- hypotheticality

The essence of the notion of consequence is that knowledge can be extended in the absence of additional information. *Given* certain things to be the case, it is possible to conclude that some *other* thing must be the case as well. Clearly, all logical deduction thanks its existence to this principle, but *defeasible* inferences (e.g. by *induction*) have a similar structure (Pollock, 1987). Hypotheticality refers to the option to consider states of affairs, even if they are not (known to be) actually the case. Hypothetical scenarios are interesting in particular *because of* their consequences; the first without the second would be of little worth. Their combined action I take to be what constitutes *general* judgments. ‘All men are mortal’ can be rephrased as: ‘*Let* something be a man (hypotheticality), *then* it is mortal (consequence)’.

³Some authors, e.g. Quine (1986) think the identity sign should be viewed as a logical constant.

As said above, not all principles of the understanding are contained in first-order logic. Space and time, however elementary, play no part in it. Likewise, *moral* concepts appear to be principles of the understanding (and not the least of them!), but first-order logic is silent on good and evil. To repeat, all of these principles can be added to the formalism, by adding sentential operators (cf. deontic logic), special predicates and so on, meanwhile enriching the semantics to put them to work; but these are all real additions: no single trace of these principles appears in plain first-order logic. This seems to be different for *modal* concepts. The sentences of first-order logic—unless the language is enriched with diamonds and boxes—do not themselves express this dimension; arguably, however, there are connections between possibility and hypotheticality on the one hand, and between necessity and consequence, on the other. If ψ follows logically from ϕ , then there is clear sense in which ϕ *necessitates* ψ , and even though Hume’s view that conceivability (which comes close to hypotheticality) entails *possibility* is a strongly disputed one, there is an unmistakable conceptual kinship between both. Yet, most of us will find this too limited a rendering of modality. Those of a Lewisian persuasion will object to the whole idea of modality as a principle of the understanding: they take our modal beliefs to be a (potentially) correct representation of what is in fact the case at certain worlds. And short of buying such realism unabridged, many theorists at least opt for a more substantial account of modality.

However this may be, in the approach defended here modal notions are most naturally understood non-realistically. As will be shown below, this is not to say that there is no room for possible worlds: approaching modality by quantification over them has the enormous advantage of providing a language in which modal statements can be made with great exactness. However, one might as well choose to read the term ‘possible world’ as implying that there can be no such things prior to the analysis of what it means—for anything whatsoever—to have this property: *to be possible*. These sketchy remarks will be worked out in Chapter 6. Thus, risking some redundancy, the list will be extended with:

- modality

providing for the possibility that this principle cannot be fully analysed in terms of consequence and hypotheticality.

These are the principles of the understanding that I intend to use in building my argument for transcendentalism. The cornerstone of the system is of course the concept of *object* itself. Although the list is wanting in several respects, it arguably spans a fully featured ontology. The next step is to turn to the ontology as it actually steers our world and see whether, and to what extent, both match.

1.5 The Draft Ontology

In the methodology I advocate, the raw material from which any metaphysical analysis must arise is our pre-theoretical appreciation of the world. To this I shall refer as *the draft ontology*. The idea is not new; it is beautifully captured by Strawson's description in *Individuals*:

[T]here is a massive central core of human thinking which has no history—or none recorded in histories of human thought; there are categories and concepts which, in their most fundamental character, change not at all. Obviously these are not the specialities of the most refined thinking. They are the commonplaces of the least refined thinking; and yet are the indispensable core of the conceptual equipment of the most sophisticated human beings. It is with these, their interconnexions, and the structure that they form, that a descriptive metaphysics will be primarily concerned. (Strawson, 2002, p. 10)

Indeed, there is strong affinity from my part with Strawson's intent to practise what he calls a *descriptive* (as opposed to a *revisionary*) metaphysics. This is not to say, of course, that such a metaphysics is wholly passive and abstains from assuming structure below the surface. The distinction between both methods, as Strawson explicates, is relative, but nevertheless important. In the present approach the principles of the understanding are put forward to provide a description of the fundamental structure of the draft ontology, which itself (with due caution) may be considered as given.

For whoever has in mind the principles of the understanding, plus the formalism of first-order logic by which they were mainly inspired, the draft ontology as it comes to us from real life can be quite disturbing. To start with the good news: if what has been contended above is correct, then the draft ontology must be all about objects. And indeed, for all its looseness and inexactness it is remarkably consistent in this respect. By way of example, consider the following sentence:

1. After the rain had subsided, a silence fell over the land, and its many shades of green intensified.

The objects we find in this sentence are: the rain, a silence, the land, shades of green, the colour green itself. The example has been deliberately wrought so as to illustrate that many, if not most, of the items we find in the draft ontology are objects whose ontological service record, from a more systematic viewpoint, is doubtful. There may be various reasons for this to be so. 'The rain' can be conceived as a large number of water-drops in between the process of being formed inside clouds and reaching the ground, but the borders of this object, on all sides, are so blurry that every crisp determination of 'the rain' along these lines cannot be but hopelessly inappropriate. 'Silence' is even worse, since it is a mere

absence—an absence of *sound*, moreover, and what sort of object is that? ‘The land’, since it presumably extends, by and large, to the horizon, has borders that are not only vague, but also hugely arbitrary. Finally, ‘green’, and even more ‘shade of green’ name abstract, universal-like objects; a type of entities whose existence is contentious.

The twists and turns of draft ontology, however, reach a climax with sentences like:

2. Santa Claus does not exist.

where there is an object of which non-existence is predicated. In the draft-ontology this is all fine, but this is exactly why such phrases may seem to put tension on the whole idea of the draft ontology as an *ontology* in the first place. Is it a view on reality or just an edifice of language? Traditionally, ontologists have been strongly committed to the project of ‘cleaning up’ the inventory of really existing objects, not only by filtering out myth and superstition, but also by distinguishing between genuine objects and countless metaphors, paraphrases and sundry other potentially misleading ways of otherwise legitimate reference. Quine’s idea of *regimented theory* (Quine, 2013) is a notable example of such a project. My exposition in 1.3 already betrays my reservations with respect to this idea in general—it is not that a transcendentalist ontology leaves *no* room for metaphors and other ways of indirect reference, but this type of explanation seems less often and less urgently needed than it is in a realistic outlook. If the statement goes that ‘*x* is such-and-such,’ then the default assumption must be that this means that there *is* something *x* of which the said condition can be predicated. Only pressing reasons allow us to decide otherwise.

But the case of (2) looks pressing enough, one could say. Does speaking about an object imply commitment to its existence? Quine thought it does, as long as ‘speaking’ means: speaking in earnest; not in narratives or metaphors, nor in a way that is colloquial or sloppy. But there is a caveat. By his famous dictum ‘To be is to be the value of a variable’, Quine did not mean to say that whoever utters (2) is committed to the existence of Santa Claus and thus contradicts herself. What he meant is that there is always a way of using language such that the variables lines up with the ontological commitments:

We may say, for example, that some dogs are white and not thereby commit ourselves to recognizing either doghood or whiteness as entities. ‘Some dogs are white’ says that some things that are dogs are white; and, in order that this statement be true, the things over which the bound variable ‘something’ ranges must include some white dogs, but need not include doghood or whiteness. On the other hand, when we say that some zoological species are cross-fertile we are committing ourselves to recognizing as entities the several species themselves, abstract though they are. We remain so committed at least until we

devise some way of so paraphrasing the statement as to show that the seeming reference to species on the part of our bound variable was an avoidable manner of speaking. (Quine, 1980, p. 13)

Arguably, (2) can be dealt with ('Nothing Santa-Clauses', c.f. Quine, 1948), and there are some examples of tolerably successful rephrasings. The main problem with this view, however, appears to be the feasibility of the paraphrasing in the vast majority of other cases. *Pace* Quine, there do not seem to be many adherents to such a program left nowadays; it is in fact quite doubtful if there *is* any such thing as an avoidable manner of speaking. Paraphrases of sentences, cases of S' , non-identical to S , but guaranteed to have exactly the same meaning, are notoriously hard to find. Consider:

3. The odds are against us.

Many of us have reservations against committing ourselves fully to 'odds', but reformulating this quote in a way that will take all the worries away, if possible at all, is a highly specialized exercise, bound to bring a substantial theoretical context. The latter will bring controversial presuppositions and the whole 'paraphrase' will in all likelihood end up at a large distance from natural language. All this makes it hard to maintain that the original sentence were 'an avoidable manner of speaking'.

But the fact that it is almost impossible to avoid reference to 'odds' (or something much like it) should also make us wonder if we are entirely *uncommitted* to 'odds'. What this word refers to seems less real than a tree or a headache, but is no fantasy or metaphor. And in (2) there is even a sense in which we are committed to Santa Claus: after all we say something true about him. My proposal, therefore, is that what has to give way is not the logical form of the sentence, but the absoluteness of ontological commitment. The only type of occasion on which we speak about things without any commitment to their existence appears to be when we are engaged in activities like, say, singing a song or telling a story. A sentence like:

4. Long ago, there was a young girl called 'Cinderella'.

is uttered without the least intention to relate a fact. But this is already different for statements like:

5. Santa Claus lives in Greenland.

where some real information is delivered about a fictional person (I shall return to this point in Chapter 6).

What this means is that—slightly ironically—from a transcendentalist viewpoint Quine's 'To be is to be the value of a variable' can be taken *more* seriously than from a realistic one—in fact more seriously than Quine himself would have

allowed! If, as I would suggest, ontological commitment is understood more as a gradual matter (cf. also Turner, 2010; McDaniel, 2017), then the dictum can be read as stating precisely what is the link between first-order logic and ontology. Whatever can be the referent of what is quantified over and subject to predication *are objects*. In a transcendentalist ontology, this match is remarkably close to perfect.

1.6 On Where to Go

Acknowledging grades of commitment relieves the metaphysician from the grotesque stance that a good portion of the things common people bother about is unreal. But lest the temptation of practical advantage take over, it must be said that Quine’s original proposal still has its appeal. Let many of the objects as they figure in natural language be ‘unavoidable’, this does not mean that they elude all further explanation. We do not have to embrace two-layer supervenience, or re-embark on the old quest to find an ideal language in the sense of Leibniz or Russell, to resist the idea that all the objects we are in some way committed to are autonomous, monolithic entities; things that just are. Even though there may hardly be any convincing case of an ontologically suspect object which is *reducible* to more dependable ones, they are not unanalysable. They are open to inquiry, and the language we use to speak about them can be made more compact by pointing at the way they are related to others. Ideally this leads to a less subjective, less context-dependent and more detailed outlook on the world—essentially a more scientific one.

If what has been said above is on the right track, then one thing that will *not* follow from an analysis along these lines is that the less fundamental objects lose their ontological legitimacy, or even their relevance. Turning the draft ontology into a more systematic one does not involve purging it of the ‘creatures of darkness’—in fact I would recommend an almost infinite tolerance towards the things Quine found worthy of this characterization. Propositions and properties are perfectly real, or so I would suggest, and those creatures living in even gloomier areas (earth rays, spontaneous generation) are only just like Santa Claus. As Meinong famously claimed, despite their non-existence there *are* such things: they too have somehow contrived to earn their humble degree of commitment from our part. A more systematic ontology, therefore, is not one with fewer types of objects, but one in which the web of relations is more transparent and intelligible.

Here we arrive at the point where some things can be said about the method available for metaphysics as the reflexive activity it has been declared above. Many of the ways in which objects can be related are, in more or less direct ways, ultimately based on *perceptual* criteria. Being similar in some respect, or almost similar, or dissimilar; being bigger than something else, or cleverer, or of a lesser cardinality; being next to something else, above, inside, or historically preceding

it. Often they are supervenient on monadic properties which are directly perceptual; anyhow, the way to find out about such relations is to use our sense organs or measuring instruments. There are myriad such relations; they are the daily bread of science and generally fail to arouse great enthusiasm from the part of the community of metaphysicians.

But there is another class of relations—it is these which are the latter’s darlings—whose conceptualization does not in any way involve space, time, or perceptual categories. For an easy label I shall call them *purely ontic* relations. *Identity* is one of them, but the class of *mereological relations* is their flagship. Since it is straightforward to view identity as a special case of a mereological relation, it is fairly common to think that these are the only ones that are purely ontic in the sense defined here. *Grounding* relations, as they have gained more attention in recent years, are also purely ontic, but, partly due to their broad range of application, have not yet attained the status among metaphysicians that mereology has. One of the main ideas to be worked out in this book is that mereological relations are just one class among a larger category of purely ontic relations, and on a par with these. All of them are involved in ways in which objects can be rooted in others.

If conceived realistically, it is hard to make progress in analysing purely ontic relations. Take a simple case of mereological composition: a table made of a board plus four legs. Clearly, the table is rooted in these parts, but to understand what this means, focusing at the objects themselves, no matter how simple and transparent the setting appears to be, provides amazingly little in the way of a clue. One could look at the scene for hours and not attain any new insight whatsoever. All the time there is just a table—and also there is a board, and legs. However, if cases like these are studied within a transcendentalist framework, new anchors for explanation become available. Finding out what purely ontic relations are and where they can be expected should proceed by studying the working of our cognitive capacities, or so my proposal goes. We will have to touch on what lies, cognitively, behind the principles of the understanding themselves.

Kant’s Categories were established by *transcendental deduction*: the argument based on the idea that, for there be sensible experience at all, it must conform to the Categories. By only slightly modifying this thought we could say that, for a cognitive system to be able to function at all, it must impose a certain order on the subject matter it is concerned with. The principles of the understanding must be viewed in this light. Today we have methods and knowledge at our disposal which were unavailable to Kant: we have some insight into the actual working of our understanding ‘from the outside’. We know something about the way our brains function, and about the achievements of cognitive systems—including artificial ones—in general. There is a large body of knowledge about computational processes *in abstracto*. Therefore we have a ‘backdoor’ allowing us to look behind the veil that always kept the machinery of our understanding hidden from sight. Hence, there might be a chance to take Kant’s project, in slightly different form,

further than would have been thinkable in his time.

A disclaimer is nevertheless in place. In an overly optimistic mood one might come to believe that, once we take a glance behind the veil and ‘drop our prejudices’, something like a ‘final ontology’ might be accessible. Yet, in entering through the backdoor mentioned above, even though new insights to be gained there do have ontological relevance, our understanding itself remains the same. We can reflect on our inbuilt practice of viewing the world in terms of objects, but we cannot switch it off! The conditions under which our understanding acknowledges objects, along with what is predicated of them—*these* conditions we shall still describe in terms of objects, properties, and relations. In this setup there is an undeniable circularity. Fortunately, this circularity is not detrimental to all explanation, since it does reveal highly non-trivial aspects of why our ontology has the structure it does. Yet, the pessimistic bottom line as it appeared in Kant’s thought stays in place: even though our ontology can be made more systematic, we cannot go beyond a world of objects—not in philosophy, but not in science either. We can discover new species of plants and animals, and new types of particles (and cross out others), but, *pace* the insistence of Ladyman and Ross (2009) to do away with the old-fashioned concept of object, we cannot discover fundamentally new arrangements of reality. This is why science, even the most fundamental science, still talks about objects. They can be reinterpreted in terms of others, and the vocabulary to establish the relations between them can be trimmed and improved. But they will not go.

Unravelling the concept of object is, among other things, revealing its logic, of which, as was argued in Chapter 1, first-order logic is a fair approximation. Rather than being put to work as a formal calculus, this logic—or rather, an extension of it that grants ‘objecthood’ to a greater variety of symbols—is used as an interpreted symbolism: a (non-ideal) representation of the draft ontology. As for the reality behind the symbolism, it is argued that the human window on objects is an interplay of thought and perception, attested by language. Armed with this conceptual framework, some fundamental notions are reconsidered: *identity* (both in the sense of ‘thisness’ and ‘sameness’) and *plurality*.

2.1 The Logic of Objects

In the first chapter I have introduced the idea of the draft ontology, our pre-theoretic understanding of the world we find ourselves to live in. I have claimed that this world is best understood transcendentalistically, as a configuration of objects, ordered by what I have called the principles of the understanding. Also, I have tentatively identified the core of these principles with first-order logic.

So this is the contention, let us see if it is true. Or rather, let us see if intuitive statements about the world allow for a tolerably natural interpretation along these lines. First, it should be emphasized once more that, when there is mention of objects, they need not be of any specific *type*. The domain of objects ranges over teacups, games, couples, electrons, regions of space, ideas, unicorns, spirits, chances, trends, universals, numbers, sets, absences, magnitudes, facts, and what not. With an eye on the principles of the understanding, we can say that objects are whatever:

- can be quantified over, and

- can be subsumed under properties and/or connected by relations, and
- can be the same as another.

These conditions can be tested by linguistic means (by demonstrating felicitousness and compositionality), and also by applying logical notation, which here I shall use to paraphrase sentences in English. The natural choice for a language in which to express the draft ontology is, indeed, *natural* language, but formal language has the ability to make visible levels of structure that otherwise remain more or less concealed. This, at least, is the intuition that since Russell has been accompanied by the use of the term ‘logical form’, and will remain important for whoever does not take natural language to be the final word on our cognitive life.

The logical notation to be used starts with first-order logic, but to do justice to the rich variety of types of objects just mentioned, the language will be extended. The result, which I shall simply call FOL^+ , is obtained by taking the language of first-order logic and extending it with the following extra rule:

Whatever refers to an object, can assume the formal role of an individual symbol.

Needless to say that the resulting system is no longer first-order logic. Yet, no fundamentally new principles are added: all change is brought about by a rearrangement of the elements already present. As the deduction rules do not (yet) go beyond what first-order logic yields, FOL^+ can by no means be regarded as a full logic. So far there is no model for the symbolism but the draft ontology itself, which, due to inevitable ambiguities, is a shaky basis. FOL^+ is therefore no more than a symbolic language, whose expressions are intended to bring out the structure of the draft ontology.

Let us have a look at some cases:

6. This teacup is blue. $B(t)$

The only object referred to in this sentence is a teacup. Teacups, as physical items, are archetypical objects. There are teacups on the tray. Teacups can have many properties and relations, and the predicate blue could also apply to a kingfisher. This teacup is the one I bought last week (identity). Here everything is arranged in the most characteristic of ways.

7. George plays football. $P(g, f)$

George is fine as an object, but football is not a material object. Nevertheless, predication is straightforward: football is a sport, and it is popular. George plays some sport. It is the same sport that Carol plays. A quite similar case:

8. Barbara smells cinnamon. $S(b, c)$

Predication works: cinnamon is light brown. There is something that Barbara smells. If Barbara hates cinnamon, there is something that Barbara smells and hates. Cinnamon is the same as *kayu manis*. There is a difference, however between (7) and (8). Cinnamon *can* be construed in the same way as football (universal-like), but it is also possible to read that Barbara smells *some* cinnamon (e.g. when she smells that some cinnamon is left in a box), in which case the cinnamon is concrete *stuff* (I shall use this word as a term of art in the usual way). If she smells cinnamon in that sense, it is infelicitous to say that she smells ‘an object’.

Stuff is a special case. Every *amount* of stuff is an object, but just ‘stuff’ does not behave in quite the same way. Stuff is more like a plurality of objects. In quantification both show the same behaviour. Compare:

9. Barbara smells cinnamon.

10. There is some cinnamon that Barbara smells. $\exists x : C(x) \wedge S(b, x)$

11. All the cinnamon has been used. $\forall x : C(x) \rightarrow U(x)$

with:

12. Barbara sees trees.

13. There are some trees that Barbara sees. $\exists x : T(x) \wedge S(b, x)$

14. All the trees have been felled. $\forall x : T(x) \rightarrow F(x)$

Stuff also comes in more abstract guise, as in:

15. Barbara hears music.

16. During secondary school I did some physics.

These examples I shall take to be evidence that stuff is object-like enough to count as such in a wider sense (cf. also Laycock, 1972 and Nicolas, 2008). Also, statements about stuff pose no problems for being handled with first-order logic (the sentences (10) and (13) I take to have exactly the same meaning as those directly above them).

So much for concrete objects. Numbers, and more theoretical entities like sets, tensors, and operators, are evidently objects: in formal theories they are routinely subjected to quantification, predication, and naming (a practice that has strongly intensified with the coming of set theory). The harder cases for an object based ontology are mostly found in very natural expressions, one of which we have seen before:

17. Santa Claus does not exist. $\neg E(s)$

Despite philosophical misgivings about the existence predicate, it is readily accommodated by first-order logic. But there is an alternative available. For (17) Quine (1948) would have preferred:

18. Santa Claus does not exist. $\neg\exists x : S(x)$

(‘Nothing *Santa-Clauses*’). Notwithstanding the reasonability of the Quinean argument, there is little doubt that (17) is by far the most natural rendering when it comes to depicting the structure of this state of affairs as it appears in the draft ontology. Suppose Carl believes in Santa Claus, then, contrary to Quinean intuition, it is not infelicitous to say that *there is a person* who does not exist, but in whom Carl believes.

In Chapter 1 I have argued that *facts* should be counted among the objects. Description with the help of that-clauses is a mode of expression characteristic of facts, e.g.:

19. That Napoleon is short is a fact.

Facts can be quantified over (‘There are some facts we learned during the search.’), they have properties (facts can be serious or well-known), and identity (‘What you tell me now is the same Paul told me yesterday.’). Attempts to express (19) in first-order logic face the problem of fitting the that-clause into the formalism. The draft ontology is somewhat equivocal as to the objectual nature of such a clause. In (19) it is fairly convincing. However, in:

20. Soldier Luc does not know that Napoleon is short.

what is being stated looks more like a relation between Soldier Luc and Napoleon. That the objectual classification of ‘that Napoleon was short’ is nevertheless natural, can be seen when considering:

21. There is something Soldier Luc does not know. $\exists x : \neg K(l, x)$

Subsequently, using the extra rule of FOL⁺ on page 28, we could write:

22. $x = S(n)$.

to express that this thing Soldier Luc does not know is that Napoleon is short.

Allowing expressions of this sort is stretching the syntax of first-order logic pretty far, but does not lead to interpretational ambiguities. One might have expected something like $x = [S(n)]$, but that would equal x to some *sentence*, which in turn expresses the fact that Napoleon was short. As $S(n)$ expresses this fact *directly*, it seems legitimate, if unusual, to write (22). Thus we get:

23. There is something Soldier Luc does not know, which is that Napoleon is short. $\exists x : \neg K(l, x) \wedge x = S(n)$

24. Soldier Luc does not know that Napoleon is short. $\neg K(l, S(n))$

And also:

25. That Napoleon is short is a fact. $F(S(n))$

This embedding of propositional expressions in object positions is what is obtained by following the seemingly natural behaviour of the draft ontology. There are other notational moves which can be motivated in this way. Consider:

26. Soldier Luc is courageous. $C(l)$

27. Courage is a virtue. $V(C)$

Proceeding in this way would make for a (quite atypical) sort of higher-order language.

In 1.4 I have mentioned *modality* as a principle which is not (wholly) included in first-order logic. Consider an extension of FOL⁺ that includes the modal operators. Here it is possible to show how the reification of *possibilities* takes place:

28. Napoleon might be dead. $\diamond D(n)$

29. Soldier Luc knows that Napoleon might be dead. $K(l, \diamond D(n))$

30. There is a possibility that Napoleon is dead. $\exists x : x = \diamond D(n)$

Notice that logical expressions like these do not enforce one ontological view to the exclusion of others. For whoever wants to stick to modal realism: to express possible world semantics in the same style is easy:

31. There is a world at which Napoleon is dead. $\exists x : T(D(n), x)$

(where $T(x, y)$ should be read as ‘ x is true at world y ’).

So much for the formal regimentation of the draft ontology. But then, are there examples of things that are *not* objects? Waiving the clumsy paradoxicality of the question, it is clear what is intended. And indeed, there are such things. What about:

32. It is freezing.

33. Nothing is both blue and tasty.

The only word in (32) that has a chance of referring to an object is ‘it’, but does it? ‘It’ could also be snowing, but that does not mean that there is something that is both freezing and snowing, nor that there is a thing that is freezing, and another that is snowing. There is not even so much as *anything* that is freezing. It is just *freezing*. In first-order logic (32) has no interpretation that preserves the logical form of (32), not even approximately. Both:

34. $*F(i)$

and:

35. $*\exists x : F(x)$

are wide of the mark. The reason why (32) poses no threat to the objectual constitution of the draft ontology is that nothing ontologically extraordinary is at issue. The phrase is easily elucidated by something like:

36. The air temperature is below zero.

where the air temperature and zero are regular objects, connected by a regular relation. And as for (33), surely:

37. $*B(n) \wedge T(n)$

will not do, *pace* Heidegger. Here we *must* go Quinean and use existential quantification in the negative. ‘Nothing’ is *not* like ‘Santa Claus’!

2.2 Inferential Reasoning

So much for singular statements in FOL^+ . For FOL^+ to constitute a *logic*, however, more would be needed. To the extent that languages respect logical form, they support *inferential reasoning*. Formal languages are mostly designed with exactly this in mind: there is a certain harmony between their semantics (what they are about) and their syntax (how they are built up), thanks to which there can be a set of rules to generate new sentences that retain certain semantic virtues, in particular truth. All this is of course commonplace; the reason why I spell it out here is to lay the groundwork for a closer look at the phenomenon on inferential reasoning itself, insofar as transcendentalism may have something to say about it—as should surely be expected. This task will be taken up in more depth in Chapter 3; in this section the focus is on exploring the ground.

Any investigation into the nature of inferential reasoning—which is possible thanks to what in (1.4) was called the principle of *consequence*—should, in my opinion, start with the realization that the whole phenomenon is little less than a miracle. There is really nothing obvious or trivial about there being such a thing as a valid inference in the first place. The primary function of sentences in any language, including symbolic ones, is to state *what is the case*—why would they have the additional option to proceed from there to what *else* is the case: extra knowledge without extra information? And what could make it possible? Kant’s whole epistemological survey started with his perplexity about the possibility of ‘synthetic’ knowledge *a priori*, but even the remaining category of ‘analytic’ judgements may be far less trivial than it seems—a stone lion, after all, is *not*

a lion (Partee, 2007)¹. And yet such options are deeply inherent to the draft ontology. The human cognitive system, in which it is based, appears to provide them as of necessity.

In the course of history, the study of inferential reasoning has succeeded in explicating several of its manifestations by turning them into rules figuring in formalisms. Many of these reflect our most indubitable convictions: the principle of non-contradiction, the rules of syllogistic and propositional logic, those regarding predication and quantification, and so on. It seems fair to read this as evidence that these rules mirror fundamental aspects of our cognitive architecture. But this, I would suggest, is not the only reason that they are the ones singled out in particular: another one is that they belong to a privileged class of rules. This can be explained in the following way. A system of logical derivation as commonly understood is a set of rules which can be implemented as *mechanical* symbol manipulations of *modest* size and complexity. The latter means that the rules are comparatively simple: simple enough to be handled by a human interpreter, armed with nothing but a sheet of paper and a ballpoint. No supercomputer needed! Furthermore, that the rules are ‘mechanical’ means that they are fixed, stable, and insensitive to external influence.

First-order logic is a system in which a notable set of exactly this sort of rules come together. As the inner shell of FOL⁺ is first-order logic, the continued validity of what is inherited from it may, in all likelihood, still be counted on, even where the sentences are no longer well-formed according to the old standards. E.g. arguments like:

Whatever soldier Luc writes, is not so.	$\forall x : W(l, x) \rightarrow \neg x$
Soldier Luc writes that Napoleon is dead.	$W(l, D(n))$
<i>therefore</i> : Napoleon is not dead	$\neg D(n)$

look pretty safe. But even if we would have proof of some such calculus for deduction in FOL⁺, then this set of computationally uncomplicated prescripts would not even start to cover all the options for concluding facts from other facts in the draft ontology. There appear to be always *more* inferential principles than have been described at any point in the history of logic. In part, this is due to the introduction of new symbolic resources—but only in part. Many (and not the most insignificant) of the remaining entailment rules are not strictly deductive: they turn out to be relative or *defeasible* (cf. McDermott and Doyle, 1980; Maher et al., 2001; Governatori et al., 2004).

What I would argue for is that such rules are not at all unexpected. Why could there not be inferential principles that we cannot handle with pen and paper, but that our brains can put to work? They could be just like the familiar rules for

¹Of course the analytic-synthetic distinction as presented by Kant has been fundamentally criticized, notably by Quine. But this only strengthens the case against the triviality of logical inference of whatever sort.

logical deduction, but so complicated, and demanding so much computational power, that they have not made it to the logical canon. What makes them of special interest, however, is that many of these ‘big rules’ appear to have ‘small rules’ that go proxy for them, but are simple enough to function as rules in a quasi-deductive system. Let me illustrate this idea with the following example. In surveys or tests one can find multiple choice questions like:

What is the capital of The Netherlands?

- A. Rotterdam
- B. The Hague
- C. Brussels
- D. None of the answers is correct

Those who know their geography know that D is the right answer. Strictly speaking however, that D is true would mean that one of the answers *is* correct, and D would be false. To appreciate that D has to be chosen anyway, one has to embrace the silent pragmatical condition that ‘the answers’ only refers to A through C. By far most of the participants do this without thinking, not aware of doing anything irregular. The small rule says: ‘If none of the answers is correct, D is true’; the big rule provides a pragmatic way to escape the paradox involved in the question, following some sort of meta-instruction that demands that there be a sensible outcome, come what may.

As said, big rules *could* be like ‘normal’ deductive rules, only a bit more complex. But arguably some big rules are way bigger than that: so big that they are more like a highly intricate natural pattern, or even so big that there is no rule in any customary sense—only a huge pile of particular facts with some regularity in them. Any way around, often the big rule is too big for any cognitive system to grasp in its entirety. But—and this is what still makes it worthwhile—there *is* a small rule approximating it! Often something like a one-liner, imperfect but by no means useless. Life is full of small rules—henceforth I shall refer to them as *intuitive rules* or *intuitive reasoning*. If this is correct, then classical logic, the set of ‘privileged rules’ mentioned above, consists of those rules that are simple *and* infallible at the same time, whereas ‘small rules’ are simple, yet vulnerable for revision. Clearly, numerous proposals have been made for new logics about all those topics in life that classical logic fails to handle in a convincing way, some of them successful. But small rules never grow big: they cannot be made part of a real deductive logic. This need not be a problem: intuitive rules, plus the (non-classical) logics that formalize them, can be useful in exactly the same way that they are in daily life.

Intuitive rules are interesting because they provide insight into the workings of our cognitive system, and therefore, into why the world appears to us the way it does. Although education or training may affect people’s abilities and

dispositions with respect to intuitive reasoning, it seems that their content is not commonly transferred in verbal form. It is not learnt like the grammar of a foreign language, but like the grammar of one's mother tongue, and its imperative can be strongly felt. Nevertheless, in contrast to the principles of the understanding, intuitive rules form a system of convictions—beliefs. We do not have the option to sidestep the principles of the understanding, but, even though they are often of a stubborn, almost instinctive kind, we do seem to have the ability, without contradiction, to liberate ourselves from at least some of the inclinations and persuasions of intuitive reasoning.

2.3 Perception, Conception, Conversation

On the world of objects, human understanding has three vantage points. They are:

- perception
- thought
- language

It is these in which our contact with the world manifests itself. We can speak about an object, think about it, and experience it. Of course we can also wish for it, fear it, create it, and so on; but for describing our cognitive relation to objects I take these three abilities to be fundamental. Of them, language is easiest accessible to investigation. Understanding thought is, at least in philosophy, typically pursued by studying its representations in language, be it natural or artificial, but thought and perception are also part of the subject matter of psychology and neuroscience, whose results thus provide another philosophically important source of information.

In the history of philosophy all three have known their champions. In the Enlightenment a good deal of the debate was between empiricists and rationalists, while the twentieth century saw the Linguistic Turn, much to the detriment of the ontologically inspired interest in thought and perception. For those, including myself, who strongly believe in the importance of both of the latter, it can be tempting to view language as derivative: as though linguistic utterances were mere passive representations of our mental life. This is of course not so. Language, formal language in particular, has a measure of independence from the individual user: it has its own structure, grammar, rules, and norms. As being part of a social phenomenon, no linguistic utterance is ever exclusively one's own. Having said this, however, I take the draft ontology itself to be something that resides first and foremost in thought and perception; hence, when studying it, language is a source of information, rather than a constitutive factor in itself.

The window that language offers is not perfectly transparent. Yet, it seems reasonable to place the burden of proof with whoever wants to hypothesize ‘cracks’ in the analogy between, in particular, thought and language: the default assumption will be that parallelisms in language indicate parallelisms in thought. Speaking about perception, on the other hand, is arguably uttering *thoughts about* perception, since a direct way for perception to translate into language does not seem to be available. Introspective reports about perceptual experiences are therefore error-prone in two ways: not only can language misrepresent thought, also thought can be wrong about what is actually perceived. Philosophical discussion about what perceptual experiences ‘really are’ (which is not at issue here) aside, intuitive ideas about what we perceive are often fallacious in quite trivial ways. The human brain is known to fill in many details to ‘complete’ the available perceptual information (e.g. Ramachandran and Gregory, 1991). But even though this makes it hard to distinguish between ‘pure sense data’ and additions and interpretations, much of what is described by subjects can be taken seriously enough, especially if it is not too complex. Reports about hearing an interrupted beep or seeing a purple line against a white background, will by most of us be accepted it as reliable.

Introspective reports, however, are just one source of information about thought and perception. The other, which for the transcendentalist is at least as important methodologically, is the study of the architecture of cognitive systems in general and the human brain plus sensorium in particular. Here thought should be associated with the process of computation, and perception, with that of introducing information from outside on which these computations are performed (Fodor, 1983; Evans et al., 2019). Considerations springing from this approach may help to settle certain issues that introspection will never be able convincingly to deal with. For one example: the introspectively awkward notion of ‘pure’ perception just mentioned can be given a clear meaning in functional terms, since the external information (from the senses) must be introduced into the computation at some point, previous to any calculation. Notice, however, that, although objects are ubiquitous in thought and perception, there is no obvious way in which the concept of object makes itself known among the functional notions pertaining to cognitive systems. This, then, will be one of the main questions that a transcendentalist ontology will have to come to terms with. In Chapter 4 a proposal to this end will be delivered.

In natural language the subject-predicate construction does not discriminate between *types* of objects. Although we may feel there to be an ontological distinction between ‘Sarah is tall’ and ‘Three is prime’, the syntactic form of both expressions is the same. However, Sarah is an object that can be perceived in a very direct way: she can be seen, heard, and felt, all of which is not in the same way possible with the number three. There is also a difference between what is predicated in both cases. The property of being tall is determined by experience: it is, and will therefore be called *perceptual*, contrary to that of being prime, which

is *conceptual*. One might think that this has to do with the nature of the object, but this does not seem to be so. Sarah is tall, but also self-identical (a conceptual property), whereas three, apart from being prime, is also the number of pencils on my table (a perceptual property). As I shall argue below, it is the distinction between *properties* which is fundamental. The parallel distinction as we feel it to apply to the objects themselves (three is a conceptual, Sarah, a perceptual object) is the one which is derivative.

Yet, this distinction, if fundamental, is by no means sharp. That there is a role for perception in concluding that Sarah is tall is obvious, also that to determine that three is prime is a pure computation, with no need for external data. But there is a conceptual element in almost every perceptual observation. If Sarah stands upright, then her length is the height of the part of her body which is highest above the ground; not so if she lies on the beach. Hence, it takes some pattern recognition when processing the visual information so as to determine whether Sarah is tall. Also, the decision whether a certain length must count as ‘tall’ is probably taken against a background of past experience from which something like a norm is extracted (Kamp, 2013; Toledo and Sassoon, 2011; Van Rooij, 2011). But even if the perceptual source of some observation is pretty obvious, the transition between perception and conception appears to be gradual. Consider seeing leaves on a tree. Let there be a branch with two leaves of equal size, growing in two opposite directions. Having seen the leaves themselves, we almost immediately see *symmetry*, a property that ‘should’ be conceptual (it comes on top of the bare visual impression of shape), but strikes us as occurring prior to any reflexive thought. But now take a branch with 139 leaves. We quickly see that there are lots of leaves (more than twenty and less than thousand of them), but, unless we make some additional effort, we cannot see that their number is 139, let alone that it is prime. Clearly, the latter properties are conceptual.

Even though it is difficult, for these reasons, to find out what is in ‘perception proper’; of certain aspects of what we experience is it fairly obvious that they do not belong there, since they could be left out with perception staying in place. Peeling off the conceptual layers in this way, we may conclude that at the very outset objects are given by their qualities only. Also, at that stage every object is *this*, but such ‘thisness’ is not the same as what is commonly called *haecceity*, since it has no permanence beyond the momentary experience. If Sarah walks out of the room and returns a moment later, our perceptions surely give us *reason* to think that it is the same person on both occasions, but do not deliver this fact directly. A very similar doppelgänger would not have caused a different experience. Furthermore, perception proper appears to be largely devoid of tense: all we see, hear, feel, and so on, is in the present moment; the past has to be recalled, and the future, surmised. Likewise it entirely lacks *modality*: although we can think that Sarah *could* have been short rather than tall, experience gives no clues in this respect: things are as they actually are, and that is all.

It should be noticed that I have called ‘perception proper’ is the result of an intuitive thought experiment and does not have a lot to do with the idealised notion of ‘pure sense data’. The utility of this exercise lies not so much in understanding perception itself, but in understanding the layers of conceptualization that turn it into the experience of objects.

2.4 Being

Of the principles of the understanding, there are three that in English are expressed with the help of the verb ‘to be’: *existence*, *predication*, and *identity*.

- Sugar is sweet. (predication)
- There are numbers between zero and one. (existence)
- Scott is the author of *Waverley*. (identity)

This threefold role of ‘to be’ cannot boast full generality, since things are not always expressed in the same way as they are in English. In fact, ‘to be’ and its relatives show quite a diverse behaviour throughout languages: many of them have expressions for existential quantification which do not contain the customary analogue of ‘to be’ (French: ‘*Il y a des nombres...*’, German: ‘*Es gibt Zahlen...*’), many tend to leave out the analogue of ‘to be’ in predication and identity (Russian, Latin), and some (like Spanish and Italian) have more than one verb translatable as ‘to be’. Observations like these may make one sympathize with Frege (1892a) and Russell (2009, p. 81-82) as to their opinion that ‘to be’ is essentially ambiguous. Here I shall nevertheless defend the idea that, for languages that behave like English, it is no coincidence that this verb plays all of these roles at once. It appears to function as a sort of cement securing some very elementary joints in the linguistic edifice: joints that betray their analogues in thought, and thus, in the draft ontology.

Predication is what has been discussed in the last section, where it was identified as being the result of perception and thought. In a transcendentalist ontology properties *just are* whatever is predicated to objects. There is no conceptual distinction between the two. This is in sharp contrast to the realistic view, by which it is altogether reasonable to think that there could be properties forever hidden from our (or anyone’s) grasp, and not entirely unlike Michael Dummett’s (1991) anti-realism (albeit that, *pace* Dummett, the activity that makes us distinguish properties is based in cognition, and only derivatively in language).

The usual, realistic way of thinking about predication is that there are objects first, and once they are in place, something can be predicated of them (‘Socrates is bald’ presumes the existence of Socrates). From a transcendentalist point of view, however, the most natural way to understand objects is to see them as mere *nodes in the conceptual network*—the network of predications as our cognition maintains

it. If this idea is correct, then predication comes first, conceptually. Properties and relations are not free-floating, though: they form a network in which they are held together in certain configurations. According to this picture, what does the holding together is what we call objects. Here it becomes clear how this construal of objects is helpful for understanding the role of ‘being’ in the draft ontology: it marks the occasions where, in the conceptual network, predications are active and different predications are connected to one another².

One such occasion is the cognitive analogue of existential quantification. The architects of first-order logic have chosen not to provide for means of introducing individuals ‘barely’ ($\exists x$ as a wff), even though it is neither impossible nor unsound for this option to be allowed. This decision makes first-order logic follow natural language, but also makes sense because, in models for first-order logic, there is always presumed to be a non-empty domain ($\exists x$ would be trivially true). For our present purposes this situation is fortunate, since, psychologically, it seems true that objects never enter our awareness entirely out of context. As was remarked in 2.3, when objects make their very first appearance in perception, they can only do so couched in a sense modality. And as for linguistic form: even an expression like ‘There is something,’ is not as blank as it looks. What it technically says is that something is *at some place* (‘there’), which could be interpreted as revealing a pre-theoretic conviction that to exist means ‘to be located’³.

Thus, the way in which new objects are introduced is by a predication:

38. $\exists x : P(x)$

As said above, the standard conception goes that there is, first, an object—we might call it a —and of it property P is predicated, resulting in:

39. $P(a)$

So something like (39) would be the reality underlying the truth of (38). It is important to see, however, that from the perspective of the cognitive agent, x has no identity whatsoever over and above its being P . Transcendentally, only this is relevant. There might of course be several objects which are P , but, by the epistemology that the transcendentalist picture entails, to know $\exists x : P(x)$ does *not* mean of one of them, a , to know that $P(a)$; it just means to know that *there is something* P ⁴ (cf. Aguilar-Guevara and Zwarts, 2010).

²The *structure* of this construal is similar to various brands of structuralism, in mathematics (Benacerraf, 1965) and in science (French, 2014; Ladyman and Ross, 2009)

³Although French and German do not include this aspect in the same way, the expressions in these languages (*Il y a...*, *Es gibt...*, litt. ‘It has ... there’, ‘It gives...’) suggest the original intention to place the newly introduced object in some minimal context (it is part of ‘it’: the world as a whole?).

⁴With the qualification that it is cognition rather than language that is at issue here, it should be remarked that Quine’s dictum (‘To be is to be the value of a variable’) is a better representation of this fact than the neo-Fregean motto ‘To be is to be the referent of a singular term in a true sentence.’ One of the problems with the latter is that, so far, there is no obvious justification for the being *singular* of any term. This point will be given due attention below.

This leaves us with the obligation to give an analysis of what a , the individual constant, stands for. If objects are mere nodes in the conceptual network, it would be nice to know what accounts for our intuitions about the *identity* of objects, of which naming is a major expression in natural language. No matter how the meaning of existential quantification is analysed, what a in $P(a)$ intends to express seems very natural: we know a by its name, ‘ a ’. Names constitute a separate concept in first-order logic. Expressions containing them suggest that we have a way of knowing some thing whereby its introduction as a *new* object, as in $\exists x : P(x)$, is bypassed, or at least put at a distance.

But, as the quote goes, *What’s in a name?* Names can be viewed as predications (*being Sarah* is just fine as a property), but a predication like *being red* seems to be importantly different from one like *being Sarah*. That something is red can be checked using the senses; also that something has the shape of a tulip, contains sugar, or is radioactive—even though the senses need some help from artificial extensions in the latter case. All such properties, i.e. properties (or relations) that can be determined with senses plus auxiliary devices, I shall refer to as being *qualitative*. Being Sarah, however, is being that very person—in fact, *that very object*. Now, what names purport to stand for (and here I follow the Kripke’s (1972) orthodoxy) is exactly this: a unique addressee, rather than a unique complex of qualitative properties. This view has much to its recommendation intuitively, but at the same time it entails an enigma. If *being Sarah*, as a haecceistic property, involves aspects that absolutely transcend what can be caught by the senses, then *how are we ever able to tell that the person before us is Sarah?* And for those who are willing to trade false certainty for true reliability: how could it even be as much as probable? Exact likeness is not a great compass here. Whenever I see a wasp I am positively sure that it looks and sounds the same as *every* other wasp I have ever seen, but I feel no inclination to think I saw it before.

For this problem there is a solution on the market, the doctrine of *content externalism* (Putnam, 1975; Fodor, 1994), which can be applied to both language and thought. According to this view, predications do not find their content in the human mind, but ‘out there’, in the things themselves. ‘*Meanings just ain’t in the head!*’, in Putnam’s famous one-liner. The general idea behind externalism is more or less the opposite of the transcendentalist view argued for in this book, yet it does seem to provide an answer to the problem just posed: the content of our thought about Sarah is simply the (unique!) person that causes our sensory experiences right now. Even so, this initially convincing picture has problems of its own. Externalism has been criticized in a number of ways, but the most relevant issue in the present context is that this haecceistic property of being-this-very-object is entirely *epiphenomenal*, since an object is causally efficacious only thanks to its qualitative properties.

The (in my view) misconception causing this difficulty is the widespread conviction that in situations like the one described by $P(a)$ it is possible to take

away the property P —in fact, to take away *every* qualitative predication—and thereby isolate the object a , which has something like a ‘residual identity’ (Jubien (1993) speaks of ‘the fallacy of reference’ in this regard). In Aristotelian terms it is the identity that belongs to the *substance* alone, separated from its form. But taking all form away, nothing is left that has any cognitive import whatsoever, which, at least transcendently, means that *nothing* is left. (Realistically, the assumption of fully non-qualitative properties is not incoherent, but their apparent uselessness in any explanation of human behaviour or the evolution of our cognitive abilities does not speak in their favour as theoretical assets.)

Since names are evidently not empty, something else must be responsible for their functioning. As the idea that names refer haecceistically seems to stand firm, the possible way out might be to reconsider the way haecceity itself is understood. It is common to place haecceistic properties opposite qualitative one’s, not allowing for overlap. What I want to propose instead—to be worked out in more detail in Chapter 6—is that the mechanism by which names ‘find’ their referent is not fundamentally different from the way predicates like *red* do, but rather at the opposite extreme of a gradual scale of increasing relationality. *Being Sarah* is being, say, my neighbourhood girl, whom I came to know years ago, when she was introduced to me by existential quantification under some very intricate combination of description and/or perceptual experience. Seeing that something is red requires very little context; seeing that something is Sarah requires facing a vastly extended web of other related objects, in which Sarah occupies a certain position. In the interpretation advocated here, to me Sarah *just is* what occupies this position.

This high level of relationality has an unexpected effect. Since the conceptual web changes in time (the world changes and so does our knowledge and focus), the exact definition of the nodes that are accessed by labels like ‘Sarah’, or other names, is unstable in time. Whereas ‘red’ sticks to a relatively well-circumscribed conceptual area, ‘Sarah’ marks a very complex aggregate of conditions, which is also, unlike ‘red’⁵, *flexible* and therefore open to change. The labelling of this aggregate of conditions is not necessarily there for the sake of this specific aggregate. Presumably it is there, as far as circumstances allow, to secure some continuity. Through time, the label ‘Sarah’ traces a path of least change, that of a node in the network as it develops in time. In our experience this has the result that the *object* thus fixed *remains the same*, even if new facts should become available that undermine the former conception. An analysis along these lines provides, I believe, an explanation for most of the Kripkean observations about the failure of fixed descriptions to account for reference as it actually happens. (I believe that this idea can be extended to include the way *pronouns* in natural

⁵The meaning of *all* words is probably subject to drift due to changing circumstances in the world, but that of ‘red’ is unusually stable: it is the only colour word that can be traced to a Proto-Indo-European origin (OED, 2020).

language work; cf. Sommers, 1982.)

This theory also explains why it would be unfortunate to treat names as predications in formal expressions, e.g. to view $P(a)$ as a shorthand for:

$$40. \exists x : P(x) \wedge a(x)$$

where a is a (very complex and strongly relational) qualitative description, putting constraints on x . A predicate like a would suggest a combination of properties and relations that, however complex, is constant rather than changeable. To extend the logic with *inherently* unstable predicates, however, seems like a doubtful move, in particular if the predicate's function is to represent some object in a stable way. Therefore the use of individual constants like a , appears to be preferable. If the above interpretation is correct, naming is not a primitive, but analysable in terms of predication.

There is, however, one highly important aspect of the matter about to slip through the cracks: what makes the referent of the name Sarah *unique*? To address this point, we first have to look at the third item in the above list: identity. It is a relation, thus falls under the general heading of 'predicate', but it is special enough to deserve separate discussion. It is the only predicate for which specific deductive rules are almost standardly included in implementations of first-order logic. In the draft ontology the fact that object x can be the same as, or different from, object y , is an inalienable aspect of the nature of objects. And with respect to the conceptual network identity is, more than anything else, what endows it with structure. It is what knots (bundles) the predicates together.

Ontologically, identity brings its own particular challenges. Is identity primitive? If so, we run into the same problem as that with respect to haecceity: if it is not perceptual, not even indirectly, then it appears to be epiphenomenal. But identity *can* be perceptual. Consider a melon: just to perceive it to be itself is not the best example, since very little appears to happen psychologically, or it must be the quite intellectual and superfluous action of actively comprehending this to be so. Normally seeing an object neither causes nor requires the operation of the concept of identity. However, what perception does connect in a non-trivial way is *this round object* and *this green object*. When we pick up the melon, we extend the identification to *this heavy object*, which is even less trivial, because a new sense modality is invoked. Also, I see someone's lips move and hear the sound of words, and I accept both as belonging to one and the same entity: a person speaking. Or I look at a scratch on my hand and feel this also to be the origin of the pain that I experience. The latter identifications are particularly interesting because they are immediate *and yet defeasible*: one *might* be mistaken about the one being the other, as when a ventriloquist makes it look like someone else (or a puppet) is speaking. Also, the pain could have a different origin. Such experiences, where the brain performs quite a bit of interpretation on the 'raw' data, nevertheless strike us as perceptions: we hear *that* this person speaks; we feel *that* the scratch hurts.

Identity can be in this way *immediate*, but in many obvious cases it is not: think of seeing Sarah return into the room or realizing that Sarah is the same person as the one who once was a classmate at primary school. The ability to appreciate that different experiences may be encounters with the same object appears to be pretty much what stands between William James's 'blooming, buzzing confusion' (the putative state of the newborn, who have no clue how to interpret their sense experience) and the most basic experience of the world as an independent state of affairs⁶. It is by this that objects are liberated from their confinement in one moment and one mode of sensory presentation. It opens the possibility that the object is perceived in different ways, on different moments, by different perceivers, or not at all.

This *mediate* manifestation of identity is slightly enigmatic. Unlike perceptual identity, which is immediate and therefore easier to accept as primitive, it appears to be in need of some sort of rule: a condition that secures its obtaining. But identity looks like a very fundamental notion, so fundamental that it is hard to see how it could be reducible to something else. There is, nevertheless, an account of identity that regards it as derived, i.e. the *identity of indiscernibles* (the converse of Leibniz's Law, sometimes called by that name as well) which has it that, for arbitrary x and y :

$$\text{If } \forall P : P(x) \leftrightarrow P(y) \text{ then } x = y.$$

In a trivial reading of this principle haecceistic properties also partake in the quantification, so P could stand for *being* x . If objects share this kind of properties, then they are obviously one and the same. In the light of the sceptical remarks that have been made above about fundamentally non-qualitative properties, however, let us assume that P is always qualitative.

Notice that, to make sense at all, the identity of indiscernibles must get its main boost from immediate identity. If not for the latter, the principle would not achieve too much, since, as there is typically *more* than one property involved, the condition on the left-hand side will be a biconditional between two conjunctions, one of which is:

$$P_1(x) \wedge P_2(x) \wedge \dots$$

Now, it is the very appearance of x more than once in the formula which reveals a prior operation of identity: it is *the same object* that has to be P_1 , and also P_2 , and so on. Thus, to evaluate the condition for identity according to the identity of indiscernibles, one has to have prior access to identity, which would seem to make the definition circular. Immediate identity provides such access. Sensory experience presents us with complex combinations of properties and relations, with the bundling already in place. This might be exactly what is needed to get

⁶It is not suggested that babies really experience the world in this way, since it is by no means certain that they lack any notion of identity.


the principle to work in a non-trivial way. Even though there are quite a few issues to be resolved yet, the principle of the identity of indiscernibles is therefore a good candidate for grounding mediate—conceptual—identity.

I shall take up this point again in Chapter 6. Meanwhile I would like to note that there are regions of the draft ontology where the principle appears to be accepted straightaway. In certain domains, in particular those of abstract objects like numbers and sets, formulating a description is considered enough to secure the existence, not just of a referent, but of a *unique* referent. In these domains, if x falls under the same description as y , then the assumption is that they are the same object: a clear application of the identity of indiscernibles. In the formal sciences such descriptions are created by stipulating *identity criteria*. Thus, $\{0, 1\}$ and $\{1, 0\}$ denote the same set, but $\langle 0, 1 \rangle$ and $\langle 1, 0 \rangle$ are different ordered pairs. Now let the existence of some such object—let us say the set $\{0, 1\}$ —be granted, but how can we be sure that there is *just one* of them?

To scale down the level of complexity, let us take the case of numbers for an example. Most of us will in some way or other believe in numbers, at least in the natural numbers. Consider the following: could the world not be such that there is a larger-than-one plurality of exactly identical copies of every natural number? For providing truth-makers⁷ for mathematical theorems this state of affairs would work equally well, even though the interpretation would become somewhat more involved. An expression like:

$$2 + 3 = 5$$

could then be interpreted in a way paralleling that of a chemical equation:



If there are no problems with *many* interacting molecules functioning as truth-makers for this description of methane combustion, so will the truth of the addition be adequately supported by many copies of 2, 3, and 5. All that is required is that the right relations obtain between the members of any suitable triple of some 2, some 3, and some 5.

Some would reject this scenario appealing to Occam's Razor, but that would be too quick. Theoretical parsimony is never *only* about how many entities there are. It is about conceptual simplicity, and if numbers would come in groups of duplicates, that would only make them more like concrete objects, removing a rather 'brute' discrepancy between the behaviour of concrete and abstract objects. Apart from this—and by way of reassurance—duplicate numbers are not a doctrine meant to adhere to. They are a thought experiment intended to show that there is nothing absurd about multiple reference; and therefore, that there is nothing too obvious about uniqueness of reference, irrespective of the domain in

⁷The mentioning of 'truth-makers' serves to sustain my point in the argument; it does not imply commitment to the general usefulness of truth-makers from my part.

question. Any good argument to sustain this idea would probably also support the identity of indiscernibles in the general case.

A final observation to make about identity is that it often behaves in ways that make it seem *predication relative*. The draft ontology strongly supports intuitions like:

41. What you see there is the same river as you saw yesterday, but not the same water.

Some theorists (e.g. Geach, 1967) have taken intuitions like these seriously, thereby endorsing *relative identity*. According to this doctrine it is possible for objects x and y that x is the same P as y , but not the same Q . So we get identity statements like:

42. $x =_P y$

(x is the same P as y), where P is some *sortal predicate*. Wiggins (2001) deduces the falsehood of relativity of identity by extending ‘the formal properties that logicians associate with “=”’ to the relation $x =_P y$. His licence to do so is obtained from the idea that $x =_P y$ means $x = y$ as restricted to the domain of P ’s. This allows him to conclude to the existence of a relativized counterpart to Leibniz’s Law:

43. $\forall x, y, \phi : x =_P y \rightarrow (\phi(x) \leftrightarrow \phi(y))$

after which the deduction runs smoothly.

This, however, seems to be too easy a way to rid oneself of relative identity. If $x =_P y$ entails $x = y$ within the domain of P ’s, that simply means that every case of relative identity under some sortal P can be de-relativized. After all, $x =_P y$ can only be the case if both $P(x)$ and $P(y)$, that is, if x and y are both in the domain of P ’s! Such a reading of relative identity, however, is hardly to be accepted as doing justice to intuitions like the one expressed in (41). It is of no help to answer this riddle by saying: ‘Look, within the domain of rivers it is true that what you see today is what you saw yesterday. So, yes, both *are* the same. In particular, they are the same water!’ By Wiggins’s definition the gist of relative identity as an explanation of cases like this one has been given away at the outset. But if this notion is to perform anything of interest, it must render the *non-relativized* expression $x = y$ at least problematic, and probably meaningless, just like every non-relativized relative concept is (cf. ‘Peter is an inhabitant.’). This appears to have been Geach’s (1980) opinion:

I maintain it makes no sense to judge whether things are ‘the same’, or a thing remains ‘the same’ unless we add or understand some general term—“the same F”. (p. 63-64)

There is no denying that relative identity is an intuitively attractive way to interpret locutions like (41); yet, the fact that every sortal will bring its own identity relation complicates matters severely. To surrender the extremely simple notion of (absolute) identity and exchange it for an uncertain journey past all its relative counterparts is quite a price to pay. Also in Chapter 6 a proposal with respect to this issue will be made, based on identity *simpliciter*, with the intention, at the same time, to respect the intuition that (41) illustrates.

2.5 Plurality

In the quest for the true nature of objects, one inevitably (and very soon) comes across the phenomenon of their tendency to organize into *pluralities*. Pluralities of any type of object readily turn up in the draft ontology—indeed, our routine appreciation of this concept could make us miss its great significance and non-triviality. One might want to make a case for counting plurality among the principles of the understanding—the main reason for not doing so is that an ontology without plurality is neither impossible nor unsound. Also the concept owes much to the principle of *identity*, viz. the possibility of its failure (if objects are non-identical, then they together form a plurality). Getting some hold on its specifics is probably best done by trying to formulate the main intuitive rules guiding it. The first that I would suggest describes the phenomenon itself:

Plurality For every property P , there is also the property of *being a plurality of P 's*.

For this property of being a plurality of P 's I shall write $P^\#$. Such a plurality is itself an object. It must be: if $P^\#$ is instantiated, then what does the instantiating has to be an object.

The draft ontology seems to support pluralities, irrespective of the type of object that constitute them: there are groups of people, languages, architectural styles, elementary particles, and logical theories. P need not even be a sortal predicate: a pile of rubbish instantiates the property of being a plurality of useless things (and *being useless* is not sortal). In fact the world teems with objects which, short of being unambiguous plurals, wear this characteristic on their sleeves (arrays, woods, collections, archipelagos, vocabularies), and for many of the remaining ones it takes just a moment's reflection to reveal it, if only because most objects are made of parts. Pluralities are self-contained objects on all accounts: the more informal term 'group' makes this immediately clear (I shall use the terms 'plurality' and 'group' interchangeably, but the latter is sometimes unfortunate due to its everyday connotations).

As a consequence the term 'plurality' will be used so as to be neutral with respect to the *type* of the unification that the resulting object will represent, be it a *mereological fusion*, a *set*, *class*, or other kind of manifold. There do not seem

to be highly reliable rules as to which gathering of objects is acceptable, and which of the options will result in which case. For mereological fusions, whether or not there are restrictions to the joining of parts is a much-disputed issue. The draft ontology seems inconclusive at this point. As for intuitions, by Armstrong's (1993) lights the Sydney Opera House and the square root of -1 have their fusion, even though the gut feeling of many of us is uncomfortable with such options. For sets the rules are far more liberal: their power of joining is as unrestricted as can be while avoiding contradiction.

The intuitions surrounding plurality are such that we tend to believe we grasp them effortlessly—which is quite misguided! A couple of unreflexive convictions about plurality appears to be grounded in the—overly simplistic—intuitive theory I call *objectual autonomy*, consisting of the following two rules:

Autonomy Let x be an object, and let y be an object. Then there are two possibilities:

- x and y are identical
- x and y are different. Then it might have been that x existed, but not y .

Extensionality Let p and q be pluralities of P 's. Both are the same iff every member of p is a member of q , and *vice versa*.

Objectual autonomy renders pluralities of objects very unproblematic, since they are guaranteed to have a fixed cardinality. Here I shall sketch the (intuitive!) proof for this fact, i.e.:

If objectual autonomy is true, then, for any plurality of P 's, counting the P 's will always, irrespective of the order in which it is done, yield the same result, provided that it yields one at all.

The provision that a result is obtained means that we may assume that the process of taking away the P 's one by one terminates at some point (this makes the plurality *finite*). The notion of *counting* is taken for granted insofar that, if the taking away is done in a *fixed order*, the result will always be the same. Thus, what has to be proved is that the order of taking away makes no difference.

Proof. By induction.

Base case. Let p be a plurality of P 's. Let x be any member of p , and let p' be p without x . Going from p to p' can only be done by removing x , since, by autonomy, taking away any other P instead will leave x unaffected, in which case, by extensionality, the remainder will not be p' . Thus, there is only one order of removal. The count is 1.

Inductive step. Assume that the statement is true for any removal of P 's from a plurality of P 's whose count is k at most. Let p be a plurality of P 's, and let p' be the remainder after removing a certain group q of P 's from p , such that

the count is k . Now remove one more P , call it x , leaving the remainder p'' . Going from p to p'' requires *each* of the P 's in q to be removed and also x , but no others. For, if one of these would *not* be removed, then, by autonomy, it would still belong to the remainder, which, by extensionality, would not be p'' . If some *other* P would be removed, it must be one belonging to p'' , and, by extensionality, the remainder would not be p'' either. Note that removing the P 's in q —or any sub-plurality of them—can be done in arbitrary order, since their count must be k at most⁸. There are three alternative procedures:

- Remove all the P 's in q and then x : this will make the count $k + 1$.
- Remove x first and then the P 's in q . This will, however, make the count of all of the P 's in q 1 higher, so that the last becomes $k + 1$.
- Remove some of the P 's in q , then x , then the remaining ones. This will make the count of all P 's in the second series 1 higher, so that the last becomes $k + 1$.

Hence, the last count will always be $k + 1$. As the process terminates, we can continue until *one* P is left as the remainder p' . Add 1 to the count thus obtained and we have the cardinality of p .

So far so good; however, objectual autonomy, even though it is an indispensable step in the dialectic about plurality, is actually *false* as a general rule. There are many interesting ways in which it is false, but let us consider some examples. Objectual autonomy holds for many 'archetypical' pluralities: eggs in a basket, colours in a painting, people in a cinema. For a case where it is false, consider a geometric pattern of squares (Figure 2.1). The plurality consists of five squares (a through e), but had b not been there (right scene) e would not have been there either. Other examples where objectual autonomy fails can be found with stuff-like pluralities. Bits of sugar are not autonomous: if one bit of some larger amount would not have been there, numerous other bits would not have existed either. What this means is that pluralities do not always behave according to our arithmetical expectations.

All these observations are of course very inexact. However, none of the two major systematic approaches to plurality, set theory and mereology, are too convincing with respect to modelling cases like that of the squares. In set theory the five squares would just make up a set $\{a, b, c, d, e\}$, without any hint of the dependence of the fifth square on all the others. This is not surprising, since objectual autonomy holds for all sets—sets are pretty much *made* to secure the objectual autonomy of their members. In mereology the dependence relations among the squares are perfectly represented, but if parthood is to play the role of membership, then the plurality has a bewildering multitude of 'unintended'

⁸Here I omit the proof that the count of a sub-plurality cannot be greater than that of the plurality to which it belongs.




Figure 2.1: A plurality of squares

members. Where mereology is too generous in granting membership of the plurality of squares, because it is willing to ‘cut the pie’ in every conceivable way, set theory behaves too artificially in placing the cuts. After all, there are lots of sets which are fully acceptable descriptions of the situation depicted in Figure 2.1, e.g.:

- $\{a, b, c, d\}$
- $\{a, b, c, d, e\}$
- $\{a, b, c, d, \{a, b, c, d\}\}$
- ...

but whoever wishes to find out which one is to be preferred is left clueless.

Would it be possible to craft a systematic representation of pluralities more faithful to the behaviour of pluralities in the draft ontology? What seems to be needed is a ‘natural’ way of determining membership, but it is easy to show that, more often than not, membership is inherently ambiguous. Let us say that a *dancing couple* is just any combination of a man and a woman. Then a plurality of dancing couples (to be found in every ballroom when in use) is any group of people of at least one woman and one man in it. Every new person who joins the group amounts to the addition of as many new couples as there are members of the other sex. Here every combination of a man and a woman constitutes a fairly natural member of the plurality, but only if the group is considered a plurality of dancing couples—of course any such group cannot fail to be a group of *persons* as well, which comes with a different membership relation. Relativizing membership might seem to be a solution, but hardly an elegant one. There are even more tricky problems with membership. Compare this case with a plurality of pairs of gloves, all of the same size, colour, and so on. Now, if we add one *single* glove, do we still have a plurality of *pairs* of gloves? Here we are unwilling to

allow the gloves a *change*: we want to see a full decomposition into complete pairs of gloves. Hence, ten gloves is five pairs at most, not 25! Such problems, because of their pragmatic subtleties, might well elude the toolbox of formal logic. This is no reason to lose interest in pluralities as objects, as will become clear in the coming chapters; but for now we will have to make do with the conclusion that there are objects which are pluralities, and that they have other objects as members.

I want to end this chapter with a short note on the phenomenon of *totality*. Many of the most interesting pluralities are totalities, and totality is part of what is involved in *universal quantification*, an aspect of first-order logic that has not yet received due attention. The intuitive rule can be stated as follows:

Totality For every property P , if it has an instantiation, it also has a unique *totality*, i.e. the plurality to which *all* P 's belong.

For the property of being a totality of P 's we shall write P^\forall . One of the salient aspects of this condition is that, for every P , there can only be one unique x such that $P^\forall(x)$. Because of this, we can write $\iota x P^\forall(x)$ to denote *the* totality of all P 's.

What will strike most logically trained reader as problematic about the above rule is that it invites Russell's paradox. This is true, but remember it is a 'small rule' we have here. 'Totality' appears to be a clean example of the pragmatic nature of intuitive rules—they will not stop functioning for such technical reasons. A great lot of properties after all *do* define *bona fide* totalities—one could easily spend a lifetime working with sets without ever running into a problematic case. But should troubles turn up anyway—should someone pass by with the description 'the property of not belonging to oneself', then intuitive reasoning simply abstains from considering this specific totality.

Universal quantification, or so my proposal goes, is cognitively associated with two *different* principles: generality and totality. Many languages have ways to distinguish between both (also known as the *distributive* vs *collective* reading). If the area in square miles is the relevant factor, then the sentence: 'Every American could live in Maryland,' is true beyond any doubt; this is less certain, however, for the most obvious reading of 'All Americans could live in Maryland,' which is that the entire population of the US could be housed in that state. With the help of the notation introduced above we could paraphrase these sentences in the following way:

44. Every American could live in Maryland. $\forall x : A(x) \rightarrow M(x)$

45. All Americans could live in Maryland. $M(\iota x A^\forall(x))$

General statements are those that can be made by considering every item x in a certain category *without* using information about the category as a whole (e.g. its size, shape, succession, dispersion, and so on). Or rather, the conclusion

is reached by *abstracting away* from the peculiarities of the individual instances, using only those properties that make these instances qualify as members of the category. Hence, it is reached by inference: the underlying principle is that of *consequence*. If the statement:

$$P(x) \rightarrow Q(x)$$

can be shown to hold, irrespective of what x refers to, then:

$$\forall x : P(x) \rightarrow Q(x)$$

has been proved true. It is therefore that most general statements one comes across have the form of a universally qualified inference.

Totality statements do not have a distinctive quantificational form. They are more powerful than generality statements: there is no way to express a statement like (45), which contains collective predicates, as a generality; on the other hand, every general statement can be recast as a statement about a totality, even though it takes a new (collective) predicate to do so. The decisive advantage of general statements appears to be that statements of this kind are the ones we can be certain about. Totalities are of theoretical interest, but they tend to be beyond our grasp. How can we know whether all swans are white, if whiteness does not *conceptually* follow from swanness? But every white swan is a white bird.

Chapter 3

Properties, Procedures, and Images

In the Introduction the prospect was raised of an analysis of properties that allows them to be non-immanent to the objects they apply to. To make good on this promise, the present chapter starts with a discussion of Bradley’s regress, a notorious issue in ontology. After examining some of the solutions that have been suggested, the (realistic) assumption of *objectual compositionality* is identified as the source of confusion, whereby a new and more specific argument for transcendentalism becomes available. In subsequent sections the first elements of the transcendentalist approach are worked out in the form of what will be called *image theory*, a theory centred around the modules of the human cognitive system (*images*) that perform *procedures* on what is perceptually available to us. It is these that determine properties, and thereby effect the emergence of the objects themselves.

3.1 A Conflict Within the Draft Ontology

In the first lines of the quotation on page 21 Quine makes the important point that saying, ‘Some dogs are white,’ does not imply commitment to the existence of whiteness. Saying this is saying that some things which are dogs, are white: the bound variable ranges over dogs only. To arrive at this conclusion it might *seem* to be syntactic form, not content, which has most clout. Yet Quine—correctly, in my view—sees an *ontological* distinction. Compare:

46. Some dogs are white. $\exists x : D(x) \wedge W(x)$

where the only objects figuring are some dogs, with:

47. There are some dogs whose colour is white. $\exists x : D(x) \wedge I(x, W)$

where the colour has become an entity in its own right (*I* stands for the relation of *instantiation*). In (46) there is commitment to dogs only; in (47) there is also commitment to the colour white: it has become an object alongside dogs.

Properties, as we have seen in 2.1 as well, can become objects themselves. This is a very general phenomenon in the draft ontology. The intuitive rule can be formulated as follows:

Abstraction Whenever predication takes place, the predicated property can also be treated as a new object, bearing the relation of instantiation to the one that was the target of predication.

That there is a pervasive tendency in our thought and speech to treat properties as objects is one of the most well-documented facts in modern ontology. It is, to the best of my knowledge, a universal phenomenon, found in all cultures and all languages, and should it be anything else than a veridical appreciation of the way things are, there must be deep reasons why human language and cognition behave in this way, for it is unthinkable that it is an accident. If belief in (as the term of art sounds) *universals* like whiteness, virtue, and electric charge is an ontological error, it is an unusually tenacious one. It is no trouble at all to learn to avoid reference to unicorns or phlogiston, and with only a little effort we make ourselves stop speaking about the movements of the sun around the earth; but a discourse free of properties as self-contained existents hardly seems to be within our power. Reflecting on this practice and analysing it comes down to querying, in Wittgensteinian terminology, the *bedrock* or our cognitive capacities.

In Chapter 1 I have advocated a general tolerance towards objects, including those to which we are not committed to the strongest degree. Things like whiteness cannot be just dismissed. You, a nominalist, say whiteness does not exist? Of course it does: look at that wall! This is not a particularly hard contention to gain sympathy for; yet, reflecting on these commitments and reasoning on the basis of them may place us before quite puzzling conclusions.

Here is one—it is a well-known issue in metaphysics. By the intuitive rule ‘Abstraction’, whenever an property is attributed to an object, the property itself can be turned into another object. Clearly, the latter is related to the former, and by re-running Abstraction we can turn this into an object too: the relation of instantiation. This third object then relates to both that came earlier, and so there has to be a fourth object—and so on and on. This scenario is known as *Bradley’s regress* (Bradley, 2016); it is an inescapable consequence of Abstraction, and very few metaphysicians feel comfortable with it. The draft ontology itself does not bother about the seeming absurdity of Bradley’s regress, just like it does not bother about Santa Claus who does not exist or the growing absence of fresh water; but those who reflect on the draft ontology have enough reason to do so. Bradley’s regress at the very least *looks* problematic: why should the simple fact of a dog’s being white give rise to a myriad of new, ever more complex objects?

Since we cannot neutralize the regress, all we can try to do is show that it is in fact innocent. Let us see what options have been put forward. First, there is the vintage idea of viewing properties as *universals*. The common way of describing the relation between a universal and the objects to which it applies is to say that the former *is wholly present in all its instantiations* (Armstrong, 1989). Hence, the strategy for escaping Bradley's regress is to frame the relation as some sort of *composition*, which type of relation, just like that of identity, is thought to be so 'thin', so devoid of real marrow, that a Bradley's regress only consisting of it and its derivatives is not vicious. All the additions to that one fact of a dog's being white do not add up to anything to worry about. David Lewis's (1991) term for this condition is *ontological innocence*.

The essence of this proposal is the notion of composition. A composition is not necessarily the same thing as a *mereological* composition, for this term is reserved for the type of composition of which, say, the bricks constituting the wall, or the cells constituting the body, provide archetypical examples. Mereology, the formal theory about parts and wholes, was devised to describe this type of composition, but there might well be other types. Yet, it seems safe to say that every compositional relation is about *components*: objects that, in some way or other, come together to give rise to a *new* object, which is the composition. Now, to secure its ontological innocence, it is important that composition be a very 'passive' relation. Nothing much is supposed to 'happen' by the coming together of the components: *no addition of being*, only a bringing together of what already was. In terms of commitment: no *new* commitment, nothing but things whose commitment was already established. So the coming together of the dog and its whiteness would not require any other 'component' (instantiation) to take place.

But all the same, how are we to conceive of whiteness as a *component* of a white dog? At least this would imply that any two white dogs share one of their components, which seems wildly counterintuitive. This intuition might be misguided, though. Universals need not be parts like legs and tails: maybe they do not exist in time and space or have a way of being that allows them to belong to two locally separated dogs at the same time. This is the idea behind *bundle theories* (Van Cleve, 1985; O'Leary-Hawthorne, 1995). If we divide the animal into parts in the more usual way, what we get is, below the level of legs and tails, organs, tissues, cells, molecules, atoms, elementary particles, and nowhere will we encounter whiteness; but whiteness could be the result of a different way of taking the animal apart. Should this be so, however, then still, when following the first way of breaking up, the whiteness should go somewhere—after all, the dog is supposed to be a composition: a going together of whiteness plus lots of other things. Now the legs and tail are all white, so whiteness is wholly in all of them (as it could be in different dogs). But down the path to elementary particles it sort of dissolves, since it is certainly not to be found in any of these.

Yet, it could be that the elementary particles of which the dog consists contain lesser components of whiteness. Universals might well be composite, and

there is some reason to believe that their components are themselves universals¹. Although it seems weird to think of whiteness as a compound entity, it is possible that intuition leads us astray again and whiteness is really a complex of smaller universals. To examine this possibility, it might be wise to switch to a property of our dog that is structured in a more obvious way. Take its *quadrupedality* (the condition of having four legs). This is a universal that dissolves well at the visible level, so its components must have gone into the macroscopic parts of the dog, in particular its legs. And here we run into a real problem. By the logic observed for universals so far, the universals to be found in the dog's legs must be the properties of those legs or of their parts, and these are the building blocks available to make up the universal quadrupedality. The candidate of choice seems to be *leghood*: quadrupedality just is what it is thanks to its containing *this* universal four times. Should the creature lose one leg, then its quadrupedality is replaced by tripedality, which, analogously, is that thanks to its containing *three* times leghood.

Remember, however, that universals were supposed to be wholly present in all of their instantiations. This means that the leghood that is in all four (or three) legs, is *one and the same* universal! Just like one and the same whiteness is in all white dogs, one and the same leghood is in all legs. But then there is no difference between quadrupedality and tripedality—contradiction. This is a sketchy presentation of one version of the *repetition problem*, famously discussed by Lewis (1986b) in relation to the structural universal *methane*. Even though this difficulty need not be the final word on the fate of universals, it should be clear that the simplicity and elegance that this type of solution displays at first sight, soon grows pale when looked at more closely.

To counter these (and other) problems with universals, there is a competing family of theories about the objecthood of properties, based on the idea of *tropes* (Williams, 1953; Campbell, 1981; Schaffer, 2001). Tropes also represent properties, and, like universals, tropes are also in the things themselves: *in rebus*. Hence, just as with universals, the whiteness trope can be construed as a component of the white dog. There are, however, separate whiteness tropes for all white dogs, in fact, for all white things. Consequently, whiteness is not one single object, but is scattered among the things that are white, and what unites such a class of tropes is that they *exactly resemble* each other. Clearly, trope theories have no troubles with the repetition problem. The draft ontology makes room for tropes as well: it was Snow White's *own* beauty that made the Queen jealous, not just beauty. But now we are faced with a different problem. Remember that it was the relation between the *object* whiteness and the *object* white dog we sought to clarify; but almost without notice, the first of them has gone missing! If we speak

¹The main reason for this is theoretical parsimony. Universals are not material objects, since wholes of material parts are surely not universals; 'mixed' compositions are difficult to conceive, and the introduction of yet another type of entity is unattractive (Lewis, 1986b).

about *the property* whiteness, what is the object that we speak about? Universal theories had a plain answer to this: the universal. What do trope theories have to offer here?

Trope theorists will not necessarily be impressed by this objection, since they can argue that their model sufficiently accounts for the structure that can be discerned in the network of objects and properties. However, our present agenda is a bit more ambitious: we also want to explain the objectual nature that properties so conspicuously display in the draft ontology—or at least *seem* to display. A possible answer is that it is the *group* of similar tropes, taken as one object, that is responsible for this appearance. Accepting the validity of this identification would do a lot to bolster the credibility of trope theories. What this move does not do, however, is obviate the need for a primitive resemblance relation, or else any collection of tropes could flock together to become ‘a property’, which would undermine the notion of similarity altogether. On top of that, there would have to be a primitive relation of *compresence* (the gathering of tropes to constitute one object). Without a relation managing the composition of tropes, it is hard to see how a component of the dog’s tail (tailhood) could fail to be a component of the dog itself, making the dog a tail! Compresence should take care of this, lest everything should have all the properties of its parts. The conclusion must be that, just as with universals, trope theories cannot make do with composition as the only primitive. The putative *parsimony* of trope theories—they need only two inscrutable principles to make the whole machine run—does little to enhance their attractiveness. It succeeds in reducing the *number* of problematic conceptions, not in making them less problematic.

In response to the problems with making composition work for, in particular, universals, several authors have tried to revise the composition relation itself (e.g. Kalhat, 2008; Hawley, 2010; Mormann, 2012; Armstrong’s (1980) notion of ‘constituent’, rather than mereological part, to bind universals to particulars, should also be counted among these approaches). Alternative accounts for composition are logical theories, often showing considerable similarity with classical mereology, but ‘enriched’ with concepts and axioms so as to make the options for resulting compositions fine-grained enough to match the diversity of real objects. If some minor adjustments could enable the ‘innocent’ phenomenon of composition to do all the structuring work that is needed to get the objects and universals into their proper places, the resulting theory would boast an irrefutable claim on meaningful parsimony.

The problem with any such alternative account of composition appears to be that there is no guarantee that ontological innocence is retained. There is some reason to believe that classical mereology is the only theory that does justice to the utterly simple principle of ‘bringing together’ that is the core of composition (I shall discuss mereological composition in Chapter 4). To achieve what is needed, alternative theories of composition add a considerable amount of complexity to this core. This is not to deny the possible value of such systems

beforehand: if they succeed in describing certain aspects of the way the world is structured from a small number of first principles, they must be welcomed just like any successful theory. Yet, whatever their merits, the extra structure they require is what weakens their claim to the ontological innocence of what keeps the constituents together. But something like ontological innocence is mandatory for doing this—*that* was the line of defence against the viciousness of Bradley’s regress!

3.2 A Provisional Diagnosis

My suggestion is that Bradley’s regress, along with many other stubborn issues in ontology, are illustrations of *the failure of objectual compositionality*. In the last section some of the main approaches with respect to the ontological analysis of objects were briefly discussed, and they all exhibit elements of compositionality. I have identified Lewis’s ontological innocence with the condition that the compositions in question are *just* the coming together of things, without recourse to additional principles. This type of compositionality, where composition does the job all on its own, I would like to call *free* compositionality. Objectual compositionality, as intended here, is free compositionality. If compositionality is not free, it is *embedded*.

Let us have a closer look at what this ‘just coming together’ amounts to. In 2.5 I have described *objectual autonomy*, a hypothetical intuitive theory—and an overly simple-minded one—about pluralities of objects. Objectual compositionality, or so I claim, is another intuitive theory; more refined than objectual autonomy, since it provides a solution for the most conspicuous defect of objectual autonomy: the neglect of mereological relations. It is based on two intuitive rules: the aforementioned ‘Plurality’ and ‘Extensionality’:

Plurality For every property P , there is also the property of being a plurality of P ’s. (Call this property $P^\#$.)

Extensionality Let p and q be pluralities of P ’s. Both are the same iff every member of p is a member of q , and *vice versa*.

Clearly, compositions are pluralities, so what we prove about the latter, applies to the former. A prominent consequence of ‘Extensionality’ is that *there is no difference without a difference maker*—a well-known one-liner, often used to defend the principle of uniqueness of composition in mereology (e.g. Varzi, 2008). A difference maker, then, is a member of the composition that sets it apart from other compositions to which it does not belong.

With these two rules, the following conclusions can be drawn. Take two objects, x and y , and let there be some P , such that $P^\#(x)$ and $P^\#(y)$. Then just one of the following options must be applicable:

- *Distinctness* If none of the members of x belong to y , then x and y are *distinct*.
- *Identity* If all of the objects belonging to x belong to y , and no others, x and y are *the same*.
- *Inclusion* If all of the objects belonging to x also belong to y , but not vice versa, then x *includes* y .
- *Overlap* If some of the objects belonging to x also belong to y , but not all of them, and vice versa, then x and y *overlap*.

Hence, according to objectual compositionality, any x and y are either distinct, or identical, or one includes the other, or they overlap. This state of affairs will strike the reader as suspiciously similar to classical mereology. As said above, however, the right appreciation appears to be the more general idea that every object can be analysed as *just being some specific aggregate* of elementary constituents. Same constituents, same object, different ones otherwise.

Applied to objects this yields objectual compositionality. Objectual compositionality is a very attractive idea, and my suspicion is that it is this conception which runs in the background of pretty much all realistic theories about objects. And not without reason: if the aim is to bring about a successful analysis of *whatever* domain of interest, few things could be more desirable than compositionality: the situation whereby every item in that domain can be thought of as being the result of recombining some basic elements. Compositionality can boast great victories: in chemistry, particle physics, genetics, linguistics, even art, it is the very light to lighten the chaos. Things can be taken apart into more fundamental things, where the number of the latter is smaller than that of the former.

Now in practically every *subdomain* of objects compositionality is embedded: it is comfortably couched in a larger context. Chemical substances are combinations of elements; but it is the atomic bonds that do the combining, and they are not among the elements themselves. Particles need spatio-temporal connections to compose larger objects. Genes recombine in chromosomes, nuclei, phenotypes. Linguistic elements need expressions, speakers, and communities; and artistic styles need artists and their products to meet. Compositionality thrives when there are pre-conceived ‘slots’ in which to insert the elements to be combined. But among these pre-conceived structures are precisely the ‘difference makers’ that stay *outside* the composition. They make a difference, but short of being one of the elements! This is what defines embedded compositionality. If there is no need for the domain of interest be all-encompassing, difference makers may stay out of it.

Exactly for this reason, this option is *not* available for objectual compositionality. A working version of objectual compositionality would have to provide,

on the basis of merely combining elements, if not full understanding, then at least *prospects* of understanding each and every object as just a composition of basic ones. Prospects would suffice—omniscience is not required: we may take ourselves pardoned for, say, not knowing whether the world is atomic or not, or what the ultimate constituents look like, or by what sort of relations they can be connected—but to say how the analysis would proceed in various ‘what if’ scenarios should not be asking too much. At least it should be possible to perform the analysis with respect to those things we are very familiar with: tables, white dogs, tails, whiteness.

And such an analysis does not look like being forthcoming. Mereology works wonderfully for concrete objects, but when extended to include, e.g. universals, the picture turns bleak. The problem for theories of a more general scope is invariably with difference making. Bradley’s regress itself is about additional objects (instantiation, instantiation of instantiation, and so on) that *fail* to make a difference; the ontological theories meant to counter the problem typically produce differences for which *no* difference maker can be appointed. Bundle theories need compresence (or bare particulars or any other such principle) to ‘regulate’ the composition of particulars out of universals (Denkel, 1997; Sider, 2006), which principle is a difference maker, but not itself a universal. Trope theories need similarity and compresence to organize the tropes. These are difference makers that prevent the world to look like a pointless mass of free-floating tropes, but neither of them are tropes themselves².

For a more general argument to the same effect: objects can be analysed in terms of parts and also in terms of properties (particularized or not), but there does not seem to be a recipe to join *both* types of composition into a unified (non-embedded) scheme. When we say, after Aristotle, that whiteness is *present in* a white dog—and how could *that* be false?—a mereological understanding of objects generally tempts us to liken whiteness to some part that all white things have in common; which is something common sense prompts us strongly to resist. Armstrong (1997) gives a unified account in terms of *states of affairs*, but at the price of having to postulate more than one notion of part (‘mereological part’ vs ‘constituent’). Such a construction is hard square with free compositionality. It seems impossible to analyse objects as *mere* aggregates of certain more basic constituents. The amount of structure thus allowed is simply insufficient to replicate that found in the draft ontology.

Or so it seems. I cannot lay down final proof that any successful implementation of objectual compositionality is impossible; all I can (and will) do in the text below is contribute my bit to the troubles weighing down on the existing ones. But, one might wonder, is (free) objectual compositionality really required? In a *realistic* ontology I believe it is. If all objects are out there, in reality, independent of human intelligence, then so are all the difference makers. *Where else* could

²This has been contended, however (Maurin, 2010).

they be? Keeping the difference makers ‘behind the screen’ is cheating. In *reality* there are no screens, only objects. Saying that difference makers are not really ‘things’ is not helpful either. *What else* could they be? A realistic ontology is supposed to round up *all* the elements involved in object making; the only excuse for leaving some of them out being that they are beyond our sight or grasp (like the possible lesser constituents of particles, space, or time). But the different makers mentioned above are wholly within our grasp! What remains wanting is a way to make them partake in the composition in a credible way: as just another component, not on a regulatory meta-level.

To my estimation, therefore, giving up on objectual compositionality is giving up on plain realism. As it turns out, the above considerations with respect to Bradley’s regress have given even more specific reason to do so. The transcendentalist alternative, uncommitted as it is to free compositionality, might open new perspectives for understanding purely ontic relations between objects, allowing for more variation than is given by composition only. This, however, requires the formulation of a new theoretical framework, centred around the main explanatory principle of transcendentalism as put forward here: the human cognitive system.

3.3 Procedures

One important hunch behind the transcendentalist understanding of properties argued for in this book was this: if we feel that two white dogs have something in common, but it is not anywhere interpretable as a ‘thing’ just like the dogs themselves, then maybe what they have in common has to do with our cognitive reaction towards them. It is a thought well in line with the general style of reasoning that has been followed so far. Here is a second thought, one that I have not put forward so far. Many of the properties acknowledged in the draft ontology do not seem to be an ‘aspect’ of the object, just waiting to be discovered, but are better understood as being the result of some *procedure* that a cognitive system *performs* on it (Moschovakis et al., 1994; Moschovakis, 2006; Duží et al., 2010). Add one to the other, and the outlines of a transcendentalist approach to properties begin to form.

The second thought may sound like an odd claim. From a realistic point of view properties are just there: the dog is white, whether or not there is anyone around to perceive, state, or conclude that this is so. It is, however, a conviction which is easily shaken. In fact this example provides a good introductory case for the sort of argument that will be developed below. For it seems clear that the dog is also white when no one looks, but when there should be no one left with the common visual capacities of humans (or if no one would ever have had them), it is hard to maintain that things would all be the same, even if nothing would have changed with respect to the dog’s fur. There *is* a sense in which there is no whiteness without (human) eyes. Even though the properties that I want to

discuss presently are arguably less subjective; here also, what is at their core is not so much in the instances themselves, as in the procedures used to determine whether they obtain. They that do not just gather what is out there, but perform something like a computation on it.

One such procedure is *counting*. Rather than the legs of a dog, we shall this time be concerned with the trunk segments of a centipede. The property we are interested in is the *being even* of this number. Although there are many species of centipedes, and many varieties in numbers of segments, none of them has this property: all of them are—let us call it—*odd-segmented*. Let us think of a biologist who tries to find *even-segmented* centipedes. A procedure that determines this feature could work in the following way. What is needed is:

- the ability to recognize a segment of the centipede,
- the ability to recognize a starting point (the one next to head of the centipede),
- a way to label any segment that has already been counted,
- the ability to know that the last segment has been reached.

The first condition is dependent on whatever it takes to determine an *underlying* property: being a segment, in particular, a segment of *this* centipede. This could again be a procedure, but in other cases it could also be a basic perceptual ability. After one of the centipede's segments has been counted, it could be labelled with ink, which would change it in a way that is recognizable (this should be included in the third condition), but the more usual method is that whereby the counter changes its own position towards the segment that has just been counted, as is done when starting at the head of the creature and working one's way towards the rear end. Finally, having arrived there, the counter should not go on searching for new segments forever, but terminate the process. A TRUE/FALSE verdict on even-segmentedness should then be returned.

Here is the flowchart of what could be the procedure of even-segmentedness, in a very condensed language, to be called IMG³:

start	step	observe	yes	no	
1	1	end	FALSE	+1	2
	2	end	TRUE	+1	1

In IMG-code every step has five fields: first its starting-point (**start**), then its unique identifier (**step**) and the observation (**observe**) by which its action is determined. The fields **yes** and **no** contain the actions taken if the observation has a positive or negative outcome, respectively. These columns can be split to

³This is an abbreviation of 'image'; see next section.

provide the *external* (left) and the *internal* action (right). External actions are, e.g. moving up one segment of the centipede (+1), but also outputting TRUE or FALSE. In the latter case the column is not split, since the procedure terminates here. Internal actions consist in moving to another step in the procedure. Any impossible action will result in the output FALSE.

The procedure always starts with the first step. Thus, in the example, determining even-segmentedness is done by starting at the segment immediately next to the head (**start** = 1). Should this be an end-segment, then the outcome FALSE is produced (the poor creature is one-segmented), otherwise the procedure moves up one segment and goes to step 2. If the new segment is a terminus, then the outcome will be TRUE, otherwise at step 2 the procedure moves up one segment again and goes back to step 1. As long as the segments keep coming, the algorithm keeps jumping between step 1 (odd number so far) and 2 (even number so far), until it produces an outcome at the final segment.

Finding an even-segmented centipede would cause something of a sensation in the arthropodologist community. In all likelihood this property would correlate with other properties of the organism. It might reveal interesting genetic and embryological facts—possibly even with medical applications! This bit of rhetoric from my part serves to make the point that even-segmentedness is a *genuine property*, no less real than being red or being an insect, and so a universal or trope theory should be able to account for it. For universal theories this is only a more daunting version of the repetition problem, but for trope theories it brings trouble that was not there before. What does an even-segmentedness trope look like? Since the property supervenes on properties like 28-segmentedness or 34-segmentedness, tropes for the latter properties should somehow be sufficient for the even-segmentedness trope to be present. But 28-segmentedness tropes cannot be *similar* to the 34-segmentedness ones, since this would erase the possibility of 28-segmentedness to be a property in its own right. Remember that trope theories have only one relation of similarity—dropping that condition would leave the theory bereft of most of its attractive economy.

It is not hard to see that identifying properties with procedures as carried out by the cognitive system is an appealing strategy to overcome the troubles with realistic theories. However, to turn this approach into a real contender to these, more is needed; after all, the realistic theories do have their explanatory virtues, and a transcendentalist alternative should have something to say about why this is so. In the interest of this, rather than to be built up around the concept of procedure, the theory will have as its central concept that which, as part of the cognitive system, allows the procedure to be performed. This I shall call its *image*.

3.4 Images

The approach that follows from giving procedures centre stage in the understanding of properties is in many respects similar to that of Moschovakis (2006), whose theory is based on *algorithms*. The main difference between images as understood here and Moschovakis's algorithms is that the latter are *abstract* objects, an assumption that would compel me to accept such objects as given. One of the aims of the present study, however, is to explain what abstract objects are and why, alongside concrete ones, they are among the inhabitants of the draft ontology. Hence, the search for a new home for properties will stay on the concrete level. On this level algorithms are embodied by two things which are not clearly distinguished by the term. First, we have the *procedures* as concrete interactions between scanning devices and objects; and second, the *devices* themselves, running them. Only the latter are material objects; and even though there is nothing evidently primitive about material objects, it is the material constitution of the cognitive system which will be at present be the anchor for explanations with respect to the way the world looks to us.

The flowchart in IMG-code on page 62 is an example of what could be a specific type of image. It is a piece of code that a suitable computing device could use to scan for even-segmentedness. Note that it is not itself a procedure, but encodes one. Even though procedures can go on for unpredictable lengths of time, their images are always of definite, finite size. The functional modules of our brain will also be ranked among the images. Of course neurons, axons, and dendrites look very different from a table of IMG-symbols or even the machine interpreting it. Yet, the idea is that the brain contains subsystems which are *functionally* similar to IMG-flowcharts or like machinery. The typical function of images is determining, by certain criteria, if some condition obtains: to this end the image for whiteness generates a positive outcome when brought into contact (in some suitable way) with a white object.

Since images are functional units of a cognitive system with a specific ability of processing information, in humans they must be thought of as being involved in thinking and perceiving. Thanks to this, they organize the flow of information: they are the nodes in the computational network, and those nodes—this is the key claim I want to submit in this work—are ultimately responsible for the way in which we experience objects and their properties. In universal and trope theories properties are conceived as being *in rebus*: they are where the objects are, and on a par with them. Above I used the term 'immanence' for this state of affairs. In contrast, in what henceforth will be called *image theory*, properties are *ante res*: they are there 'before' any concrete object is even in sight.

That this is in accordance with the transcendentalist conception of objects endorsed here, may require some clarification. In the foregoing I have insisted that objects as they appear in the draft ontology should be taken seriously, even though some of them deserve a stronger commitment from our part than others.

The treatment of whiteness as it is presented here *could* be read as a betrayal of this position. It may look as though white dogs and other white things are just real, as opposed to whiteness, which is a mere mind-dependent addition. Clearly, this would be another version of two-layer supervenience, an idea that was rejected in 1.3. But this is not what is intended. Images, or so I shall argue, are needed to account for *any* type of object. To reiterate the point made before: a nominalist who has rid herself of whiteness will, sailing on the same compass, soon find herself troubled by cracks in the ontological solidity of the white dogs themselves. Removing the abstract property alone is, by all available accounts, insufficient to clean up the ontology so that only ‘real things’ remain. There are issues about the composition of the dogs in question (is a dog still the same dog once its tail has been lost?) and about their enduring existence in time, as opposed to that of a succession of temporal stages; to name some (Johnston and Forbes, 1987; Varzi, 2000). It may look as though these problems are unrelated to the point about whiteness, but by the present proposal they all share a common root. All of these matters are related to the fact that no objects are entirely mind-*independent*, whereas, on the other hand, we cannot without contradiction deny reality to what we find ourselves committed to.

But how, and why, do images ‘give rise to’ objects in general? Images were introduced as having the function of distinguishing between objects by scanning them for certain properties—this does not appear to include the emergence of *new* objects. A (tentative) answer to this question must wait until the end of the next chapter. The coming three chapters will be dedicated to different aspects of the functioning of images, but the general thing to be said now is that the effect of images organizing our view of the world is that we have two ways of experiencing properties. As *predications* (this dog is white) properties are ways of being: when seeing (or thinking about) certain objects, they strike us as being that way. As *objects* (whiteness) they arise as a novel way of carving up reality.

Now where, according to this picture, has Bradley’s regress gone? As image theory is not constrained by the demands of objectual compositionality, the inexorability of Bradley’s regress is no longer there. Dogs can be white without the need for any extra object associated with them. Whiteness is *rooted* in white dogs, rather than that white dogs are composed of (among other things) whiteness. In image theory, the property as an autonomous entity thus originates from a *shifting perspective*: that there is whiteness where white things are. The proliferation of new items takes place *in thought*, not in the world, and therefore need not go on forever. Notice that this view on the matter fits tolerably well with the otherwise slightly esoteric expression *no addition of being*. Whiteness is an object different from all white things, but it does not bring about a new reality. Yet, the freshly arising perspective justifies the impression that something new has come about anyway; which provides an elegant vindication of our intuitive feeling that there is an addition after all, albeit a puzzlingly insubstantial one. It takes the transcendentalist position to benefit from this idea, but perspectives,

leaving what is within their reach essentially untouched, have a rare claim on genuine ontological innocence.

Image theory provides some distinctive explanatory advantages. None of the nasty little objections against universals and tropes, some of which have been touched on above, apply to images, whereas at the same time there is no well-specifiable property of any concrete object that *cannot* be understood as the outcome of a perceptual-computational test. To the extent that we know the conditions for a property to apply, we can think of it as implemented in the guise of a procedure executable, and therefore definable, by means of an image—whose details are of course the business of those working in cognitive science and artificial intelligence. As said before, to be a worthy alternative to the realistic theories about the nature of objects, image theory will have to do more; in particular, it will have to reproduce the explanatory successes of the latter. Therefore, in the process of working out the theory, there will be comparisons with these from time to time. As for now it seems clear that image theory at least provides sound prospects for incorporating (embedded) compositionality: after all, intelligent systems, composed out of smaller modules, with several layers of nesting, have been with us for some time yet.

3.5 Grades of Sparseness

It has long been appreciated that properties differ in the extent to which they have a claim on being ‘natural’. The classic example of an ‘unnatural’ property is *being grue*, as devised by Goodman (1965). Being grue is a property which can be had by e.g. emeralds: an emerald is grue if either it has first been examined before 1 January 2000 and is green *or* has been examined later (if at all) and is blue. Properties like these strike us as artificial, contrived, gerrymandered. They lack the nobility of a property like *being blue* or, better still, *being gold*. In fact it can get importantly worse: consider all the objects there are, and let their number be α . Then there are 2^α sets of objects and there is a sense in which all of them are the extension of a property; it is this sense for which Lewis (1986a) coined the term *abundant properties*. In this terminology natural properties are *sparse properties*—whether both designations refer to the same class of properties is a point of dispute, but the intuitions behind them seem to be largely the same.

I prefer to use Lewis’s terminology for making a distinction between being natural in a broad sense—uncontrived, elegant—as opposed to being fundamental from a scientific point of view. Properties from the former class I shall call ‘sparse’, saving ‘natural’ for the scientifically basic ones. In my view the terms sparse and abundant mark both ends of a scale, the distribution on which is strongly asymmetric. In fact *no* properties that we humans have the faintest conception of are anywhere near abundant: the properties that we know all have their *grade* of sparseness. Accordingly, when I speak of a non-sparse property, I just mean

that it is sparse to an unusually *low* degree.

It is easy to recognise a non-sparse property when seeing one, but which, then, are the sparse ones? How are they told apart? There is a fair amount of consensus that the properties which are most basic in science, in particular in fundamental physics, are good candidates, but there are more of them. Biological species, even though there is little hope that they can be reduced to a set of conditions in the vocabulary of fundamental physics, seem to have some right to be called natural kinds as well (Schaffer, 2004). Inevitably we will have to face the deeper question as to what specific criteria can be found for the sparseness of properties.

Universal and trope theories are often thought to be in a favourable position to say something about this. The idea is that there could be a *blue* and a *green* universal or tropes for these properties, whereas making up a condition like ‘grue’ will not bring about a new universal or a new kind of trope. However, by taking properties like blue and green as the examples, it could be that we make the problem look easier than it is. As was discussed above, the main problems with universals and tropes have to do with the difficulty of reconciling these entities with kinds that manifest themselves in similarity of structure and composition, rather than ‘brute nature’. *Being gold* and *being water* are examples of such kinds: gold atoms and water molecules owe their nature to the way their quarks and electrons are *arranged*, not to consisting of different *kinds* of these particles. However, properties of this sort surely do not suffer from lack of sparseness: *being gold* or *being water* are archetypical natural kinds.

How does image theory fare on this score? How should it handle the grue case? First, observe that the most extreme cases of non-sparse properties—the abundant ones in Lewis’s sense—are forever inaccessible to images of remotely feasible size. Presuming that it is possible at all to distinguish every (material) object from every other (a condition that is by no means guaranteed); to encode a procedure that picks out just one of all those subsets of all objects (not in a random way but reliably), the size of the image would have to be gigantic. This is because this size has a lower bound, which has to do with the *Kolmogorov complexity*⁴ of the condition in question (Kolmogorov, 1963, 1965). This consideration is already enough to show that, rather than being mere extensions on essentially equal footing, in image theory properties differ strongly in the extent to which they are practically usable in cognitive processes.

Thanks to this it is possible to evaluate images by their *efficiency*. Although efficiency of image is not the same thing as sparseness of property, the ideas figuring in thoughts about efficiency I take to be valuable to understand what sparseness is. Let us return to the property *grue*: this property has been deliberately given a fragmented description, but the core of its non-sparseness is that it is *disjunctive*. Let us therefore first consider disjunctions in general. Let P and

⁴The Kolmogorov complexity is the length of the shortest algorithm that would pick out the objects in the extension.

P and Q be two predicates, and let $P|Q$ be the predicate of being P or Q . If P and Q are very different from one another, then any procedure for detecting $P|Q$ must include a sub-module for detecting P plus one for detecting Q , which means that its complexity is slightly more than the *sum* of both. Hence, the image for $P|Q$ must be at least as big as those for P and Q together. Maintaining such an image is a costly matter, only affordable if it can be balanced by its benefits.

So what does ‘benefit’ mean in this context? An organism with cognitive abilities will not only judge objects by the way they, so to speak, *look*, but also by the way it will (actually or possibly) be *affected* by them: what their impact will be on its well-being, including chances of survival, reproductive successes, etc. Encounters with objects may take various forms, ranging from seeing them to being attacked by them, and organisms will typically develop abilities to experience the former prior to, and preferably instead of, the latter. What is needed, then, is the ability to use knowledge of a certain property to gain knowledge of *other* properties, in particular, properties that are (more) directly relevant for the interests of the organism. From sound to danger, from shape and colour to mating opportunity, from scent to the satisfaction of hunger.

Hence, a cognitive being will make use of the inferential relations that obtain between properties, e.g.:

- Any prime-segmented centipede is an odd-segmented centipede.
- Any triangular-headed snake is poisonous.

In the interest of this general sketch we shall only discuss monadic inferences (of course there can be polyadic ones too). These have the general form: for any object x : *if* $P(x)$ *then* $Q(x)$. To express this we shall use the following notation:

$$48. P(x) \Rightarrow Q(x)$$

Some inferences are necessary, like the first (if we rule out the possibility of a two-segmented centipede); many more, like the second, are contingent, and of the latter class, most are only probabilistic (also, the relevance for well-being of the second inference is evidently more direct than that of the first, but that is inessential for my present point). Inferences typically form chains that may run from visible (audible, tangible, . . .) properties to information about well-being. If an object is both green and flat, and has a certain scent, it is likely to be a *leaf*. If it is a leaf, it is, under certain further conditions, *food*, and if it is food, it will satisfy your hunger, thus furthering your well-being. Hence, the interesting properties are those that are in the antecedent of interesting inferences, and interesting inferences are those that put you on track to those of the inferential chains that are eventually relevant for well-being. Now a terser expression for ‘being in the antecedent of interesting inferences’ that I will use is *having predictive value*.

Predictive value is an important aspect of the benefit of being able to distinguish a property, and some properties have much more of it than others. But

predictive value is not everything. To know that cobras are dangerous is only profitable if there is a non-zero chance to run into one. Thus, the frequency of a certain property, the relative frequency of its instantiations in the environment of the observer, is another feature that contributes to the benefit of maintaining a image for it. Weighing cost and benefit on a balance should produce a verdict on the efficiency of a image.

Considering all this, what can be said about the predictive value of the disjunction of two properties, $P|Q$? We have seen that its cost is roughly the sum of that of P and Q , but in favourable cases the same is true of its frequency. If so, both effects should keep each other more or less in balance. The bad news comes when considering predictive value. Clearly, whenever, for some R , $P|Q(x) \Rightarrow R(x)$, it must be that *both* $P(x) \Rightarrow R(x)$ and $Q(x) \Rightarrow R(x)$. Thus the predictive value of $P|Q$ is never greater than that of either P or Q , and, if P and Q are of a very different nature, with hardly any property R following from being P *and* from being Q , it will be close to zero. This explains why an image for $P|Q$ is likely to be very inefficient, which makes $P|Q$ an uninteresting property to have an eye for. This, by the present theoretical framework, is the main reason for the bad press that disjunctive properties have.

With respect to disjunctive properties, it must be added that some of their images do possess some measure of efficiency. If P and Q in fact stand for green and blue, then a visual system unlike that of humans might fail to distinguish them (or only under favourable conditions), since they are adjacent in the colour spectre. Here the cost of having a image for $P|Q$ would be less than that of having images for P and Q separately. If, at the same time, the predictive value of P or Q would hardly differ from that of $P|Q$, the latter would have the upper hand after all, thanks to its greater frequency. For a more graphic illustration of this phenomenon, consider images for scanning necklaces strung from beads of different colour. Let us use the predicate RA for necklaces that alternate between red and blue beads, starting with red, likewise BA for those starting with blue. This is the IMG-representation of the image for RA :

start	step	observe	yes	no
1	1	red	— 2	FALSE
	2	end	TRUE	+1 3
	3	blue	— 4	FALSE
	4	end	TRUE	+1 1

(The image for BA is of course obtained by switching the colours.) The disjunction $RA|BA$: *alternating* tout court, ‘feels’ quite natural for a disjunction. The likely explanation for this intuition is that it has an implementation, let us use the predicate A for it, that is *not* a combination of the procedure for RA and BA , and is only one line longer than that of either of the disjuncts:

start	step	observe	yes		no	
1	1	red	—	3	—	4
	2	red	—	3	FALSE	
	3	end	TRUE		+1	4
	4	blue	—	5	FALSE	
	5	end	TRUE		+1	2

Thus, it is sometimes possible to replace disjunctive images for extensionally equivalent but smaller images, which consequently come at a lower cost. This is not a general phenomenon. Alternative implementations for disjunctions are only available in certain specific cases, viz. those where both properties involved are close relatives: blue and green, alternating with red or blue first.

It is often thought that *conjunctions* of sparse properties have better prospects of being sparse as well, and it is certainly true that conjunctive images *can* be more efficient than their conjuncts. On the cost side there need not be much difference with disjunctions: if the conjuncts are P and Q , then in the general case both sub-modules will be implemented in the procedure for the predicate—as we shall write— $P \cdot Q$. However, $P \cdot Q$ has the added predictive value of its conjuncts, plus, interestingly, often something extra. In conjunctions properties may work *synergistically*: their conjunction may support inferences that the conjuncts themselves do not support. An illustration of synergy can be given with the help of the following example. Let BE be the property (of, again, a necklace) of consisting of only one blue bead at the end of the chain, with all the others red:

start	step	observe	yes		no	
1	1	red	+1	1	—	2
	2	blue	—	3	FALSE	
	3	end	TRUE		FALSE	

BH will be the property of being a homogeneously blue necklace:

start	step	observe	yes		no	
1	1	blue	—	2	FALSE	
	2	end	TRUE		+1	1

Finally, One is the property of consisting of one bead only:

start	step	observe	yes		no	
1	1	end	TRUE		FALSE	

Now it is an *a priori* truth that $BE \cdot BH(x) \Rightarrow One(x)$, even though neither $BE(x) \Rightarrow One(x)$, nor $BH(x) \Rightarrow One(x)$ are valid.

To appreciate the specific contribution of image theory to this argument, consider the following objection against this difference between conjunction and disjunction. Let us call it *the triviality argument*. The objection goes that the inferences that disappear by forming disjunctions are not a real loss, since others take their place. Suppose:

$$P(x) \Rightarrow A(x)$$

and:

$$Q(x) \Rightarrow B(x)$$

Then even though it need not be that $P|Q(x) \Rightarrow A(x)$ or $P|Q(x) \Rightarrow B(x)$, it is nevertheless still true that that $P|Q(x) \Rightarrow A|B(x)$, which can also be expressed as:

$$P|Q(x) \Rightarrow C(x).$$

Likewise, the synergistic inference:

$$P \cdot Q(x) \Rightarrow S(x)$$

(where neither $P(x) \Rightarrow S(x)$, nor $Q(x) \Rightarrow S(x)$) has its precursors in inferences supported by P and Q , viz. $P(x) \Rightarrow S|\bar{Q}(x)$ and $Q(x) \Rightarrow S|\bar{P}(x)$ respectively, also known as:

$$P(x) \Rightarrow T(x)$$

and:

$$Q(x) \Rightarrow U(x)$$

so, when counting the number of conditions to be inferred, there is no real gain in predictive value.

The assumption, however, that this line of reasoning depends on, is that all such conditions are ‘equal’. To break the spell of the argument, it would be necessary to show that this is not the case—that they are genuinely different in ways that transcend what extensions they happen to have. Exactly this is what image theory provides. On this account it must be expected that, for certain predicates P and Q , the cognitive system has images for them separately, but not e.g. for $P \cdot Q$ or $P|Q$, since they would become too big. Thus, the availability of innumerable complex predicates, taken for granted as just so many alternatives in the triviality argument, is not warranted: $A|B$ or $S|\bar{Q}$ cannot be counted on a par with A or S , since it could be that, for the cognitive system, the latter are maintainable, but not the former.

Let me end this section with noting that this analysis does not make use of the notion of *projectability*. Projectability has been used as a criterion to distinguish

natural properties from unnatural ones (Goodman, 1965; Quine, 1969), but it is surprisingly easy to show that, depending on circumstances, unnatural properties can be projectable just like natural ones! Suppose in Goodman's heritage a large collection of stamps is found, all of which turn out to be *reallow*, i.e. either red or yellow. Folder after folder is inspected, and only reallow stamps. Then a small booklet appears, with the words 'Most valuable stamp of the collection'. Now, on opening the booklet, what colour do we expect that stamp to have? Should it not be reallow—either red or yellow? (Cf. also Maher 2004; Hempel 1945.) The non-projectability of *grue* is incomplete (it projects *grue* before 1 January 2000, and also afterwards, from observations after that date), and where it obtains, it was built-in in the description: after 1 January 2000 something unexpected had to take place.

Chapter 4

Mereological Fusion as a Case Study

The subject matter of this chapter is the part-whole relation between objects. The Lewisian claim of ontological innocence is evaluated, leading to the conclusion that it has *mereological nihilism* as its almost inevitable consequence. A transcendentalist approach to mereological fusion, however, is shown to bring new insights into the working of procedures: *vicinitive relations* between objects allow there to be loops in procedures, which are important for the representation of pluralities, and for their hierarchical arrangement. This line of reasoning culminates in the Central Assumption of Image Theory, a general hypothesis about the conditions for objects to emerge.

4.1 Parts and Wholes

The objects that we find in the draft ontology are related in many different ways, but one of the most fundamental of these relations—and at first sight most reassuring and fathomable—is that in which one object has others as its *parts*. Mereological fusion is among the favourites of metaphysicians: if there is any promise at all ever to understand an aspect of being, it must be here.

The best-known formal theory about parts and wholes is called *classical mereology*. The towering advantage of classical mereology is its simplicity and (up to some point) its intuitiveness. There is no consensus that it correctly describes our mereological intuitions, but plenty of agreement that it describes *some* of them as well as anyone could wish. One could say that the theory pictures the relation between parts and wholes as it would be in a simple world, which can be discussed in a simple language without too much subtleties. In this simple world *every* combination of objects has a whole, which is another object, and for every such combination there is only one whole. Maybe the part-whole relation just *is* as simple as that—if it is not so simple, it is probably not simple at all.

It does no harm to have a look at some parts of the formalism. According to classical mereology, the relation *is part of*, for which we shall use the symbol \leq , is a two-place relation which is reflexive, antisymmetric, and transitive (Hovda, 2009):

- $\forall x : x \leq x$
- $\forall x, y : x \leq y \wedge y \leq x \rightarrow x = y$
- $\forall x, y, z : x \leq y \wedge y \leq z \rightarrow x \leq z$

Contrary to the more usual approach, my preference is to build up the theory from the concept of *whole* rather than *part*. In this approach the part is what is given first, whereas the whole (also known as the *fusion* of the parts) is derivative. Formally, this order of explanation does not pose more troubles than the reverse one. To capture the core of classical mereology, and all we need for now, it takes 5 axioms. As we will include *uniqueness of fusion* among these axioms, we can use the convenient functional notation $x + y$ to denote *the fusion* of x and y :

1. AXIOM. *For any pair of objects there is a unique object that is their fusion.*

$$\forall x, y \exists! w : x + y = z$$

2. AXIOM. *Fusion is idempotent.*

$$\forall x : x + x = x$$

3. AXIOM. *Fusion is commutative.*

$$\forall x, y : x + y = y + x$$

4. AXIOM. *Fusion is associative.*

$$\forall x, y, z : (x + y) + z = x + (y + z)$$

5. AXIOM. *Any fusion is the fusion of itself and one of its parts.*

$$\forall x, y, z : x + y = z \rightarrow x + z = z$$

Given this piece of theory, we can define \leq in the following way:

4.1.1. DEFINITION. $x \leq y =_{def} x + y = y$.

from which it is easy to prove the three properties of \leq mentioned above.

Axiom 1 is probably the only one of the five that is not wholly uncontentious. In fact it is quite controversial: in one go it states both *uniqueness of fusion* (given the parts there is only one fusion) and *unrestricted fusion* (given the parts there is always a fusion). Both principles have been accepted by some authors and rejected by others (Varzi, 2008).

4.2 Pleading Innocence

The intuition that there is nothing problematic whatsoever about considering the fusion of a given couple of objects as one object seems irresistible. Lewis (1991) identifies *ontological commitment* as the basis on which this conviction rests:

(...) [G]iven a prior commitment to cats, say, a commitment to cat-fusions is not a *further* commitment. The fusion is nothing over and above the cats that compose it. It just *is* them. They just *are* it. Take them together or take them separately, the cats are the same portion of Reality either way. Commit yourself to their existence all together or one at a time, it's the same commitment either way. (p. 81)

The ‘It just *is* them. They just *are* it’ is a concise summary of the idea that fusion is a concept very close to identity. For any cognitive being, if it is committed to the one, and also committed to the other, how *can* it not be committed to the fusion? This epistemic construal of the doctrine is also known as the *ontological innocence* of mereological fusion.

But it is rhetoric. As has been argued by many theorists since, there is nothing obvious about the transition (Van Inwagen, 1994; Yi, 1999; Sider, 2007; McDaniel, 2008). The identification of mereological fusion with identity is in fact hardly defensible. The problem is not that the fusion is *different* from its constituents (although it obviously is!); it is that the fusion need not exist at all, and here the contrast with identity could not be greater. Given an object, there is *trivially* something to which it is identical, but given a plurality of objects, that there is always a fusion is highly dubious. Yet fusion as identity would seem to presuppose the truth of unrestricted fusion to be every bit *as* trivial. And then again: can we be sure that mereological fusion does not ‘add’ something to the mere gathering of objects? I have already mentioned Armstrong’s (1993) awkward fusion of the Sydney Opera House and the square root of -1 . Awkward or not, even the most fervent opponent of unrestricted fusion is hard-pressed to deny that there *is* a sense in which—well, if you *insist*—this ‘thing’ exists. There seems to be little to stop fusion in this minimal sense to be unrestricted. That this *seems* to be so is an interesting observation, and one that might be closely related to the fact that mereological fusion *seems* to be innocent. But even if, for any couple of objects, there is something that ‘just is’ them; that alone is not guaranteed to make the latter their *fusion* as the table is the fusion of its board and legs. Is any old lumping together of things to be regarded as such?

In 2.5 and 3.1 I have already hinted on a negative answer. No doubt there is something which is *the Sydney Opera House and the square root of -1* , but what kind of thing that is, is another matter. There is a set that fills the bill, and for those who are sceptical about sets, there is the plurality of these two items—who could be sceptical about pluralities? Pluralities abound in the draft ontology, only their type is not always clear. Fusions are different from sets: the

latter have a depth structure that the former lack. If I speak of the bricks in a wall, the books on a shelf, or the numbers below 10; although I have mentioned a plurality in all three cases, I have not yet given enough information to decide which of them I have in mind. Pragmatically, the bricks are most likely to be understood as a fusion, the numbers, as a set, and the books can go both ways.

In 3.2 the possibility was contemplated that the innocence of mereological fusion (insofar as it obtains) should be attributed to the working of *free compositionality*, rather than to its putative kinship with identity. And so the following strategy suggests itself: take the shortcut and say that the fusion, unlike other plurality-based types, just is the ‘bare’ plurality. Sets, to get their layered structure, must be more than that, but for fusions the option looks like being available. Ontological innocence would follow immediately. But this seemingly easy solution in fact opens a Pandora Box. The town just is the houses. The houses just are the bricks, which just are the grains of sand and concrete and the fibres of wood, which just are the chemical atoms of various sorts. Rather than setting the stage for accepting compound objects as trivially entailed by their parts, the argument could take an entirely different course. Are the new objects, forged from a collection of old ones, genuine new *objects*, or are they merely a new way to refer to what there already was? If every fusion just is its parts, then it is only a small step to say that only the parts exist. This idea leads us to the doctrine of *mereological nihilism*, a highly insightful contribution to the problem of the relation between parts and wholes.

4.3 The Challenge of Nihilism

If it is not true that every plurality of things has a fusion, maybe *none* of them does, and all compound objects are, as single objects, really some sort of delusion. *Mereological nihilism*, the view that there are no composite objects, only mereological *simples* (Van Inwagen, 1995; Dorr, 2002; Rosen and Dorr, 2002; Williams, 2006; Sider, 2007, 2013b) is, for all its weirdness and counterintuitivity, a perfectly intelligible one. The canonical wording of the nihilist creed, as it can be found in Rosen and Dorr (2002), is as follows:

There are no protons or galaxies or houses of cards. There are rather billions of simple particles arranged proton-wise and galaxy-wise and house-of-cards-wise. (p. 152)

The idea that there are no tables, only particles arranged table-wise, does not necessarily facilitate our discourse about these particles, but has an elegance of its own and arguably makes for a more *parsimonious* theory about the world. For, if composite objects are essentially nothing over and above the particles they are made of—if, in particular, they have no causal powers of their own—what

could assuming their existence possibly contribute to the explanatory power of the resulting model?

As regards the domain of physical objects, where the simples are supposedly mereological (not chemical!) ‘atoms’, there is one potentially fatal objection to nihilism: the world could fail to be atomic. The main troubles with atomism seem to be twofold: (1) the hypothesis has few prospects of being either verified or falsified, and (2) atoms are very hard for us to picture. The latter problem looks surmountable: experience suggests that we can continue cutting things in ever smaller pieces well beyond the limits of our visual imagination, but the fact that we have no clue as to where this should stop is no evidence that it does not. And with respect to the first: even though physics has so far found most ‘elementary particles’ to be complex, when approaching the molecular level the variation declines with smaller scale, and around the Planck scale becomes only a minute fraction of what it is at the visible level. The more reasonable inductive expectation would be that it bottoms out not too far from there, rather than that electrons and quarks turn out to be as complex as rabbits and continents.

Atomism about the physical world is easily misunderstood as giving in to an all-too-intuitive picture of the world. Ladyman and Ross (2009) launch a blistering attack against such predilections. The authors use the term ‘micro-bangings’ to tag the ‘neo-scholastic’ metaphysics that is the target of their wrath. The idea behind ‘micro-bangings’ is that elementary particles—the most basic ones, that physics will one day unveil—behave more or less like everyday objects and that, banging around like that, they constitute all objects great and small, where the latter’s properties just follow from those of the particles. This, however, may be suspecting more naivety among metaphysicians than is really there. Physical simples need not be like that at all; all it takes for an object to be a simple is that it is not analysable as a plurality of other things, connected in some way or other. Commitment to such simples does in no sense imply commitment to microscopic globules that persist in time, move quickly, and are causally effective by bumping into each other just like medium-sized things. But apart from this, the appeal to actual facts about the micro-physical world is not quite to the point here. This is not because such facts are inconsequential, but because the challenge that nihilism poses is less about if and where to find the simples, than it is about the relation between parts (whichever they are) and the wholes they make up. And that is hardly a physical relation. If everything is built up from a certain level of basic constituents, then there can be a version of nihilism taking *that* level as its ‘simples’, giving rise to exactly the same questions about existence. Therefore I shall not make too much of the ‘real’ atomicity of the particles discussed below.

If every object is itself a fusion of lesser ones, forever and ever, then no one should be surprised to find that objects can be fused to produce bigger objects: these would not add anything fundamentally new to the ontology. However, if there *are* simples, then there is room for different reactions. First reaction: If everything is composed of absolute physical simples, then, if *they* would not

merit to be called objects, nothing would. But then the objecthood of everything compound would be a reason for utter puzzlement. If we have a collection of irreproachable objects (the simples), then by what justification would we mess up this beautiful class by allowing in *pluralities* of such objects as well? Are they not completely different ‘things’, only interpretable as some sort of metaphor? Second reaction: No, there is no problem at all. Objects just are, by their very nature, pluralities of simples. A plurality of one thing is no less a plurality, therefore a simple is also an object. Thus, *if* there are simples, then the question before us is what to make of composite objects. Are they real objects or just a convenient way to refer to simples? And if the latter should be true, whence the convenience?

4.4 Nihilese and Mereolese

Let us first scrutinize this latter scenario. Here is the intuition. Apparently, a person who says: ‘There is a table,’ where another would prefer: ‘There is a board with four legs attached to it,’ does not thereby incur a different ontological commitment. There do not seem to be deep metaphysical facts behind there being five objects there or one with five parts (of course referring to the parts separately is not the same as referring to the table: to mention the way in which they are connected is essential). If we could likewise demonstrate that speaking about tables were equally consistently interpretable as just an alternative mode of speaking about the simples they are made of, the ‘conflict’ being as innocent as that between calling the thing on my nose *eine Brille* or *a pair of spectacles*, would that not be enough to explain what our commitment to compound objects really comes to? In this way the dispute between the nihilist and the universalist¹ would be scaled down from an ontological discussion to a disagreement about vocabulary. The contention is that on the ontological level nothing interesting is going on: all there is, is a change of wording while speaking about one and the same thing.

Notice that this is not the same as ontological innocence—which, as the name suggests, is an *ontological* view, not one about mere wording. Arguments of this sort deserve to be examined most critically: whenever, in the context of philosophical elucidation, it is suggested that X is ‘just a convenient way to speak about’ what could also be referred to as Y, the sensible thing to do seems to be to compare X with Y on the linguistic level, to test whether:

- using X holds advantages over using Y (*convenience*);
- X is not informationally stronger than Y (*informational parity*).

¹A universalist is someone who does believe in compound objects.

Let us, to this end, consider two languages, *Nihilese* and *Mereolese*. In the first one can only speak about simples; in the second, about tables and sundry in the usual way. If the first condition is not met, then speaking Nihilese is just as easy and effective as speaking Mereolese. Even though this seems abysmally improbable, to gain insight into the nature of mereological fusion it is probably very useful to have knowledge as to why this is so. If the second condition is not met, i.e. if speaking Mereolese allows the speaker to deliver *more* information than speaking Nihilese, then, convenient or not, it cannot be just a way to say what could be said in Nihilese as well.

For present purposes it is enough to consider the subset of *existentially quantified sentences without individual constants*, since the discussion is about objects, and every conceivable object can be spoken of in such sentences. The task of spelling out what makes a (logical) language a legitimate claimant of being Nihilese or Mereolese is one that we shall face bit by bit. In the first approximation, where Mereolese is unconstrained by further quantificational limits, Nihilese only quantifies over simples.

Taking convenience for granted; when understood in the most straightforward sense (expressions of finite size; quantification over *singular* simples) informational parity can be shown to be highly problematic. A sentence in Mereolese, like:

49. There is a table.

should be translated in Nihilese as:

50. There are simples arranged table-wise.

However, although (50) looks like a statement in Nihilese, I shall show that it cannot be expressed in Nihilese as envisaged here. The reason is that, in this version of the language, quantification is *singular*, and the predicate *being arranged table-wise* is very hard to be understood as applying to singular simples, no matter how many.

Let us use subscripts to indicate whether a formula is in Nihilese ($_{nih}$) or in Mereolese ($_{mer}$). Nihilese-speakers can say things like:

51. $\exists x : N(x)_{nih}$ *There is a simple which is negatively charged*

In Mereolese one can say this as well:

52. $\exists x : N(x)_{mer}$ *There is a simple which is negatively charged*

for, since simples are also objects, and a procedure working on simples is also a procedure in the general sense, every expression in Nihilese is also an expression in Mereolese. However, there is no trivial way to turn the Mereolese sentence:

53. $\exists x : T(x)_{mer}$ *There is a table*

into Nihilese. The similar looking sentence:

$$54. * \exists x : T(x)_{nih}$$

does not provide a satisfactory way to express: ‘There is a table,’ since no simple is a table; arguably, the concept ‘table’ is not even part of the Nihilese vocabulary. Of course for any specific table-configuration Nihilese speakers can sum up all the properties and relations that some group of simples must have to form a table. Something like²:

$$55. \exists x, y, z, \dots : P(x) \cdot Q(x, y) \cdot R(x, z) \cdot \dots_{nih}$$

Now, is this a translation of (50)? It is fairly obvious that it is not, since, even though (53) follows from (55), the converse does not hold. The condition (55) is just a one of an immense multitude of possible implementations of (53).

What options do we have to devise a real equivalent of (53) in Nihilese? First, we could replace the conjunction of predications in (55) by a very lengthy *disjunction* of all alternative conjunctions that qualify as simples arranged table-wise. Apart from the extreme unnaturalness of this solution and the problem of how to decide which disjuncts qualify, it remains to be seen if the expression could be a finite one. Assuming that the number of arrangements for tables up to a certain size is finite, then there still does not seem to be a clear upper bound. It is true that tables cannot have just *any* size, but every crisp border to the size of a table must be highly arbitrary. A related problem concerns the number of simples involved in making up a table. Every expression of the sort of (55) presumes commitment to a *fixed number* of these, even prior to disjoining all the ways in which they can be arranged, which is quite a preposterous demand. The only way to avoid such commitment is to use a disjunction *of* existentially quantified sentences. This, however, only aggravates the unnaturalness of the translation, since it wipes out every resemblance in logical structure between the original expression and its suggested equivalent.

A far more natural way to proceed is to use *plural quantification*. This would yield the following translation:

$$56. \exists xx : T'(xx)_{nih+pq}$$

Yet, this switch to Nihilese plus plural quantification is not an insignificant one. Plural quantification, after all, makes for a *richer* language than plain first-order logic: not every plurally quantified expression has a singularly quantified equivalent (Boolos, 1984). Moreover, it has been shown (Uzquiano, 2004; French, 2016) that a language allowing plural quantification over tables and chairs—one which can be used for logical paraphrases of sentences like:

$$57. \text{The chairs outnumber the tables.}$$

²For typographical reasons I prefer ‘·’ instead of ‘^’ in many-conjunction formulas.

—is *not* sufficiently strong to do so when quantification must take place over the simples making up the tables and chairs. The chairs could outnumber the tables even though there are fewer chair-simples than table-simples. One could think that pluralities of pluralities just boil down to pluralities, but that is giving away essential structure, since:

[S]ingular quantification over composites can be paraphrased as plural quantification over simples, but plural quantification over composites cannot be paraphrased as plural quantification over simples. (Uzquiano, 2004, p 434)

This is because in a language with plural quantification the pluralities themselves (the referents of the xx -type variables) cannot in turn be collected to form pluralities *of* pluralities.

Providing the formal tools to translate (57) is technically undemanding: all that is needed is higher-level plural quantification (where xxx -type variables are introduced which denote pluralities of xx -type pluralities). But that, it might be argued, would go badly against the spirit of plural quantification, where reification of pluralities is explicitly avoided (McKay and MacKay, 2006). This point is particularly weighty with respect to Nihilese, where the referents of quantification are by definition singular entities. In fact, there are reasons to be critical about *any* use of plural quantification in Nihilese. If this language is extended with quantification over whatever is *not* just simples (no non-unit group of simples is itself a simple, after all), by what justification may it be called Nihilese? It seems to make perfect sense to regard the referents of plural variables like xx , as they appear in e.g. (56) as objects in their own right—but if so, then they are clearly *complex* objects. As said before, this is not enough to make them mereological fusions, but plainly enough to undermine the claim that all the objects referred to are simples! And this looks like taking the heart out of the nihilist persuasion: the idea that there are no mereologically complex objects loses much of its appeal if complex objects *tout court* have to be acknowledged.

The mereological nihilist is not forced to accept this conclusion. She can stick to the intended meaning of plural quantification and simply deny that groups of objects are objects themselves, appealing to ‘ways of speaking’ once more. With only formal logic as the arbitrator this dispute cannot but end in a stalemate. It is time to see what image theory can do to solve this problem. To appreciate what the differences between Nihilese and Mereolese really come to, we must look at procedures once more.

4.5 Complex Procedures

Under the assumptions of the nihilist, all the world consists of simples, and so *they*, in the wording of image theory, are the final target of the cognitive procedures that

underlie predication. Let us for the moment waive the practical (and principal) impediments to perceiving individual simples in the interest of the theoretical promises that this option brings, for it facilitates our analysis significantly. If every object is a simple, then, given the objects x and y , the only two mereological possibilities are that they are the same or (fully) distinct, so the—in the eye of the logician slightly awkward—condition of overlap, partial, or near-identity, does not occur. (As a matter of fact these are exactly the conditions that would follow from what, in 2.5, I have called ‘objectual autonomy’). Also, it is plausible that simples always have their properties unequivocally, without room for vagueness.

Before proceeding to this part of the argument, let me stress that I do not in any way propose to *believe* in simples as such. Simples are an assumption that is put to test, to see where it leads. Apart from this, it is likely that some things, in certain respect, behave like simples to an appreciable extent, thanks to which the thoughts developed below have a broader application as well. Simples probably allow for an invaluable *simplification* of theories about how procedures work, and simplifications are often instructive.

In Chapter 3 we have looked at procedures in a loose-handed way. When considering properties of necklaces with the colours of beads as the fundamental input, we have taken for granted the underlying procedure that should recognize the beads themselves and determine their colour. We have also helped ourselves to an easy case, for beads behave like simples in important respects: they also do not overlap. With simples themselves being studied, no such hand-waving is necessary anymore, since they would provide us with the perceptual end-point of the procedure: the sub-procedures that scan individual simples. Let us call the latter *base procedures*. All they do is collect a Yes/No-answer: no computation is involved.

The most elementary of the base procedures test one simple for one monadic property. Obviously the Nihilese-speaker must have words for such properties, but to arrive at a concept like *arranged table-wise* more is needed, to wit, relations. To arrange themselves as a table, simples must be related in all sorts of manners; in particular, they must be spatially positioned toward each other some way or other. Traditionally, two main classes of relations are distinguished: *internal* and *external* relations. This distinction has been defined in several ways, but in image theory the most natural understanding (here we stay closest to Armstrong’s (1978) interpretation) is that the internal relations are those for which procedures are possible that only use the output of procedures scanning the relata separately. A shorter way to say this is that the relation *supervenes* on the monadic properties that the relata have separately. The most basic among such relations are mere combinations of two monadic properties: it will be the case that $R(x, y)$ whenever $P(x)$ and $Q(y)$.

Slightly more interesting is the class of relations to which *similarity relations* belong. Let us assume that simples have ‘colours’: all simples have one of 100 colours, C_1 through C_{100} ; and let S denote the binary relation of having the same

colour. This time, there is not one single combination of properties of x and y on which $S(x, y)$ supervenes. What the procedure for S could do is first determine for which i it is the case that $C_i(x)$, and subsequently, whether $C_i(y)$ holds as well³. So, contrary to the type of relation to which R belongs, relations like S do not decompose into a combination of monadic properties. Yet, they are internal: it is possible to build procedures whose outcome is wholly dependent on that of certain underlying procedures for monadic properties. External relations do not supervene on monadic properties at all. To find out if two simples x and y are *adjacent*—let us write $A(x, y)$ — within some bigger structure (e.g. a ‘molecule’ or just space itself) it will not do to collect properties of x and y separately. As there are no procedures underlying the ones determining external relations, they can be seen as more primitive than internal ones. The most convincing external relations are spatial and temporal ones, plus all relations that somehow involve them.

All such relations must be part of the vocabulary of the Nihilese-speaker. Base procedures are primitive, and they combine to form *complex procedures*, where scanning goes hand in hand with running an algorithm. Hence, it will usually be possible within complex procedures to discern *sub-procedures* that produce their own outcomes, functioning as signals which are being processed further in the course of reaching the overall, *unified* outcome. A theoretically simple option for combining sub-procedures is truth-functional recombination: the joining of signals by logical gates (AND, OR, XOR, ...). This, however, is a very simple way of creating complex procedures out of more basic ones, and not too profitable from the viewpoint of information processing. Below we shall consider more sophisticated ways of building procedures out of sub-procedures.

Procedures for properties of compound—thus complex—objects should be expected usually to be themselves complex. The details of how they work are pivotal for understanding the relationship between the predications under which the parts fall and those applying to the whole they constitute in virtue of its being so constituted (e.g. being even-segmented). To be relevant for the comparison between Nihilese and Mereolese it is inessential that the parts are simples, and since exactly similar goings-on may be found at any level of complexity, we may as well take an intuitive case to study them: the aforementioned table will do. If B and L are predicates for *being a board* and *being a leg*, respectively, then let us compare the expression in first-order logic of there being a table (T) with that of there being a board plus four legs properly connected (C).

$$58. \exists x : T(x)$$

$$59. \dot{\exists} a, b, c, d, e : B(a) \cdot L(b) \cdot L(c) \cdot L(d) \cdot L(e) \cdot C(a, b) \cdot C(a, c) \cdot C(a, d) \cdot C(a, e)$$

(where $\dot{\exists}$ should be read as ‘there are *distinct* ... such that’).

³Of course we cannot be sure if this is what actually happens: it might be that nature allows for more efficient ways to determine whether $S(x, y)$.

For the sake of argument, we shall take the meaning of the predicate T to be such that it refers to exactly the kind of table that is obtained by fixing four legs of type L to a board of type B according to relation C . Clearly, (58) is logically equivalent with (59). Now it would be interesting to know if (59) is exhaustive as an analysis of (58) or not. (Of course we could go deeper, take legs and board apart and so on, but that would be a repetition of a similar type of analysis.)

The compositional character of predications both in formal and natural language appears to reflect the key feature of complex procedures, viz. that they break down into sub-procedures. An expression like (59) clearly displays this trait. Here is how we could interpret (59) as a complex procedure:

- Find separate objects a, b, c, d, e $\exists a, b, c, d, e$
- Test whether a is a board $B(a)$
- Test whether b is a leg $L(b)$
- Test whether a and b are properly connected $C(a, b)$
- (etc.)

Completing this operation for all four legs, subsequently joining the results by an AND-gate, would yield an overall procedure with a unified outcome. Notice that it is hardly conceivable that anything like this actually happens, however, or it would be a Herculean task to perceive a table, since the outcome of the conjunction would virtually always be negative. Collecting arbitrary board and legs, to test them for being properly connected only afterwards, would rarely yield so much as *one* leg connected to a board, let alone the very lucky combination of hits that would bring the table as a whole within sight.

This does not seem to be how the procedure works. Perceiving compound objects—or so I would conjecture—typically starts with gaining sight of *one* of the parts, after which the cognitive system starts to look for the other parts. Seeing a board, after all, does not *guarantee* the presence of a table, it only strongly increases its likelihood. And not only does it increase the odds of finding the other parts of the table; the initial observation also involves important hints as to *where* to look: in the immediate *vicinity*. To get full grasp of what this means requires a separate analysis.

4.6 Vicinitive Relations

There is, or so I would propose, a distinction to be made among relations that by and large coincides with their being internal or external, but is of a very different nature. In order to get some sense of the difference, consider the following examples. Suppose I want to know if there is someone looking exactly like me, my

doppelgänger. There could be two of them, or one hundred, but to know if this relation is instantiated I will have to do some *searching*. If I find a doppelgänger of mine, I am done: the relation is realized, with me as one of the relata. But if there is *no* such doppelgänger anywhere, how will I know? After having checked every person on Planet Earth, I may still ponder over remote galaxies with planets where evolution has brought forth Homo Sapiens as well, with my doppelgänger possibly among them. So there could be a positive answer for me to be found, *but not a negative one*. Compare this with the situation in which I want to know if there is someone, not *looking exactly like* me, but *standing next to* me. In reasonable circumstances one glance beside me will do, and I can equally well find out whether someone is standing next to me *or not*. Be it positive or negative, an answer will be forthcoming. Finding an object under such circumstances could be described as *picking* rather than searching.

Relations of the latter sort I shall call *vicinitive*. Vicinitivity is a condition that is strongly tied to the transcendentalist approach; in particular to the way procedures work⁴. Many vicinitive relations have the feature of being *uniquely determining*. If everyone is holding hands with the person next to them, then there can be only one person left of me, and one right. If Sarah is left of me, and the mayor is left of me as well, then Sarah must be the mayor. Such a thing is impossible in the case of the doppelgänger. If on Monday I hear that a doppelgänger of mine has been found in Brazil, and on Thursday I learn about one running a hamburger restaurant, there is no particular reason to believe that they are the same person.

The dichotomy may suggest a cleaner split than there really is. Some relations are extremely vicinitive (being adjacent links in a chain); some, entirely non-vicinitive (being of the same shape or colour), while others are somewhere in between (being family members). Some vicinitive relations are uniquely determining and some are not. Although it is not conceptually impossible for internal relations to be vicinitive, there seems to be no reason to expect it: why should the availability of an object with certain intrinsic properties ever be guaranteed? On the other hand, almost all external relations that we know appear to be vicinitive to some extent. This is true for spatio-temporal relations, and if Hume is right and causality is external, then it is no exception, if only because causes and effects are always connected spatio-temporally as well.

⁴In computer programming, there is a parallel in the way data can be accessed. A computer application has only got one way to really ‘pick’ a piece of information, viz. going to its memory location. If information has to be accessed by any *other* type of description, then it has to be searched for. This is typically the case: what we want to know is Mr. Cardoso’s email address, not what is the value of the first fifteen bytes starting from 0x81A32:0F25. Searching, in this case, means that the memory locations have to be run through one by one to see if the information is there. Much of the technology of database searching is therefore about making this process more efficient by establishing convenient relations between content and memory location (so-called *indexing* of data files).

Let us return to the composition of the table. What is missing out in (59) as an analysis of (58) is that the availability of a vicinitive relation—in this example it is the relation C , of being ‘properly connected’—makes a huge difference in how objects available for cognitive processing are being *found*. To find a table is not to find four legs and a board and subsequently to test whether they fit. It is to find one of these parts first, and then to look for the other parts *where they are to be expected*. But that is only possible thanks to the fact that C involves spatial relations, and these are vicinitive *par excellence*. Thus, the procedure behind (59) now becomes:

- Find an object a $\exists a$
- Test whether a is a board $B(a)$
- Pick an object b , properly connected to a $\dot{\exists} a, b : C(a, b)$
- Test whether b is a leg. $L(b)$
- (etc.)

Rather than collecting five objects to see if they have the required properties and relations (which, to repeat, would be a very improbable event), only the first object has to be searched. The others can be picked, whereby the efficiency of the procedure is vastly improved.

Efficiency, however, is only one reason why vicinitive relations are real game changer for complex procedures. An even more fundamental consequence of having the option of ‘picking’ an object rather than searching for it is that procedures can contain *loops*. For a loop to be functional it is essential that it can *terminate*, and, since picking can result in a *conclusive* failure (no new object), it has this option in a straightforward way. In the IMG-flowcharts discussed in Chapter 3 this is evident from the presence of external actions like ‘moving up one segment’ (of the centipede) or ‘moving up one bead’ (of the necklace), which are possible only thanks to the spatial connectedness of segments and beads, respectively. Thanks to this feature there can be a procedure for *being a homogeneously blue necklace*, (cf. the image on page 70) which has no restriction as to the *number* of beads to be considered. The loop is run through as many times as is necessary and terminates at finding a non-blue bead or the last bead in the chain.

Complex procedures of this kind have no analogue in first-order logic. Clearly, it is possible to write a logical expression for homogeneously blue necklaces of any number of beads. For three beads it looks like this:

$$60. \dot{\exists} x, y, z : B(x) \cdot B(y) \cdot B(z) \cdot T(x) \cdot T(z) \cdot A(x, y) \cdot A(y, z)$$

(where B indicates being a blue bead; A , adjacency, and T , being a terminus in the chain). But it is not possible—without adding profoundly new expressive resources—to have an expression that *abstracts away from the number* of beads

in the necklace. This property of predicates—their abstracting away from the *number* of constituents of an object, while retaining their properties and the way they are connected—I shall call *plurality abstraction*. What predicates of this sort express is *being a such-and-such plurality of P's*, where the ‘such-and-such’ part handles the ways in which the *P's* are linked to one another.

Plurality abstraction is a feature of many predicates in the draft ontology—think of *ash grove*, *chain*, *people*, *constellation*—but also of stuff terms, like *sand*. In the interpretation I advocate the job of plural quantification in logic is precisely to convey plurality abstraction, thereby mimicking the expressive power of images encoding loops in the procedure. And these are the ones employing vicinitive relations. Now, loops can contain sub-loops, containing sub-sub-loops, and so on, down to any level of nesting; hence, what requires ever new types of pluralities in the logical formalism (*xx*, *xxx*, *xxxx*, ...), comes naturally in images. This is what makes the latter essentially richer semantically than the former.

The conceptual watershed that is marked by exploiting vicinitivity appears to have much to do with our ability to perceive compound objects. Without the possibility to explore reality step by step, one part after the other, the latter could only be encountered in a haphazard way, and their shared structure could only be appreciated afterwards. Vicinitive relations are not only detected: they actively lead the way in the exploration process. It may be overstating to say that, without them, complex procedures would not exist at all, but they could not reach far beyond mere truth-functional recombination of properties of parts. Vicinitivity makes complex procedures worthwhile: thanks to it they can be efficient enough to have value for the cognitive system.

So much for the relevance of vicinitive relations for the feasibility of complex procedures; in the next section I shall use these ideas to launch a general hypothesis about the conditions for objects in general to appear. Meanwhile it is worth noticing that the meaning of this phenomenon in our cognitive life far outreaches matters of mereological fusion only. Vicinitivity is about how objects—genuinely new ones, unrooted in others—enter our consciousness in the first place. In 2.4 I stated that objects never appear to us in ‘bare’ form. Now we can be more specific: their first anchor of identification is often their spatio-temporal positioning relative to what is already known. That this relation is vicinitive is what allows us to seek out these new objects *actively*, and, what is more, to make certain fair predictions about what they will look like. We find our homes back thanks to the fact that the roads and tracks, along with landmarks like trees, buildings, and bridges, are connected vicinitively. We find the bridge by walking down the street until the end, then arrive at the brick road by crossing the bridge, eventually to reach the front door. If the latter would have to be found with its qualities as our only guideline (rectangular, green painted, with a nickel knob and a small round window) we would be clueless. Following vicinitive chains may also bring surprises, but nothing compared to what we would experience if new objects would enter our attention entirely unrelated to what has come to be known so far.

4.7 Emergence

Complex procedures are built up of sub-procedures which can be built up of lower sub-procedures in turn. Let us, as we did in the previous section, assume that this goes on all the way until the level of base procedures scanning the simples. Looking at the matter in this way makes it fully transparent how compound objects are scanned: it is by scanning the simples they consist of. All this *can* be expressed in Nihilese, but only if the language has the expressive power, without restriction, to organize pluralities as groups of lesser pluralities—that is, only if plural quantification up to any level is allowed. All this *can* also be expressed in Merelese, but then the results of all the sub-procedures will be interpretable as applying to compound objects, composed of the objects associated with the lower sub-procedures. The conclusion must be that the differences in descriptions as they occur between Merelese and Nihilese *do not mark any difference in what lies behind the description in terms of procedures*. Therefore, they do not mark any difference in terms of images. The image for the table is exactly the same thing as the image for the simples arranged table-wise.

This is remarkable, since there is a very obvious *ontological* difference between there being simples plus a table on the one hand and there being only simples on the other. The table is there or it is not; if only the simples exist and none of them is a table, then there is no table. In 4.3 I have considered the possibility that to be an object ‘just is’ to be a (possibly single-member) plurality of simples, but there are several reasons to be suspicious about such a scenario. It does not seem to have much to say about the structure of our mereological convictions: *why* do we unite certain groups of simples to the exclusion of the vast majority of them? Also this idea is fairly *ad hoc*: it provides a fix for the problem of mereological fusion, but does nothing to explain how other types of rooted objects come about. Reasoning from realistic premises it is indeed difficult to see what another, more promising answer might look like. Fortunately we have, in the course of this chapter, unravelled enough of the details of mereological fusion on the level of images to sketch the outlines of a transcendentalist account of compound objects. The image for the table is the same thing as the image for the simples arranged table-wise, but it is *not* the same thing as all the images of the simples taken together, since the latter are united under some specific arrangement. This being united I take to be what matters. It is pivotal for what I shall call the *emergence* of objects: their transcendental appearance before the mind’s eye.

So here comes the main hypothesis of the book. The rough idea is that emergence coincides with the points where information flows together. Figure 4.1 is a sketchy representation of the overall state of affairs with respect to the flow of information through a certain complex of images. The collection of simples (small circles—never mind the scale) can be interpreted as a table with a bench next to it. In the first layer we find the base procedures: the scanning of each of the simples separately. In the second layer the output of some number of these




Figure 4.1: The flow of information and the emergence of fusions

base-procedures is gathered by first-level complex procedures (e.g. by a loop in the algorithm), to yield a new single output. In this simplified example this part of the procedure coincides with the emergence of the board and the legs of the table. One level up, the information from both legs and board becomes united, causing the emergence of the table. Of course it need not stop there. The bench emerges along similar lines, and on a still higher level ‘the furniture’ appears as a unified whole.

Again, the ‘simples’ *need* not be taken seriously. The argument is largely the same if they are really some basic level of perception. The upshot of this slightly impressionistic exposition appears to be this: what on the image-level is the informational hierarchy of images, *in the realm of objects takes the guise of mereological fusion*. Emerging objects find their origin everywhere the output of sub-procedures is jointly deployed to be the input of a further procedures. In Chapter 3 it has been argued that the observation of properties by the cognitive system should be associated with the working of images; now it seems we can say more about what images are. Their somewhat vague characterization as ‘functional units of a cognitive system with a specific ability of processing information’ can be made more specific: images should probably be identified with those units which produce a *unified output*, for it appears to be these structures within the cognitive system which are the crystallisation cores for the emergence of objects.

This I shall call the Central Assumption of Image Theory:

Objects emerge wherever cognitive procedures reach a unified output.

The more tangible consequence of this is that objects at least emerge wherever *predication* is active⁵.

Although this seems a fair way to use the transcendentalist approach to account for part-whole relations in the draft ontology, as a hypothesis it is still somewhat *ad hoc* if it cannot be embedded in a more encompassing theory, one that also works for cases other than that of mereological fusion. This will be the subject of the coming two chapters. Before proceeding to this part of the argument, let us close this chapter by taking stock of what the Central Assumption does for the explanation of mereological fusion, for it is quite significant. Let us, to this end, go back to the viewpoint of mereological nihilism, and realize how reasonable it in fact is. Setting aside our common human prejudices, it should be acknowledged that, realistically speaking, *if there are parts in the world, there need be no wholes*. A world with only parts is, up to the smallest detail, *identical* with one with the wholes added, and this holds down to the smallest parts.

According to transcendentalism, the concept of object is a categorizing principle that human cognition imposes on reality. The emergence of objects on top of those that were already there—an idea deeply troublesome for the realistic view—is transcendently unproblematic: it is in fact just what might have been expected. If objects are essentially a matter of perspective, then shifts of perspective are bound to bring about new objects, whereas there is also a clear sense in which they are not *really* new. Compound objects are no addition to the parts they consist of—and yet they are. The ontological presumptions of Nihilese, therefore, are legitimate, even though they belong to a different (in fact profoundly non-human) perspective. And Mereolese is not another way of speaking about the subject-matter of Nihilese. It has its own subject matter⁶.

Observe that this does not only apply to mereological fusions. If the above is correct, then it is true for any plurality whatsoever that becomes unified in some way or other. Unification is *always* a conceptual manoeuvre. It is never ‘just there’. As already shone through in the equivocality of the phrase ‘just the coming together’—bearing the suggestion that nothing *happened* to whatever came together—the very idea of free compositionality, and with it, of ontological innocence, is quite illusionary: *there is no such thing*. And here we seem to have an answer to the old conundrum of whether *unrestricted fusion* is true. To argue for all sorts of arbitrary gatherings of objects having their fusion is counterintuitive, yet it is notoriously hard to draw the line. Seen through the lens

⁵Cf. also Evans’s use of the term ‘Apperception Engine’ in machine learning (Evans et al., 2019, 2020). Following Kant, the author describes apperception as ‘the process of assimilating sensory information into a coherent unified whole’.

⁶The comparison with Nagel’s (1974) *bat* comes to mind: a being from whose perceptual (echolocation), and therefore—doubtlessly also—*conceptual* point of view the world must look completely different than it looks like for humans. Due to acoustic circumstances unbeknownst to people, bats may well acknowledge objects absent in our world.

of image theory, however, the answer suggests itself: our epistemology will have a preference for those fusions that are picked out by relatively *efficient procedures*. Recognition of an object as such means that there is at least one image which can pick it out as falling under some predicate. Unrestricted fusion is true in the general sense that, for every plurality of objects which can be recognized as such by our cognition, their fusion can be recognized as well, viz. by combining their images into one image. But the intuition that some of these fusions are strongly preferred over others finds its justification in the fact that certain complex predications allow for *compression* (as described in 3.5). Objects falling under them may be highly compound, yet the regularities and repeating patterns they contain make it possible for relatively small images to capture them.

If this analysis is correct, then it highlights what was implicit in the transcendentalist interpretation of objects all along, but attains new relevance in the context of fusion: that unity is not ‘noumenal’. That the result of a complex procedure is conceptualized as perceiving a property of *one* object, rather than a multitude of them, is not the effect of the target’s *being* undisputedly singular; it is the effect of the procedure’s being a functional unit, i.e. a feature of the cognitive system. But if unity is not in the things themselves, there is little reason to believe finding an exception in ‘simples’ and hardly more to grant them independent existence or even the status of useful theoretical tenet. This is not an argument for ‘gunk’ or non-atomicity. It is still possible that physics will come up with a class of elementary particles that, with excellent arguments, can be declared final; but if so, then their objecthood is just as much a matter of human perspective as is that of my goldfish bowl. What lies behind is utterly beyond us.

Chapter 5

Abstraction

It will be argued in this chapter that *abstraction*—the emergence of objects like *the tiger* and *whiteness* out of tigers and white things respectively—is a phenomenon showing many similarities with mereological fusion. First an argument is presented against the putative *triviality* of talk about abstract objects. Abstract objects form a large and strongly variegated class; but rather than approaching this class as a whole, some of the most conspicuous representatives, viz. *kinds* and *attributes*, will be studied. By the suggested transcendentalist interpretation of these types of objects, generality turns out to be the default condition and individuality the exception, brought about by the properties of space. The chapter closes with a short exploration of the territory of ‘deeper’ abstractions and the phenomenon of perfection.

5.1 Abstract Objects

In Chapter 4 I have discussed one of the most prominent ways in which objects emerge from pre-existing pluralities: the composition of a whole from parts. I have called this a case study, since there are other ways in which pluralities give rise to new objects. Here we shall consider a different major principle of emergence, for which the word *abstraction* seems to be the most apt. Abstraction, like fusion, makes new objects emerge out of already existing ones, objects which—this will be my interpretation—are *rooted* in the latter as well. But the result of abstraction is very unlike that of fusion. One of the most striking differences is that the new objects seem to be of a different *type* than the old: they are *abstract objects*, whereas fusions, whatever their ontological credentials, are every bit as concrete as the things that were fused.

Before trying to explain their origins, let us see which objects are eligible for explanation, and what has been said about abstract objects in the course of history. In the former chapter we have seen that it is possible categorically to

doubt the existence of mereological fusions, but this scepticism, however defensible, is not widespread, nor does it have a long history. For abstract objects the situation is dramatically different. Quine's rejection of them as 'creatures of darkness' is well-known, but already in the writings of John Locke we find eliminativism with respect to abstract objects:

To return to general words: it is plain, by what has been said, that general and universal belong not to the real existence of things; but are the inventions and creatures of the understanding, made by it for its own use, and concern only signs, whether words or ideas. (Locke, 1999, III Ch. 3 Sec. 11)

'General words' are to be understood as referring to mental principles of classification:

That then which general words signify is a sort of things; and each of them does that, by being a sign of an abstract idea in the mind; to which idea, as things existing are found to agree, so they come to be ranked under that name, or, which is all one, be of that sort. (*ibid.* III Ch. 3 Sec. 12)

Also Frege notes the intimate connection between abstract objects and realm of the mental. He is unwilling to identify both, however, since that would block the possibility of common knowledge and truth about the matter:

If every thought would need a bearer, to whose consciousness it belongs, then it is a thought of this bearer only, and there is no knowledge, which would be common to many, on which many could work; but I may have my knowledge, namely a whole of thoughts that I am the bearer of, another has his knowledge. Each of us deals with the contents of his consciousness. (Frege, 1918, my translation)

Hence, Frege believes a third realm should be acknowledged: 'Thoughts are neither things in the outside world, nor ideas' (*ibid.*).

Despite the popularity of nominalism for a good part of the twentieth century, in recent times realism about abstract objects is no longer a maverick position. Of course there is Armstrong's (1980) realism about universals. Also Wetzel (2009), espouses realism about what she calls *types* (as opposed to tokens, concrete things):

I think we should conclude that the generic objects, the types, apparently referred to in the data *exist*—that is, species, genes, epigenotypes, languages, body parts like the larynx, syllables, vowels, allophones, computers like the Altair 8800, Mozart's Coronation Concerto, the Queen's gambit, the hydrogen atom, the football and so on. (p. 23)

In the foregoing I also have insisted on the reality of an objects like *whiteness*. In the terminology aptly criticized by Gideon Rosen (2018), this position would qualify as ‘Platonism’. My agenda, however, is not to argue for the reality of abstract objects as an entirely new type of objects, to be added to the stock of physical ones; nor to allow them a realm of their own. Rather, it is to show that there is no sharp boundary at all between the one and the other. To sustain this point, the following expressions might serve as examples:

- 61. The sky is blue.
- 62. If I were you, I would go to the doctor.
- 63. In the fifties, we all listened to the radio.

The sky, *the doctor*, and *the radio* all seem to be hovering somewhere in between concrete and abstract. For what is the referent of ‘the sky’: is it just the local expanse of the atmosphere above the speaker, or is it the one thing we see each time we look up, no matter where or when? Or else, is it the mereological whole of all the local expanses (on a sunny day)? The latter understanding is not available with respect to ‘the doctor’, but otherwise it remains unclear whether some specific practitioner is intended, a variable for substituting one’s own doctor, or ‘the doctor’ as a general term.

For ‘the radio’ it is also possible to read the *content* that was on the radio, and here we have another issue.

- 64. We heard the Queen’s voice on the radio.

Let us assume it is a live broadcast. In (64) it *seems* that we hear a concrete utterance of words, a concrete sound. But thanks to the mediation by electromagnetic waves and electronic devices, millions of people could be hearing the Queen’s voice. There is a sense in which none of them hears the concrete sound, and although all *copies* of this sound are concrete as well, we certainly do not take ourselves listening to *another* sound than that of the Queen’s voice, far less, to something different than everyone else is hearing right now. What we all listen to we doubtlessly accept as the one and only *Queen’s voice*. So is the Queen’s voice on the radio abstract or concrete? Similar questions come up with regard to much that has entered our world with the dawning of the digital age. Are *files* concrete or abstract? What of applications or messages on social media?

This chapter, therefore, is not about abstract objects as a general category. I have no ambition to provide an overall account of these—I believe this class to be of a highly diverse nature; maybe rather a rest-category. My aim is not to demarcate the abstract from the concrete, but to see if it is possible, from a transcendentalist perspective, to understand of certain types of abstract objects, what accounts for their appearance alongside physical ones. The exploration will start with what in the draft ontology is most immediately given, attempting, bit

by bit, to extend the known territory. Before embarking on this project, however, there is a matter of principle that has to be dealt with.

5.2 Triviality

In Chapter 4 the idea of ontological innocence was discussed (and rejected). Ontological innocence is the—*prima facie* quite persuasive—idea that there is something *trivial* about parts composing a whole: that metaphysicians see a problem where there is none. There is a clear similarity between this view on mereological fusion and the line of reasoning advanced by Schiffer (2003) and Thomasson (2013; 2014) with respect to (certain types of) abstract objects (cf. also Hofweber’s (2016; 2019) dual quantifier approach). Thomasson endorses what she calls *meta-ontological deflationism*. The idea is that *trivial entailments*, as they are found in common language, take us from undisputed claims to existence claims about (e.g.) properties. Here is an example (*ibid.*):

Undisputed claim: This bag is red

Transformation rule: If N is Q then N has the property of Q-ness

Transformed claim: This bag has the property of redness

Ontological claim: There is a property (namely of redness).

The upshot is that commitment to the undisputed claim automatically leads to commitment to the existence to the property of redness *in the only sense it has*. The existence claim is thus as real as can be, yet nothing ‘deep’ is going on. The above example is about the existence of properties, but with respect to, e.g. propositions and mereological fusions similar schemes can be built.

Although there is little doubt that this approach yields a correct description of the relations between concrete and abstract objects in the draft ontology, it stops short of providing a *reason* why this is the way it goes. The use of the term ‘trivial’ and, consequently, the rejection of the sort of intellectual toil that seeks to find ‘deep’ or ‘heavy-duty’ ontological results, suggests that such a reason is not needed: going after one is making far too much of it all. But that remains to be seen. The transformation rules themselves look very simple indeed, but the answer to the question *which are* the valid rules does not seem to be trivial in any way. Sometimes rules that, for quite a long time, were believed to work fine, turn out to perform disappointingly when put under strain:

Undisputed claim: *Being concrete* does not apply to itself

Transformation rule: If N is Q then N has the property of Q-ness

Transformed claim: *Being concrete* has the property of non-self-applicability

Ontological claim: There is a property (namely of non-self-applicability).

And so we have an ‘easy’ version of Russell’s paradox (cf. also Frost-Arnold, 2016). The conclusion must be that some such inferences work while others do not, and to see why does seem to require some effort.

There are more reasons to be suspicious about triviality claims of this sort. Thomasson does not present her account of ontologically contentious objects as *fictionalism*, nor does she claim that they are ‘merely verbal’ items:

[W]hile the meta-ontological approach is deflationary in the sense of rendering disputes about the existence of these entities pointless, the first-order metaphysical view that results (about numbers, properties, etc.) is not deflationary. Instead, the result of this deflationary meta-ontological approach is typically a first-order simple realism about the disputed entities. (Thomasson, 2013)

However simple this realism, *nominalists* about properties must of course reject it, which suggests that it cannot be so ‘simple’ that every sensible person¹ readily appreciates it as such. Indeed, although speaking Nihilese is not a serious option for any human community, speaking, say, *Concrete* (a language that only quantifies over concrete, say physical, entities; somewhat like Sider’s (2013a) ‘Ontologese’) is feasible. Maybe, as per the Quine-Putnam indispensability argument, it must be feared that the Concrete-speaking community would have a hard time developing science, but they would probably enjoy plainly enough communication to survive. Although they cannot speak of ‘redness’, they are not bereft of the ability to use predicates: a sentence like ‘Red snakes are the dangerous ones’ is well within their reach.

This is not to say that there are not excellent arguments against nominalism (for an instructive tour past the main snags, see e.g. Wetzell, 2009), only that the existence of abstract objects is not trivial, nor does it allow of an easy explanation. That said, there need not be a ‘one size fits all’ solution for the existence of abstract objects either. Some of them might be ‘merely verbal’, others, a useful fiction or metaphor, while still others could have a serious claim on ‘genuine’ existence. All these cases deserve to be considered one at a time.

5.3 Kinds

There are many objects of which abstraction is the origin, but for the explanatory route I have in mind it makes sense first to look at *kinds*.

What is a kind? This will be my working definition: a kind is an abstract object that has concrete objects as its *exemplifications*, where the latter term

¹My presumption—that I will not give up—is that there are (or at least once were) sensible nominalists.

refers to the fact that kinds are in many respects *similar* to their instantiations, and as such are something like a ‘model’ for them (Marcus, 2001, Ch. 5). Names of kinds in English are often (but not always) preceded by the definite article, and in most (but not all) cases, replacing the definite with the indefinite article will indicate a concrete instantiation. Here are some examples of kinds plus their instantiations:

- The violin – a violin
- The tiger – a tiger
- The Marseillaise – an occasion on which it is performed
- The word ‘kind’ – an utterance of the word
- The Ford Escort – a Ford Escort
- Man – a person
- Carbon – a piece of carbon
- Carbon – a carbon atom

(The latter two examples show the indeterminacy that sometimes obtains as to what should count as an exemplification of a kind.)

Kinds are a relatively easy prey for the just-a-way-of-speaking approach. Even Frege—surely no nominalist generally—in his discussion of the sentence ‘The horse is a four-legged animal’, tends in that direction:

The (...) sentence is probably best regarded as expressing a universal judgement, say ‘all horses are four-legged animals’ or ‘all properly constituted horses are four-legged animals’ (...). (Frege, 1951)

There are several problems with this type of solution—as with the various refinements that have been proposed subsequently (e.g. Goodman and Quine, 1947). Attempts to find a *general* rule for paraphrasing away all reference to kinds or other abstract objects, so as to end up with purely nominalistic discourse, have on the whole proved unsuccessful. No doubt many speakers who utter: “The Q is P,” will gladly agree that, on that occasion, they could as well have said “All Q’s are P,” but the exceptions are manifold and stubborn. Frege already seems to have been aware of the problems looming, given his second paraphrase, containing the qualification ‘*properly constituted*’.

Kinds can be placed at the one end of a scale where on the other end we find *attributes* (I shall use the word ‘attribute’ as a perfect synonym of ‘property’, in contexts where it improves clarity). There is a sense in which everything on that scale is an attribute, as all such conditions can be predicated of objects: the

difference between *the violin* and *whiteness* seems to be first and foremost that the latter is more abstract than the first (I shall be more explicit about this below). Be that as it may; there is a tradition, starting with Aristotle, to keep kind and attribute distinct (by narrowing the sense of the latter), since the differences are significant and interesting. The unity of all violins is far stronger than that of all white things, and for a given object its being a violin is a far better answer to the question ‘What is that?’ than its being white. In Aristotelian ontology this distinction can be found as that between (secondary) *substance* and *quality*. The analysis to follow will first focus on kinds.

Kinds can be predicated of objects, much like attributes: ‘This dog is a poodle’ in addition to ‘This dog is white’. But kinds also have a life of their own, which, to a remarkable extent, parallels that of their exemplifications. *The poodle* is an object in the draft ontology. It is one of many dog breeds, more intelligent than most of them, and with a curly coat. Thus *the poodle* has many of the properties that individual dogs have. There may be some doubt whether *whiteness* is itself white (just like white dogs), but not that *the poodle* barks, just like individual poodles do. As Wollheim (1980) puts it (Wollheim uses the type/token terminology):

[F]or much of the time we think and talk of the type as though it were itself a kind of token, though a peculiarly important or pre-eminent one. In many ways we treat the Red Flag as though it were a red flag (cf. ‘We’ll keep the Red Flag flying high’). (p. 50)

This, however, does not stop kinds from showing some conspicuous differences when compared with dogs and tables. They lack the latter’s spatio-temporal positioning, they are not causally efficacious or accessible to our senses in the same way, and they are unique (there is only one of each kind), to mention some major divergences with ‘normal’ objects.

5.4 The Primacy of Generality

Suppose my cognitive system has an image for tigers. If a tiger passes by, it has the ability, mediated by my sense-apparatus, to recognize it as such. Thanks to this, it is also possible for me to think about tigers, bring them to mind even before seeing one, draw inferences with respect to them. Nothing of this essentially turns on the condition that the cognitive system is a human brain—the principle is fully general.

So far this is all about processing information about *concrete* tigers; there are no other objects around. However, there seems to be a strange mismatch here between cognition and what is cognized. If the line of thought pursued in the previous chapters is by and large correct, then no object is ever seen, known, or talked about just for its own sake: it is given to us as a combination of properties

and relations, and none of the latter is essentially unique for this one object. For a designer of information storage this would not be the obvious choice: given a couple of objects, it makes sense to represent each of the relevant objects in a way that keeps them cleanly separate, and only afterwards to link all the relevant information to these representations. But to the human mind, objects are given by their qualities and their relations with other objects.

Surely this is moot. If anything *could* secure the specificity of my contact with the animal, i.e. that it is *that* tiger, then it is supposedly something much like Kripkean causal grounding (Kripke, 1972), where it is not the combination of properties that fixes the referent (i.e. not *after* the act of ‘baptism’), but the causal connections between the referent and the reference bearer. However—and this is a point that merits special attention—although the causal connections in question do determine uniquely, this information is not available to the brain’s computational activity for making a difference. In those areas of philosophy (in particular, in the philosophy of mind), where the focus is strongly on the *third-person perspective* (Dennett, 1991b; Chalmers, 2004), causal grounding appears to deliver everything that could be wished for in terms of unique reference. If I see a tiger, then, from the vantage point of someone observing me, one and only one specific animal is responsible for the activation of my tiger-image. *This* fact, however, cannot leave traces in my mind. Given what was observed in 2.4 with respect to haecceity, to me, from my first-person perspective, the tiger is given as a structured combination of qualitative properties and relations: *a* tiger is all I see².

In fact, there are not many physical objects in our perceptual experience that, from the way they are represented, have a claim on being discriminated up to uniqueness; and even fewer of them, if the representation is in thought or language. There is a fairly clear sense in which the way we connect cognitively to some individual tiger extends to all of a larger group of tigers—which, on the other hand, implies that it picks out none of them exclusively. What I hope to demonstrate is that, rather than being the problematic exception, *generality* is best regarded as the default mode of our cognition, and that it is knowledge about *individual* objects which stands in need of an explanation. I shall postpone this explanation for a little while; for the moment let us notice that, given the Central Assumption on page 90, the emergence of *the tiger*, as a *kind*, from acquaintance with singular tigers, is not at all unexpected. As Marcus (2001, p. 119) writes: ‘In many ways our mental representations for individuals are quite similar to our mental representations of kinds.’

Images are more often general than specific. There is a practical side to this as well: general images allows us to act sensibly when confronted with exemplifications of the general pattern. To classify concrete, perceptual objects, we

²See also Fodor’s (1994) discussion about the relation between semantic content and mental computation in general.

have images that give a generalized template for dealing with tables, tigers, and violins.

5.5 Abstraction vs Fusion

Kinds emerge from a plurality of concrete objects (*the tiger*, from individual tigers), and in that sense they resemble compound objects, which emerge from a plurality of parts. In both cases the emerging object is rooted in a plurality of already existing objects. But that is where the similarities end, both in the ways in which the underlying objects cooperate to make the new object arise, and in the features of the resulting object.

Let us look at these features first. Some important diagnostics often quoted with respect to abstract objects are:

- being non-spatial / non-temporal
- being non-perceptual
- being causally inefficacious.

At least for kinds, however, there seems to be only limited support for these assumptions in the draft ontology. It is very common to say things like:

65. In India the tiger is only seen in a few remnants of its original habitat today.
66. My favourite brand of toothpaste turned out no longer to exist.
67. Yesterday evening, somewhere near the City Park, I heard the cuckoo.
68. This (a violin is held up) is the key instrument of Western classical music!
69. The Plague caused the European population to be more than halved.

Even though individual tigers are evidently *more* strictly bound to place and time than *the tiger*; that the latter is completely unconnected to them is surely false. And as for perception: we can hear *the cuckoo* and see *the violin*; apparently we perceive a kind *through* perceiving one or more of its instantiation. Also, a kind is causally efficacious *through* its instantiations being causally efficacious.

This appears to accord well with the fact that, as pointed out in 5.1, in the draft ontology the border between concrete and abstract is not sharp. The examples above are illustrations of the easy conflation of kinds with their exemplifications as regards perceptual encounters, their whereabouts, and what they cause. As such they can be read as indicative of the same thing: rather than being somewhere behind the screen (like the Paris Standard Metre or an ideal Platonic Form), kinds are ‘in’ their exemplifications in the most literal sense possible—much like how compound objects exist ‘in’ their constituting parts. But

kinds are also the first step on a road leading in a completely different direction than fusion. One of the noticeable facts about abstraction is that it appears to come in grades. That kinds have a low degree of abstraction may seem like a very loose observation, but compared to *the number three*, *the pursuit of happiness*, and *the Schrödinger hamiltonian*, most of us will largely subscribe to it. Within the realm of abstractions, a biological species like *the tiger* seems to be a relatively down-to-earth type of object—which makes it all the more interesting to know in which way the emergence of kinds and fusions differ.

My proposal is that this difference has to be sought on the level of procedures, where it is based in the possibility of *parallel vs serial* scanning. Serial scanning, our mode of perceiving compound objects, means running through a group of objects (the parts), returning an outcome only after all of them have been scanned, which will typically be done by making use of the vicinitive relations connecting them. Every one of these objects will in principle contribute to the outcome. Similarities among the members of the group are helpful for the efficiency of the procedure (e.g. the beads in a necklace, see 3.5), but they are not necessary: some of the parts may be much unlike others (e.g. a board and a leg). Although it is never accomplished in practice, there is a sense in which the aim of a perceptual encounter with a compound object is to include all the parts. Consequently, the *properties* of compound objects in principle receive contributions from the properties of all the parts.

I shall refer to this process as *property composition*. Property composition is a remarkable effect of mereological fusion, where *new* properties in what is commonly called ‘the same domain’ are created out of those of the parts. Where beads can only be red or blue, necklaces can also be alternately-red-and-blue. The procedures for the parts turn up as sub-procedures of the procedure for the whole. Due to how property composition works, compound objects do not typically share properties with their parts; nevertheless, they by and large have properties in every domain (shape, colour, spatio-temporal positioning, and many others) that their parts have properties in. Since colour properties of tables supervene on those of their boards and legs, there is no definite need for ‘levelling out’ with increasing complexity. This nicely explains why mereologically compound objects are exactly as spatial, temporal, and perceptual as their fusing parts (whether they are causally efficacious is a moot point within the discussion about the reality of mereological composition—if they are, it is *through* their simples being so).

Parallel scanning, which, according to my proposal, is our mode of perceiving kinds, means approaching each of a group of objects by *the same* procedure. Figure 5.1 is a graphical impression of the differences. In parallel scanning there is no role for vicinitivity. As there is one procedure involved—scanning the same property or complex of properties—the members of the group must display a considerable amount of similarity. In this way many properties of the exemplifications are directly passed on to the kind they belong to; let us call this *property inheritance*. Thus, as a kind, *the violin* inherits many properties that individ-




Figure 5.1: Serial (a.) vs parallel (b.) scanning

ual violins generally have: *the violin* is eight-shaped, is made of wood, has four strings, and weighs around 400 grams. Compare this with the fusion of all violins, which, by property composition, is also made of wood, but has a very irregular shape and weighs many tons.

The exemplifications of a kind, however, usually differ from one another, and so inheritance is not unconditional. If we (as non-experts) see a tiger open its mouth and roar, we tend to conclude by default that this is how *the tiger* roars; but it is of course always possible that this specific tiger has a funny way of roaring, which should not be attributed to *the tiger*. Thus, in the course of learning, only properties that are to some degree consistently found in tigers, or are salient for them by other criteria, are considered properties of *the tiger*. Property inheritance seems to work like a filter that keeps certain properties intact but throws away lots of others: in particular spatial and temporal definiteness goes by the board, which is likely to be the main reason for the prejudices that kinds suffer from as being somehow ‘disembodied’. Here we find a clear contrast with compound objects, which have their properties (including spatial and temporal ones) in the same way as their parts have them (even though they are not the same properties).

The critical reader will have identified the above phrase ‘*salient for them by other criteria*’ as less than fully informative. For this I defer to those who, in linguistics, have done research on the semantics and pragmatics of *generics* (Carlson and Pelletier, 1995). Generics are sentences like:

70. The tiger is a good swimmer.

71. Tigers are good swimmers.

From an ontological viewpoint there are important differences between both statements: a nominalist could accept the second, but not the first. This should make it all the more remarkable that the ways in which their truth-conditions depend on facts about individual tigers show many similarities. What is known about the use of generics in natural language strongly suggests that the rules—to the extent that there are any—are at least as much driven by practical interests as they are a matter of metaphysical nomology, to put it mildly (e.g. Carlson and Pelletier, 1995; Nickel, 2008). Clean, logically expressible conditions for which properties are inherited and which are not, do not seem to be at hand.

Most of this appears to transfer to what the draft ontology attributes to *the tiger* as a kind. This amounts to what is bound to be quite a puzzling state of affairs in any realistic account of this object, but does not require much sophistication in image theory. The image of *the tiger* must be conceived as being the result of a psychological process, where it takes the shape that brings its bearer most cognitive profit. As a result, properties that not all of the concrete individuals share, or at least a significant majority, or at the very least a subgroup that has attracted special interest (*'the tiger is a man-eater'*), are, generally speaking, not considered properties of *the tiger*. Unlike what is the case with compound objects, there is not merely a levelling out, but even a *leaving* out of properties. Kinds need not have properties in every domain in which their exemplifications have them. Thus, if the swimming capacities of tigers generally are important enough to merit stable association with the kind, this trait will make it to the image for tigers generally. And, exactly for that reason, it will be part of the image of *the tiger*.

In this way properties that are not salient enough to belong to the kind leave 'gaps' in the property-spectre. Interestingly, these gaps are typically filled by others in the same domain. Kinds may lack the spatio-temporal positioning of concrete objects, yet *the Tasmanian wolf* was last seen in the wild in Mawbanna, in 1930. In many cases these 'compensating' properties have a *collective* character:

72. In the nineteenth century the polka was widespread in Europe and even made it to India.

73. The smartphone is used by more than five billion people.

Thus, we could distinguish *collective*, as opposed to *specific* property inheritance. At times, such properties collide with properties in the same domain that are inherited directly, to form what we could call *quasi-paradoxes*:

74. The malaria mosquito's adult life has a duration of one to two weeks. (specific)

75. In the last century the malaria mosquito made around 160 million deadly victims. (collective)

How could a creature that lives two weeks at most, pursue its murderous ways for a hundred years? The paradox is not real, of course: the facts stated in (74) and (75) are in no sense at odds with one another. The contradiction does not go beneath the verbal level.

A final important difference between compounds objects and kinds, and one that is directly related to the difference between serial and parallel processing, is that kinds are rooted in *all* of the concrete objects of a certain category. *The violin* is exemplified by all violins, with no exception. For compound objects anything like this is quite unusual: although it is possible to speak about, say, the totality of all bars of gold on earth, the typical compound object is not open-ended in this way. It is for the same reason that compound objects, unlike kinds, appear to allow for *exact duplicates*. The *gathering principle* (as we could call it) of composition is predominantly vicinity, which has the option to stop at some point. For kinds, the gathering principle is just matching. Every encounter that happens (or even *could* happen) anytime is part of the gathering. Hearing a cuckoo is hearing *the cuckoo*.

5.6 Attributes

Kinds are a distinctive class of abstract objects, and, as I hope to have shown in the foregoing, one that image theory has something to say about. But there are other types of abstract objects, and the natural step to take next would seem to be to consider *attributes*. Just like kinds are rooted in their exemplifications, attributes are rooted in their instantiations. A way to illustrate this is to consider perceptual relatedness. Seeing a tiger is at the same time seeing *the tiger*. Seeing something white is seeing *whiteness*. And if we want to teach children what the word ‘blue’ means, we typically point at blue objects and say: ‘*That* is blue, and *that* is.’

Like kinds, attributes also appear to emerge from parallel scanning, but while kinds stay as close to concrete objects as is possible while abstracting from spatio-temporal determination, attributes can move away from these at any distance, as long as they retain their discriminating function for the cognitive system. Consequently, the differences between kinds and attributes, *qua* abstract objects, are not just a matter of grade. Let us take *whiteness* as an example of an attribute again. This object is certainly more abstract than *the tiger*, but it is also another type of object. The tentative explanation that I want to put forward here is that attributes like *whiteness* can be seen as derived from conditions like *being white*. For *the tiger* there is also a variant of this type of abstraction: *being a tiger*. What is the difference between both the latter? The question is pressing, since the conceptual content seems to be very much the same.

More or less in the spirit of Frege's (1892b) terminology, *being a tiger*, unlike *the tiger*, is 'unsaturated': it has a 'hole' that must be filled by another object. *Being a tiger* is, one could say, the *relational* version of *the tiger*: it is not just there on its own: there is something that has this condition. As said, *being a tiger* is not the norm: typical attributes are far more abstract: *being white*, *being fast*, and *being dangerous*. Now, what sets attributes apart from kinds, or so my hypothesis goes, is the *lack of independence* of their procedures: it takes more effort, cognitively, to dissociate them from their instances.

Newly created images are likely to be based on images already present in the cognitive system, functioning as a template for them. Given this, to go from a concrete tiger to *the tiger* only requires dropping a couple of properties and relations which are variable and unstable anyway: space, time, movement and other highly accidental circumstances. But to forge an image for *being white* out of that for some white dog, or *being fast* or *being dangerous* out of a fast and dangerous tiger, is probably a more challenging reorganization of images, given how intricately sub-procedure for such conditions must be interwoven throughout the main procedure.

There is some reason to think that, in connection with this, the emergence of abstract attributes out of the concrete objects that instantiate them is more indirect than that of kinds: it already happens at the concrete level. As noted in 3.1, the draft ontology acknowledges *tropes*: a type of attribute that is particular rather than general, e.g. Snow White's beauty. The latter is not just being beautiful, it is *this very person's* being beautiful. The fact that such intuitive tropes of an object come apart the way they do is just what must be expected according to the Central Assumption, for complex procedures not only have sub-procedures taking care of the mereological parts of their target objects, but also sub-procedures running through all the parts, collecting only a specific portion of the available information—e.g. the centipede's being even-segmented. All of them give rise to predications, and this explains the emergence, not only of parts, but also of 'aspects' of an object. From here, the remaining step in the establishment of *abstract* attributes is an easy one. The relation between *being beautiful* and Snow White's particular being beautiful is exactly the same as that between *the tiger* and particular tigers (Moltmann, 2004). (It looks like we have found a less typical case of rootedness here. Tropes like *Snow White's beauty* (and her other properties) are not single objects rooted in a plurality, but a plurality of objects rooted in a single one. For there must be a beautiful object or person before the latter's *being beautiful* can be there, and a tiger, before its *being a tiger*.)

The explanation of attributes in image theory proposed here is quite helpful for dealing with the famous conundrum about the truth conditions of statements like:

76. Red is more like orange than it is like blue.

It is well-appreciated that its truth does *not* boil down to red things being more

like orange than like blue ones (Armstrong, 1989). Wittgenstein would probably have called such sentences ‘grammatical propositions’ (inferential rules disguised as descriptions of reality), which view comes close to explanations in terms of trivial entailments mentioned above. In a transcendentalist understanding of colours, however, the explanation is surprisingly straightforward. The colour *blue* to us is an object stripped of all of the properties ‘normal’ objects have, whereby only colour is left³. Now, if we would make a comparison between, say, two violins, to decide which of them is more valuable, the outcome would be the result of weighing many properties of the violins—in fact any property both relevant and accessible. Comparing *red*, *orange*, and *blue* is done, analogously, by comparing their properties—but that is basically *only* colour, so none of the accidental properties of concrete red, orange, and blue objects interfere.

5.7 In Search of Individuality

Above it has been argued that, rather than being the problematic case, abstract objects are *easier* to understand than concrete ones. Most people have a clear concept of ‘tiger’, but for all but a few there is no mental representation of any *individual* tiger. And even though we can imagine a tiger walking through our back garden, there is not normally a way to make this mental construct refer to any specific animal. We can also perceive a tiger—see it, or hear, or smell it—but then again, although there *is* a fact of the matter as to which one it is that we have before us, its individuality, as was argued in 5.4, does not reach our mental machinery. We observe that *there is a tiger* (Aguilar-Guevara and Zwarts, 2010). We do not observe that it is, say, Shakti, who just escaped from the zoo, unless, again, we belong to that small circle of people who enjoy acquaintance with individual tigers, like Shakti’s keeper. And this makes it all the more intriguing how the latter does it.

Individuals are remarkable objects anyway. In a standard semantics for first-order logic they are postulated as given. The common understanding of a phrase like ‘There is a tiger’ is that, when the domain of individuals is run through entirely, then on at least one occasion the individual in question will be found to be a tiger. Surely this is what must be the case if this conclusion is drawn after a sound perceptual process; however, if perception has not yet taken place, e.g. if it is only *feared* that there is a tiger—not seen, heard, or smelled—then this account fails. The thought ‘there is a tiger’ is surely activated, but before one of these creatures has been actually encountered, it is not anchored in any individual in any domain. How the general concept assumes individuality is by

³The idea that there are such images appears to find some support by what is known about the brain: rather than being scattered all around our visual representations of objects (as might be), different aspects of those objects (colour, shape, movement, and so on) are to some extent processed by specialized regions in the brain (Zeki, 1992; Cavina-Pratesi et al., 2010).

no means self-evident, even though, in an exposition like the present one, where objects are interpreted transcendently, it cries out for an explanation⁴

Where do we find individuals? Their natural habitat appears to be what I shall call *the concrete level*. This is an impressionistic term: the draft ontology appears to ‘believe in’ such a level, which is inhabited by what Aristotle called primary substances: those objects that cannot in any way be thought of as predicated of others. That is at least something, but maybe we can make further observations. We all know individual people; and also individual tables and chairs, towns and highways. What makes them that? What makes Shakti’s keeper so sure that *this* is Shakti? Probably he knows some characteristics of Shakti he is positively sure that no other tiger in the world has. But that is not enough: he might be mistaken and there could be a second tiger sharing these traits: that would not in the least make it the same tiger. What does it *mean* that the latter is another tiger?

Of central importance seems to be the fact that, on the concrete level, there is *distinction beyond dissimilarity*. The intuition that two individuals could share all their qualitative properties—that an individual has an identity that goes beyond mere properties—is quite strong, *but only on the concrete level*. This is more remarkable than is commonly recognized. In 2.4 I have made some comments about the feasibility of perfect duplicates of abstract objects, to conclude that there is nothing absurd about the idea; yet, such options are firmly rejected by the draft ontology. If two kinds have all their properties in common, then they are one kind, likewise for conditions, liquids, colours, numbers, ideologies, and so on. On the abstract level similarity works like a gravitational force: objects for which there cannot be anything standing in between them conceptually, fall together.

Now, an explanatory scheme in this spirit—the identity of indiscernibles, supplemented with immediate identity as discussed in 2.4—of course settles matters quite comfortably for kinds, attributes, and other abstract objects, but it breaks down on the concrete level. It would be interesting to know why. As for explanatory resources we have, at the moment, two important principles at our disposal. The Central Assumption puts objects on stage (emergence), while the identity of indiscernibles puts a limit to their number (identity). Could they not be the alpha and omega of the transcendentalist understanding of objects? Although it seems tempting to introduce an additional principle to provide for something like a ‘deeper’ ground for the identity of individual objects, fitting individuality in the same template would make for a more parsimonious explanation.

Fortunately, the circumstances as they prevail on the concrete level might be exactly right for this purpose. As we saw before, spatio-temporal relations are special in several ways, one of which is that they are usually thought to be

⁴See also Huang and Ahrens, 2003, who, based on findings in Mandarin, challenge the idea that individuals are semantically basic.

external: there are no underlying monadic properties on which they supervene. Or so it seems; much effort from the part of metaphysicians has been spent on speculating about the ultimate nature of space and time, one of the main issues being whether space and time are substantial or merely relational (Hooker, 1971; Newman, 1989). However, making sense of this condition in image theory may be relatively easy. Even though human beings experience objects to be ‘here’ or ‘there’, ‘in rest’ or ‘in motion’; even though there is a strong absolute perception of ‘up’ and ‘down’, and even though ‘now’ feels like a very distinctive moment amidst the past and the future, we are well able to appreciate that ‘there’ can be ‘here’ for another observer, and that ‘up’ can be ‘down’ elsewhere on the globe. Real symmetry appears to be feasible with respect to the condition of ‘being at distance D ’ (say $D(x, y)$).

But if this is so, then relations like D represent the unlikely combination of *symmetry* and *irreflexivity*. Thanks to this, they provide a ground for distinguishing x and y (objects at some distance cannot be the same⁵) *without* making them qualitatively different. Only spatial relations seem to have this character. Temporal relations are strongly asymmetric, and if Hume is right and the only remaining class of external relations is causal connections—which are also asymmetric—then these also distinguish their relata qualitatively. And so do intrinsic relations: unlike $D(x, y)$, every intrinsic relation $R(x, y)$ supervenes on monadic properties $P_1(x)$ and $P_2(y)$. For R to be irreflexive, P_1 and P_2 must be mutually exclusive, or else there could be some a such that $P_1(a) \wedge P_2(a)$, i.e. $R(a, a)$; but then x and y will be qualitatively distinguishable by P_1 and P_2 .

The upshot is that, if two things are at some distance from one another, then they are *thereby* distinguished, irrespective of how similar they may be qualitatively. Black’s identical spheres (Black, 1952) were intended exactly to illustrate this possibility. It is this property of space that opens the possibility of a predication that declares its substance a plurality of qualitatively identical objects. Such objects are, as the term of art goes, *weakly discernible* (Saunders, 2006), a condition which is often quoted in relation to the pluralities of indistinguishable particles as they are known in quantum mechanics. Clearly, moderate-sized individuals are never *really* indistinguishable: in anything like a real-life situation individual tigers are conceptually more than just what emerges from a repetition of the tiger-procedure. Still, from a transcendentalist perspective on objects the mere option to conceive of pluralities of objects without the need to distinguish them qualitatively is most relevant. My hypothesis, therefore, is that this feature of space—its power to bring about a distinction without a difference—is all that is needed to explain why we feel the concrete level of objects to have its very particular character.

⁵In a suitably curved space this would not have to be so, however (Swinburne, 1995). Although this point does not seem to have too much relevance for the present argument, the question why space is experienced as having a Euclidean structure is surely an interesting one.

If this is correct, then individuality, rather than being a pivotal ontological condition, is just the consequence of some ‘oddity’ of one class of predications, viz. spatial distances. This is of course theoretically profitable: no ‘primitive individuality’ has to be invoked, nor is there any need to contrast some conception of primitive identity with identity due to the identity of indiscernibles. But if so: if the special ontological position of concrete objects, is (to exaggerate a bit) just a remarkable by-product of their having spatial properties, then they should be suspected otherwise not to differ very much from other objects. Individuality could well be a relative condition: why could individual objects not have a general side to them too? Maybe being an individual is less special than it is normally thought to be, and so is being concrete.

This suspicion is justified, or so I believe. As argued above, real-life individuals are not like the ones we study in logic. Considering the easy-going understanding of such items there and the messy reality of the inhabitants of the real world on which they must somehow have been modelled, it is hard to miss the abysmal discrepancy between both. Theoretical individuals are primitive, either have a property or lack it, do not overlap or include one another, and, most significant of all, are *atemporal*. Given x and y they are either the same or not—the latter meaning: not at all. They conform to what I have called ‘objectual autonomy’. And that is it. Real-life individuals can overlap and be part of one another, but what complicates things most by far is that they exist in time, are subject to change, and can cease to exist.

Now, I would suggest that it is precisely this, the existence of objects *in time*, which represents their general side. Sarah one moment ago, and Sarah now, as well as Sarah when she will, in a moment, re-enter the room, clearly seem to be subject to the same mental image. That all are accepted as being the same person—i.e. the same object, enduring in time—is just what the Central Assumption predicts. Just like *the tiger* is rooted in individual tigers, Sarah is rooted in her temporal stages. The perception of objects as *enduring* in time (rather than *perduring*, cf. Johnston and Forbes, 1987) can be interpreted as a consequence of the fact that the scanning is clearly parallel, not serial.

I must stress that these comments are very sketchy and do not in any way add up to a solid account of how time must be understood within image theory. Essential aspects of endurance have not been touched yet, e.g. that *temporal contiguity* is extremely important in the identity of individuals over time (even an extremely identical twin-sister of Sarah will not blend with Sarah herself). This, however, does not affect the moral that individuals are, for the most part, a class of objects like all others: only one among the many grades of abstraction.

5.8 Deeper Abstraction

In 5.5 I have proposed that the emergence of objects like the attribute *whiteness* be associated with the procedure, involved as a sub-procedure in many more complex procedures, of detecting whether something is white. It is tempting to consider the prospects of even ‘deeper’ abstractions like, say, numbers. Could *the number three* be explained along similar lines? I believe it could.

Here is the apparent problem with the idea that the emergence of *the number three* parallels that of *whiteness*: whereas many things are white, nothing, at least at first blush, ‘is’ three. But first impressions can be deceptive. Pluralities are objects, and many of them are triples, which comes pretty close to being three. It is no coincidence that the way to teach little children what numbers are is to show them numbers of beads, blocks, and ducks. The main obstacle for a predicate like this to go proxy for a genuine number seems to be that the ways to divide up some couple of things is mostly arbitrary. To find ‘threeness’ in a gathering of three persons at a bus stop may seem trivial, but from a less prejudiced point of view one could as well have counted cells, molecules, or genders. Cf. Frege (1980):

If I give someone a stone with the words: Find the weight of this, I have given him precisely the object he is to investigate. But if I place a pile of playing cards in his hands with the words: Find the Number of these, this does not tell him whether I wish to know the number of cards, or of complete packs of cards, or even say of points in the game of skat. (...) I must add some further word—cards, or packs, or points. (...)

It marks, therefore, an important difference between colour and Number, that a colour such as blue belongs to a surface independently of any choice of ours. (...) The number 1, on the other hand, or 100 or any other Number, cannot be said to belong to the pile of playing cards in its own right, but at most to belong to it in view of the way we have chosen to regard it; and even then not in such a way that we can simply assign the Number to it as a predicate. What we choose to call a complete pack is obviously an arbitrary decision, in which the pile of playing cards has no say. (p. 28-29)

This decision, identifying cards, packs, or points—does a lot of work. Every such decision is itself a procedure, and the properties that it defines are important with respect to the idea of *sortal predicates* as well⁶. The existence of such procedures

⁶Counting does not require ‘ideal’ sortal concepts, however. Also the role of similarity, in particular of strong, visual similarity, in setting the stage for counting is easily overrated; *to be an egg* is a good example of a sortal predicate, but *to be a thing on the table* is an exceptionally poor one. Yet, in typical cases things on the table are just as easily counted as eggs—we intuitively pick an *ad hoc* criterion for ‘thing’ (often in striking agreement) and count.

seem to be entirely within the assumptions of image theory so far.

But counting *itself* is a procedure as well, and to find out about the transcendental origin of numbers, we should indeed study the nature of counting, or so I would suggest. The only difference between a procedure like counting and most of the procedures we have considered so far is that counting stands at a certain distance from perception. To count, one must first have a plurality (cards, packs, or points). Now, to keep things manageable ontologically, the pardonable inclination is to conceive of the plurality as being a couple of mereologically distinct objects, all of them instantiating some sortal concept. In the transcendentalist approach defended here, however, this is a non-starter for several reasons. First, mereological relations are in no way privileged over other purely ontic relations, and even if they were: counting is not restricted to non-overlapping objects (cf. Figure 2.1 again). Counting is not even restricted to objects to which mereological relations can be sensibly applied at all: apart from tomatoes, persons, and squares, we routinely count colours, ideas, causes, events, relations, and so on.

Thus, counting needs two things: a procedure that picks units, plus another that does the counting. No matter the enormous variety with respect to the former, the latter *is always the same*. Even though it is fully accepted that there can be issues about how one should count and whether there is a ‘natural’ or ‘obvious’ way to count, none of these issues are ever felt to undermine the idea of counting itself, or even shaking it. There just is a procedure of counting, and all it needs to be put to work is some convention as to how to cut the pie. From this, our human notion of natural numbers apparently emerges. This need not come as a surprise: as long as there is an image in our brain running a procedure with a single outcome, the Central Assumption predicts that some object will emerge, and the case of *the number three* illustrates that this also holds for procedures playing an auxiliary role in other procedures.

Procedures at some distance to immediate perception may also provide a clue as to the explanation of *facts* in image theory. Facts could strike us as being abstract, but they are not necessarily *general*. Sometimes they are (‘All ravens are black’), sometimes not (‘This raven has a round white spot of 0.35 inch below its right eye’). In fact they are typically *more* specific than, say, concrete physical objects. My motivation for discussing them under the heading ‘abstraction’ is that, e.g. the fact that Socrates is bald is only one aspect of the object Socrates; there is a lot more to tell about this man. The specific nature of the fact *that Socrates is bald* is demonstrated by the following comparison: even though this fact shows some similarity with *Socrates’s baldness* (which, like *Snow White’s beauty* is a trope), there is a subtle difference between both. Socrates’s baldness may have properties that are open to enquiry. It may be partial, repulsive, charming, authoritative, and what not. But the fact *that Socrates is bald* is just that. Facts, unlike other objects, do not have a great variety of modes of presentation, and they are informationally transparent to a high degree. The reason appears to be that the content of the fact is neither Socrates as such, nor

baldness, but only the *link* between both.

By the interpretation I wish to advocate, what sets facts apart from other types of objects is their being particularly compact information bearers, which makes them especially suitable for *being communicated*. Most objects can be perceived, and by perceiving them the cognitive system gathers information about them. Perception can be extended: we can look again, listen again, smell, measure other variables, thereby discovering more and more aspects of what is still the same object. Because of this, it is typically hard or even impossible to communicate the object *per se* by purely verbal means. Hence, the mildly paradoxical situation: when we look around we see objects everywhere, but when we speak, what we exchange is facts, no objects. *Language*, however, appears to be made for this purpose: we communicate *about* objects, by exchanging facts. They are the ‘digital’ medium for speaking and writing about the analogous world.

Facts can be thought of as emerging due to the working of a class of procedures operating on a meta-level relative to other ones. There are, to my mind, independent reasons to assume the existence of cognitive modules that detect the collaboration between procedures for objects and properties: they would cover the type of mental activity in which a *reflexive* element is involved. In that sense, facts very much appear to be a human ‘invention’. To see that, e.g. a teacup is blue, is one thing—many animals can do that as well. However, to appreciate this fact in abstraction of the other aspects of the teacup, or the special shade of blue—let alone to communicate it—requires a level of thought that at best only the smarter species master. All this would suggest a strong link between facts and language. A world without facts seems feasible: dinosaurs, mammoths and even primitive humans would be happy enough with only objects. But once there was language, it seems, there had to be facts. And once there were facts, due to their informational economy, not only our communication improved, but also our memory resources could be used far more efficiently. At the risk of speculation: the emergence of facts, besides trees, mountains, sticks, and stones, is likely to have marked a decisive moment in human evolution.

5.9 Perfection

With the study of deeper abstractions we venture further into the realm of the Platonic, a realm the inhabitants of which not only stand out by the extent to which they are abstract, but also by their putative *perfection*. *The circle* is more circular than any concrete circle. The observation that abstraction goes hand in hand with perfection is not confined to numbers and geometrical structures. Statements about abstract objects often seem to be used to express *stereotypes* or *norms*, rather than facts. Compare:

The American soldier is disciplined, physically and mentally tough,

trained and proficient in his warrior tasks and drills⁷.

with:

American soldiers are disciplined, physically and mentally tough, trained and proficient in their warrior tasks and drills.

Even though these are sentences of much the same content, the first, more than the second, describes how things *ought* to be.

It is worth noticing that, although the association between abstraction and perfection is an extremely familiar one, it is not evident or straightforward that this is how it has to be. Any realistic account of abstraction will have to make the abstract circle in some way or other arise from concrete circles, none of whom are perfect, and it is hard to see how the former—which, being perfect, by no means represents the average!—could spring from such a well. The perfection of this and other Platonic forms, however, is a phenomenon that image theory is in an excellent position to explain. A procedure to test circularity of objects can be a very *simple* procedure, since perfect circularity is conceptually simpler than any specific kind of actual, and therefore *imperfect* circularity. Hence, an image representing the perfect circle can be smaller and consume less memory.

An important observation in this regard is that the perfection of abstract objects is not perfection in any moral or economical sense. It is *conceptual* perfection: the perfect counterfeit, perfect fool, or perfect totalitarian state are all perfect because of the extent to which they satisfy a certain (simple) idea, not because they are good or useful. Meanwhile this aspect of Platonic objects opens a new perspective as well. Since no concrete circle is perfect, it would seem that the perfect circle does not in fact represent any everyday circle whatsoever. No table board, vinyl record, village green, or even lunar disc is ever perfectly circular, so what does Plato's circle have to do with them at all? Image theory should not have much trouble with the general phenomenon of incomplete match between image and object, but the metaphysical consequences of this possibility are formidable. The abstract object itself gains some autonomy with respect to its worldly counterparts. The gap that is opening between the perfect circle and concrete circles makes the former more abstract, but also *less real* (cf. Turner, 2010). This will turn out to be a valuable observation, the consequences of which will (among other things) be considered in the next chapter.

⁷Paraphrases of the *U.S. Soldiers Creed*.

Chapter 6

Being and Non-Being

This chapter is devoted to some of the remaining issues that image theory has something to say about, with a central role for *being* itself. And for its absence: transcendentalism is in quite a favourable position to deliver an explanation for the state of non-existence that so many entities find themselves in. First, however, the relation of *identity* will be discussed, a relation that, being far from trivial, it takes some effort to give its proper place in image theory. Hereafter the concept of *rootedness*, which so far has been little more than an intuition, will be worked out further. It will also be proposed as an alternative to ‘relative identity’. The conditions of being real, being unreal, and being halfway will be addressed, guided by ideas about *existence* and *fiction*. The chapter closes with some thoughts on *imagination* and *modality*.

6.1 Identity

As we have seen in 2.4 and 5.7, the concept of *identity* holds specific challenges for image theory. This may seem strange for whoever holds Lewis’s view that identity is ‘utterly simple and unproblematic’:

Everything is identical to itself; nothing is ever identical to anything else except itself. There never is any problem about what makes something identical to itself; nothing can ever fail to be. And there is never any problem about what makes two things identical; two things never can be identical. (Lewis, 1986a, p 192–193)

Lewis himself cannot have missed the circularity in his reasoning. To explain the meaning of terms like ‘itself’, ‘except’, and ‘two’, one must make use of identity or its negation. All the same, his standpoint—from a *realistic* point of view—is lucid and sensible. How could a concept as immaterial as identity be anything to

worry about? For it to be operational, nothing particular has to the case in the world: there only have to be things.

This representation of the matter, however, is spectacularly at odds with philosophical experience. ‘Two things can never be identical.’ Really? Suppose I am on a quest, hoping to find two things: a white horse and a good comrade—is that not looking for *two* things, even though they might turn out to be the same? For reasons like this, few metaphysicians have left this problem completely untouched. There only have to be things, but with them come our conceptualizations. Even if these, rather than the things themselves, are the troublemakers, to stop conceptualizing does not look like a feasible solution; and if, as the transcendentalist understanding goes, conceptualizing is more than passive reception, then the philosophical preoccupation with identity is just what should have been expected.

Two—quite dissimilar—principles have so far (cf. 2.4 and 5.7) been identified which are likely to be key elements: *immediate identity*: which applies when we see that some round object is the same as some red object (say, a billiard ball) and the principle of the *identity of indiscernibles*, which can be used to infer *mediate* identity (identity which is not perceived or thought directly, but concluded from other sources of information). In logic, immediate identity is mostly expressed by structures like $P(x) \wedge Q(x)$; mediate identity, by $x = y$. Immediate identity looks fairly inevitable: properties come together in what we experience as being objects: the redness of the ball and its roundness are inseparable; but of the identity of indiscernibles this is less obvious. Whether, and to what extent, the principle holds is an open question, and *if* it holds, it might be wondered if it is more than a tautology. Let us consider this point first.

As for formal matters: first-order logic comes in two varieties: with or without identity (equality), which means that (mediate) identity can be viewed as an extra, an addition to the core of the logic¹. What does the addition add? Let us consider first-order theories of a fully *descriptive* nature, i.e. there are no individual constants (names). The idea behind such theories is that they tell us what is to be found in the world, but not beyond descriptions in terms of qualities. Leaving out identity will deprive these theories from the ability to express the existence of *exact duplicates*; but from a transcendentalist viewpoint, the existence of exact duplicates (*really* exact: not just locally, but also relationally) is just as epiphenomenal as that of haecceistic properties (cf. 2.4).

With this in place, it would seem that identity is still needed for expressions like:

$$77. \forall x, y : \phi(x, y) \rightarrow x = y$$

as in:

¹Of course many authors take identity to be a core logical notion itself (e.g. Quine, 1986, Ch. 5)

78. If George is your only brother and the mayor is your mother's only son, then George must be the mayor.

Still, this is not a completely convincing argument for the indispensability of identity. Surely a term like 'only' *presumes* identity, so (78) does not support a valid conclusion to the effect that mediate identity is essential. If the theory is really descriptive, then the effect of expressions of the type of (77), can be mimicked by the following rule inspired by Leibniz's Law:

79. $\forall x, y : \phi(x, y) \rightarrow \forall \psi : \psi(x) \equiv \psi(y)$

(even though this would introduce a second-order element).

Does this formal trick correspond to anything interesting ontologically, i.e. could we leave out mediate identity—over and above mere indistinguishability—altogether? This is another one of those thought experiments in which the world is bereft of, if not an essential aspect, then at least one we are extremely attached to, like compound objects, whiteness, the number three, and so on. In a world without mediate identity there would only be encounters with certain combinations of qualities, but without the very notion that what I see on one occasion is 'the same' as what I see on another. There would still be a teacup in front of me, and reason to expect that, barring some interfering action, some teacup would be there one minute later as well. Can it be trusted to hold hot tea like the old one and not, say, melt or crumble? It would seem so: many previous teacups have proved up to their task. Induction does not require identity: why do I fear a wasp I have never seen before? And what about the observer, what about *myself*? Without identity, there would be thoughts, feelings, and perceptions at different moments in time very much like those that there are now. A Buddhist would not be shocked.

Scepsis about mediate identity is not the least reasonable kind of scepsis, and yet there are, I think, good reasons to resist it. Exact duplicates, after all, are no strangers in the draft ontology. The concept is by no means confused or counterintuitive. And then again, there *is* such a thing as immediate identity: it seems strange that we should take identity seriously in one case and adopt something like a deflationist stance in the other. That identity is, at the end of the day, a transcendental notion, one that belongs to the concept of object, and thereby to the human cognitive point of view; *and* that indiscernibility is the gravitational force, does not mean that indiscernibility *constitutes* identity. But if it does not, something else must.

It may be time to look at procedures once more. Not coincidentally, the issue is reminiscent of something every software developer becomes familiar with early in their career. Some variables referring to objects (=items as they are stored in memory) are assigned *by value*. The assignment can take the shape of equation with another variable. This is what the code looks like (the instruction ':= ' indicates assignment) when x is declared to be 5, and thereafter y is equated to x :

```
integer x := 5
integer y := x
```

Result: $y = 5$. New memory space is allocated for y and the value 5 is copied into it. The effect of this is that, when the value of x is changed later on:

```
x := 6
```

it is still the case that $y = 5$. This is how it commonly goes with relatively ‘small’ objects: strings, numbers, bits, and the like. But for bigger objects (consuming more memory space) this is not the default mode. Variables for bigger objects are (in most languages) assigned *by reference*². No copy is made, only a link to an already existing memory location. Let the type ‘record’ be such a bigger object. Initializing records requires some information, e.g. the first name and surname of a person:

```
record x := new record("John", "Brown")
record y := x
```

The result is that y is, just like x , a record the surname of which is ‘Brown’, just as expected. But now suppose the input for x was wrong and had to be ‘Braun’. Hence, an adjustment is made:

```
x.setSurname("Braun")
```

This time, when x changes, y changes with it. When printed, the surname of y is also ‘Braun’, not ‘Brown’.

Applied to image theory, this would inspire the following hypothesis. If x equals y only in terms of properties (‘by value’), then both have their own representation in memory: their own image; whereas if x equals y in that they are the same object (‘by reference’), there is just one image. This has practical consequences. As long as x and y are just similar, the way they are treated by the cognitive system allows them to undergo changes *independently*. Let there be two tigers, Shakti and Shanti. I may not be aware of any qualitative difference between them; yet, if Shakti turns out to have a miserable character, I will not therefore conclude that Shanti is mean, too. If Shanti breaks her leg, nothing prompts me to expect that Shakti limps. But if I should find out that Shakti *is* Shanti, then every single detail about Shakti becomes handled by the same image dealing with Shanti. Were this not so, then updating the memory content would have to be done by copying every new piece of information from one image to the other, which would be inefficient and error-prone.

This—rather uncomplicated and not entirely surprising—idea appears to give a criterion for when to consider objects to be the same, rather than just indistinguishable. Notice that the relation between objects and images is not one-to-one:

²Things are in fact a little more complicated, but that is inessential for the example.

a great deal of our cognitive handling of Shakti and Shanti will be done by the same image anyway (that of *the tiger*), but as long as we have ways to tell them apart, the images will remain distinct, if overlapping. Once there is enough reason to believe that they will be similar in every possible respect, their representations ‘collapse’ into one image.

That representations may so collapse is not, according to this picture, a matter of necessity. It is a decision by the cognitive system. The decision finds its justification in that no independent changes are to be expected anymore, and for concrete objects the trigger appears to be that objects occupy the same position in space synchronously (cf. 5.7). This would nicely explain the fact that for abstract objects there is more room for deliberation about their being the same or not. In the abstract domain, by formulating *identity criteria* we can almost ‘create’ new objects. Sets differ from ordered pairs, and *the tiger* is not the same as *the Sumatran tiger*, which in turn differs from *the Gunung Leuser tiger*³.

This, then, is the meaning of the identity of indiscernibles in image theory, or so I would suggest. To be the same object implies to occupy the same position in the conceptual network as described in 2.4, which means that there can be no deeper *ground* for identity than *perfect* (local and relational) indiscernibility. Transcendentally, however, the difference between either approaching things with one image, or with a couple of functionally equivalent ones, appears to mark the distinction between identity and mere indistinguishability. Meanwhile, it should be noticed that the failure to *actually distinguish* objects, at this occasion or that, is neither necessary nor sufficient to conclude that they are the same thing—it would be quite absurd if it were. If Shakti always walks by in the morning, and Shanti in the evening, we may, short of being able to tell them apart, yet suspend our judgment as to whether they are the same tiger. On the other hand, learning, by hearing or reading about it, that Scott is the author of *Waverley*, involves no similarity (let alone indiscernibility) between both these modes of presentation. Nevertheless, here also: what makes the ‘collapse’ into one image—i.e. full identification—worthwhile is that no more differences are anticipated between Sir Walter Scott and the author of one of the first historical novels. If the first one is blonde with a fleshy face, we trust the second is.

The latter case reveals a complication with respect to mediate identity. It is an example of the sort of knowledge that we do not acquire by perception, but by e.g. reading a book, and when the information strikes us as being sufficiently credible, some *rearrangement* must take place that will result in a single image: Scott, the author of *Waverley*. Now, merging both modes of presentation into one image sounds easy, but it remains to be seen if—and how—any straightforward meaning can be attached to this condition. To put the point in more general terms: some things are *P* and some are *Q*, and others still are both *P* and *Q*. Let us say that we have *x*, *y*, and *z*, with *P(x)*, *Q(y)*, *P(z)*, and *Q(z)*. But how

³Gunung Leuser is a National Park in Sumatra

is a cognitive system to know about z being *one* object having *both* properties? So far we have only provided for ways to know that $\exists x : P(x)$ and $\exists x : Q(x)$. Let there be a procedure that scans objects for being P and one that scans objects for being Q : how can these be combined to constitute a procedure that scans for being $P \cdot Q$ ⁴?

I hope the urgency of this problem is brought out clear enough. So far (from 2.4 onwards) I have tacitly taken this possibility of conjunctive procedures for granted or deferred to identity as some sort of primitive; but wholehearted embracement of the transcendentalist view does not leave much room for such manoeuvres. If images and procedures are the *explanans* of the world as it appears before us, then they should be put to work here as well. Notice that a cognitive system as it has been envisaged in the above is essentially like a person groping in the dark, bumping into several things with several properties. Worse: if one of those things has multiple properties, it is *mere* properties the system encounters, without any obvious way to appreciate their coinstantiation. What this entails is reminiscent of the identity of indiscernibles, but a lot less acceptable: where the latter says that a world with identical x and y is as good as a world with only one of them, here a world with only x and y (with $P(x)$ and $Q(y)$) is as good as a world with only z (with $P \cdot Q(z)$). This simply cannot be true!

How to explain, on the image level, that a world in which both P and Q apply to *one* object looks different indeed? We might speculate about some—hitherto unknown—ability to add identity directly. The procedures for P and Q are both at work, and some special faculty informs us that they happen to scan the same object: from $\exists x : P(x)$ plus $\exists x : Q(x)$ to $\exists x : P(x) \wedge Q(x)$. Notice that the unification cannot be brought about *after* P and Q have been detected to apply, for no valid *inference* could take us there. Whatever secures conjunction must happen already during the scanning process. As far as I can see, no computational principle to accomplish this suggests itself.

The solution I want to put forward here consists in the hypothesis that, in the cognitive system, the conjunctive property $P \cdot Q$ is always the *consequence* of some *single* property, call it C . The procedures for P and Q are not ‘tied up’ so that they do the scanning together, but get activated higher up, thanks to the inferences $C(x) \Rightarrow P(x)$ and $C(x) \Rightarrow Q(x)$, and therefore, $C(x) \Rightarrow P \cdot Q(x)$. A red square is not just a combination of redness and squareness: it is a specific visual pattern, that may take different size, position and angle of rotation. The idea is that this complex condition is taken in as a whole *before* the properties *red* and *square* are abstracted. Also, if an ornament in a room is described as ‘*Art Deco*’, there must be a concrete pattern of shapes and colours giving rise to the use of this term. The same pattern can be a ground for the conclusion that the ornament is (artistically) ‘abstract’, and so we have a conjunction between *being Art Deco* and *being abstract*. But this suggests that knowledge of such

⁴I.e. being both P and Q , cf. 3.5.

conjunctions in the *absence* of knowledge of the underlying C must result in a psychological activity to ‘reverse-engineer’ some such underlying condition. This may sound strange, but a good case can be made that it is what we routinely do. When we hear about an ‘abstract Art Deco ornament’ without actually seeing it (e.g. by reading this description in a book), we do not usually just store the predicates *abstract* and *Art Deco* in our memory, but adopt some ‘working hypothesis’ filling in the unknown details, including, however vague, colour (soft hues of yellow, brown and green) and repeating patterns with curved lines.

Immediate identity, according to this line of reasoning, is *analytical*: first there is an object, known by a single property C (whose procedure can be quite complex), and only then certain other properties (P and Q) are found to apply to the object jointly. Mediate identity is *synthetic*: certain properties come together to apply to one object. But to make them apply jointly, an underlying condition C must be retrieved.

6.2 Roots

On a number of occasions in the above I have raised the prospect of an explanation of the relation of *rootedness*. This is not too early: the concept, however important in the overall argument, has been left somewhat in a state of inexplicitness so far. Of course it shows, as remarked before, many similarities with another purely ontic relation, to wit *grounding*. But rootedness is intended as a concept belonging to image theory, while grounding appears to have a more general applicability.

Grounding is a term of art—and yet it is not: the idea that some things ‘come first’ ontologically is probably as old as philosophy itself, and certainly not counterintuitive. In recent thought, the preferred view by many theorists is that grounding is a relation between *facts* (Correia, 2005; Rosen, 2010; Audi, 2012; but cf. also Turner, 2016). To me it seems that this is a matter of convenience: according to the analysis in 5.8, facts are informationally transparent, which makes them a useful vehicle for expressing all sorts of relations between all sorts of objects. But if fact f is grounded in fact g , then the objects figuring *in* f are related to those in g as well—often in purely ontic ways—so a grounding relation between objects generally seems quite acceptable.

Despite superficial similarities with entailment, grounding, if taken as a relation between facts, is not inferential itself. Inferences can go round in circles, but a ground should be ‘below’ what is grounded, and also come ‘before’ it, or so the intuition insists—convincingly, however naive the picturing. In Schaffer’s (2010) words:

There is (...) the metaphysical structure of prior and posterior, reflecting what depends on what, and revealing what are the fundamental independent entities that serve as the ground of being.

Consider Socrates. Given that he exists, the proposition ⟨Socrates exists⟩ must be true. And conversely, given that the proposition ⟨Socrates exists⟩ is true, there must be Socrates. Yet clearly there is an asymmetry. The proposition is true because the man exists and not vice versa.

In a realistic ontology it is hard to avoid a phenomenon like grounding to become an ‘extra’ piece of structure: an addition to being itself. If objects x and y just exists, then what does it mean that x should be ‘ontologically prior’ to y ? What *could* it possibly mean? Let it be granted that physical support or temporal order is not the intended relation, then in what other, metaphorical sense could x ‘precede’ y ?

Yet, intuitions, even very inexact intuitions, often have their rationale. Image theory is in no bad position to deliver certain insights into those behind grounding, and the concept of rootedness will bear much of the load. Let us first round up some of the intuitions that have come along in the above chapters. Of certain compound objects we have said that they are rooted in their parts, and of some abstract objects, that they are rooted in their exemplifications. The suggestion of layers, and of order, both of which are hard to locate in a realistic understanding of the world, is allowed an alternative explanation in image theory. An interpretation that easily comes to mind is this: since, in the course of the running of procedures, objects emerge due to the *flow of information* as it passes through the images of the cognitive system, the emergence of one object is often dependent on that of others. As we have seen, the result of one procedure can be used by a procedure that belongs to a next layer of processing. This, or so my proposal goes, is the main origin of our intuitions about the ontological dependence of certain objects on others. Henceforth I shall treat this as a definition: x is rooted in y iff running the procedure for x involves running that of y .

An explanation like this should provide tools to evaluate grounding claims, assuming that the latter are correctly identified as rootedness. Consider fusion again: although it may be thought obvious that parts come before the whole they compose, Schaffer (2010) demonstrates that, for certain cases, there are good arguments for exactly the opposite view:

Consider a circle and a pair of its semicircles. Which is prior, the whole or its parts? Are the semicircles dependent abstractions from their whole, or is the circle a derivative construction from its parts?

Here it seems obvious that the circle is there ‘before’ the semicircles; but all the other cases (tables, necklaces, less natural fusions), where the whole does look like being ‘composed’ from the parts have not been forgotten. Hence, Schaffer’s nuanced position:

I think common sense distinguishes *mere aggregates* from *integrated wholes*: “that which is compounded out of something so that the

whole is one—not like a heap, but like a syllable” (Aristotle). Common sense probably does endorse the priority of the parts in cases of mere aggregation, such as with the heap. Yet common sense probably endorses the priority of the whole in cases of integrated wholes, such as with the syllable. (*ibid.*)

This appears to be in accordance with the interpretation in terms of rootedness: the concept of a semicircle (or a syllable) is probably always embedded in that of a circle (or a word in which the syllable occurs). Procedures for what Schaffer calls ‘integrated wholes’ are commonly *simpler* than those for parts of them: a algorithm performing a test for circularity is less complex than one for semi-circularity. To ‘see’ a semicircle means to add details (viz. the borders) to the more abstract notion of circularity. If, however, the whole does not answer to any simple (integrating) concept, it is a ‘mere aggregate’, and the whole can only be grasped once the parts are in place.

How is this with abstraction? It may be argued that, when seeing a tiger, one first sees the general traits that makes something a tiger, and only then pays attention to the details that distinguish the individual tiger from others. However, the account of immediate identity as it has been given in 6.1 lends support to the interpretation in which *the tiger* is rooted in Shakti and her conspecifics, rather than the converse. Even though the temporal order may be first to see tigerhood in general before recognizing Shakti, the flow of information goes from the single underlying condition (*C*: the individual tiger), to all possible abstractions that follow from it: *P*, *Q*, and also *the tiger*. Likewise for Snow White’s beauty and beauty in general. It is for this reason that concrete objects will always have primacy relative to abstract ones (Schaffer, 2016).

The concept of rootedness has interesting applications. It follows immediately from the above that *x* and *y* ($x \neq y$) can be *co-rooted*, i.e. one is rooted in the other or both are rooted in some *z*. Image theory is thus able to acknowledge as *different* objects things that realists would feel to be inseparable. One famous case was discussed in 2.4:

80. What you see there is the same river as you saw yesterday, but not the same water.

Similar situations can be found in all those famous conundrums about statues and the lumps of bronze or marble they are made of. The draft ontology is quite stubborn in withholding clean, unproblematic identity to such couples (Rea, 1997; Wiggins, 2001; Fine, 2003). In 2.4 I have uttered concerns about the solution that *relative identity* offers, but postponed further considerations. We are now in a better position to address these cases. The analysis by means of relative identity reads (80) as stating that *x* (what you see there) and *y* (what you saw yesterday) are related in the following way:

81. $x =_R y \wedge x \neq_W y$

The notation makes use of indexed equality signs, i.e. x and y are the same *qua* river ($=_R$), but not the same *qua* water ($=_W$). What I would like to propose instead is that the river and the water are two *distinct* (but co-located) objects, both today and yesterday, and although the river is the same on both occasions, the water is not. For this all we need is plain identity:

$$82. r_x = r_y \wedge w_x \neq w_y$$

(The subscripts indicate the object as it is today (x) and yesterday (y)) Hence, there is one river, but two ‘waters’.

Of course this removes any problem with Leibniz’s Law (if things are the same, they are the same *tout court* and have all their properties on common), but so far ignores the intuition that there is a fairly intimate relation between the river and the water. What is this relation, if not identity? It is, of course, co-rootedness: the river and the water are sustained by the same perceptual input (‘What you see there’). Let \parallel express this relation, then we can write:

$$83. r \parallel w_x \wedge r \parallel w_y \wedge w_x \neq w_y$$

(where $r = r_x = r_y$).

Relations of rootedness, just like identity, are *purely ontic relations*. As was pointed out above, purely ontic relations other than identity and mereological relations sit somewhat uncomfortably with realistic ontologies. As a consequence, such approaches, once the mereological options have been tried and have failed, typically run out of theoretical resources. What the above considerations show is that there are potentially far more ways in which there can be ‘overlap’ between objects. To point at an object (‘What you see there . . .’) is always to point at *more* than one object. Pointing at the river *is* pointing at the water, but the water is not the river. The mereological approach, putting the focus on whether distinct objects can have the same set of parts (‘extensionality’), misses the fact that distinct objects can coincide, e.g. by being co-located. I take this to be an illustration of the failure of *objectual compositionality*, discussed in 3.2. For, to keep our ontology at peace with Occam’s Razor, we would not like to deny that both the river and the water are made of the same molecules! Hence, *contra* the widely held doctrine of the *uniqueness of composition* (Varzi, 2008; cf. also 4.1), in a transcendentalist outlook on the world the same parts may have different mereological fusions.

Notice that, if this account of the relation of (co-)rootedness is accepted, the role of mereology as one of the central notions in ontology gets trimmed down considerably. For, just as the *same* group of parts can give rise to *different* wholes, the converse is also true, since losing or gaining parts (Varzi, 2000; Herre, 2010) need no longer have dramatic consequences either (dropping some whiskers or even losing its tail will not make its owner fall out of reach of the average cat-image: the image is flexible enough for that). But this means that parts as such

simply do not achieve half as much as they are often taken to. The *way in which* they are put together contributes far more to the making of the compound object than exactly which ones they are.

There is, to the best of my knowledge, no single natural language in which relations of rootedness are granted their own linguistic vehicle. What I mean is a simple general-purpose word or other linguistic routine catching this concept—in circles of philosophers there are of course terms of art, and every language has ways to describe the situation. They share this deprivation with concrete-abstract relations, but not with mereological ones (part, whole, comprise, constitute, compose, ...). Call it an embarrassment for the ideas behind the draft ontology, but relations of rootedness are commonly handled with the verb ‘to be’ and its analogues, which makes them look like identity relations:

- 84. See that bird? It was the first pet I had when I was a child.
- 85. This path is the border between Belgium and The Netherlands.
- 86. The baby on this picture is the president of the United States⁵.
- 87. It just *is* them. They just *are* it.

I take it to be obvious or at least highly credible that identity is not what is at issue in these examples. Assuming that it is not the very same bird in its old days, (84) mixes up the individual with the species, even though an individual bird is not the same as its species, not even nearly. Even if a walking track exactly coincides spatially with a national border, it is not that border (should, due to heavy rainfall, the path change its course by a local detour taken by the walkers, then this will not affect the territory of either of both countries). The baby on the photograph has many properties that would disqualify it for the Presidential Office. And even though commitment to the one arguably goes hand in hand with commitment to the other, a cat-fusion is not the same thing as two cats (cf. 4.2)—or so I wish to defend.

6.3 Existence

Ontological scruples notwithstanding, existence, in the draft ontology, is just a property: one which some objects have and others lack. Santa Claus, Sherlock Holmes, Pegasus, Vulcan, the luminiferous aether, and many other objects, *do not exist*. Alexius Meinong, the prominent theorist of non-existence, indeed claimed that there are objects such that *there are no such objects*. One might think this to be a problem for any ontology, since it seems to entail logical absurdities like:

- 88. $\exists x : \neg \exists y : x = y$

⁵Cf. Morzycki, 2016, p. 45

which conclusion proclaims the existence of non-self-identical objects (Van Inwagen, 1977). To my mind, however, it is a problem for any *realistic* ontology. Efforts of Quine and others to defuse the problem by a better use of language have proved mostly futile. Apparently unimpressed by realistic prejudices, the draft ontology is incurably committed to objects that do not exist, and to distinguish the one category from the other, the existence predicate comes in handy:

89. There are things that do not exist

or:

90. $\exists x : \neg E(x)$

I shall argue that (88) is the wrong way to spell out Meinong's dictum, and, consequently, what has to be done is find a sound interpretation for (90).

One tremendous difficulty with doing justice to our Meinongian intuitions is that, when granting *reality* to, e.g. fictional objects, one appears to be forced to hold that they have their properties in a different way than real objects have theirs. It is not hard to see why. Every version of Meinongianism is committed to something like the so-called *characterization principle*, which is that fictional objects must, short of existence, have all the properties they have according to the fiction. Now, the *literalist* version of this doctrine is obviously dubious. Berto and Jago (2019):

Holmes literally is a detective and literally lives at 221b Baker Street. But this is problematic. Until 2002, 221 Baker Street (there's really no 221b) hosted the Abbey Road Building Society. It has never been the home of any detective.

(...)

If Holmes really is a detective because he is so characterized (...), why can't we shake hands with him? (p. 261)

But there are other ways to make sense of fictional entities. Van Inwagen (1977, 2000) considers the likes of Sherlock Holmes to be objects that are real (albeit *abstract*) and have properties like any others. In the case of Sherlock Holmes these are, among others, being a character in a nineteenth century novel, and being very famous—*not* being a detective and living in London. The latter are properties that the fictional entity Holmes *holds*, not has. Zalta (e.g. Zalta, 2012) makes a similar distinction when he speaks of properties that an object *encodes*, rather than exemplifies. He considers fictional objects to be comparable with abstract ones.

Construals of the characterization principle of the latter (non-literalist) type, however, come at a high price. We can understand the story about Holmes because we know what a detective is, what the name 'London' refers to, how to

interpret all the details in the descriptions. The overwhelming impression is that we do so by interpreting the same language in the same way. Moreover, these accounts make it hard to deal with the range of *intermediates* between plain fiction and plain reality: some non-existent objects are *believed* to exist or considered possible. Caplan (2004) makes the important point that fictional objects are unlikely to be ‘ontologically special’ in a way that mythical and imagined objects are not. The same seems to be true for objects whose existence we are simply *ignorant* about. I do not believe in the Loch Ness Monster, but Mr. Spicer does—yet there is much we agree on, it there not? It is very big, and a good swimmer. And if, against all odds, Nessie should suddenly raise its head above the surface of the lake, will it thereby start to ‘have’, rather than ‘hold’, its size, colour, and frightful appearance?

Image theory is in a position to allow non-existent objects to have their properties in exactly the same way as existent ones: to allow their *Sosein*, as Meinong called it, to be quite independent from their *Sein* (‘being thus’, as opposed to ‘being’). The key to this interpretation is to clearly distinguish the fictional object itself (Sherlock Holmes) from *the piece of fiction* in which it figures. Of the latter there is no dispute about its reality, which is attested by books, movies, and the like. But Sherlock Holmes *is* not that fiction: *he* is a detective, living in London. Can we shake hands with him? Surely we can—but only in fiction, of course. Were I a literature teacher, I could do so in a fictitious interview to introduce the character to my students.

Accepting this: accepting that there are non-existent objects in a fairly literal, Meinongian sense, compels us to face a domain of things, all fit for being introduced by ‘existential’ quantification, some of which, however, ‘exist’ in a stronger sense (Turner, 2010). Despite the bewildering illogicality of all this as seen through realistic eyes, in a transcendentalist appreciation nothing too worrisome is going on. And nothing paradoxical, since there are, indeed, different kinds of existence. Here, I would suggest, is the split: some things can be seen, heard, or felt; while others only live as an idea. At the image level: the image is present, but receives no perceptual support. This is quite common: many scenarios never come true, either because of failing anticipations, or because they were never intended to, as in fiction. Or—even more importantly—they *nearly* come true. I planned to go for a long walk, but the last few miles were done by bus. It may be illuminating at this point to return to the Platonic objects discussed in 5.9 once more. *The circle* does not strictly apply to any concrete circle: no object complies exactly with what it takes to be an instance of *the circle*. However, many real things are good enough to trigger *some* approval: they are *almost* circular. No wedding is perfect, but some come a lot closer than others, likewise with bank robberies. Within this broad scenario there is plenty of room for images arising in the human mind that do not fit any perceptual input, maybe never will, for some of whom it is even impossible to do so. Nothing in the exposition of image theory so far blocks this possibility. But then, as an immediate consequence of

the Central Assumption, these are the non-existent objects.

Meinong coined the term *subsistence* for the condition such objects are in, as opposed to *existence*. Here I shall use the term *being real* for those things that exist in the strong sense—non-existing objects will be (pragmatically) those that are unreal. Hence, the real objects are those that can be perceived. There is, in this respect, no privileged position for the naked eye or ear: Higgs bosons and gravitational waves are real because—mediated by lots of technical devices—they have been perceived. With this the notion of (real) existence comes close to that of the verificationists of the Vienna Circle, and it is one that makes existence essentially *contingent*—nothing is necessarily real—which seems good.

Notice, meanwhile, that the existence of an object is thus established by its *potential* of being perceived. This must be read as a matter of principle: it might well be that for all practical purposes the intended perceptual event is entirely out of the question; yet, in the transcendental approach laid down here, to say that an object exists just is to say that it will be perceived when, and if, ‘looked’ at. Likewise for its properties: image theory is, by its very nature, vulnerable for the misconception that an object’s being *P* is ‘locked’ to its being actually scanned. Taking this seriously, the existence of objects and properties beyond what is immediately perceived would be cast doubt upon. Before trying to patch this up this defect in the explanation, however, let us pause to admit that it is not an artefact of the theory. It is very real: it is the old riddle of the sound of the falling tree deep in the woods. If no one was present to hear the sound, what does it mean to say that there was a sound? If an orange mushroom grows one day, then withers before anyone has passed by, what does it mean that it was orange? The only way around this problem, to my mind, is to apply the primitive notion of counterfactuality (see also 6.5): *should* someone have been there, they *would* have heard a sound, or seen something orange. (Notice that this counterfactuality only pertains to the explanation of why and how we perceive properties; it has no bearing on the obtaining of the properties themselves.)

For abstract objects the situation is less clear-cut, but this is not in any way counterintuitive. In my interpretation of abstract objects, as worked out in Chapter 5, *the tiger* can be seen (e.g. in the Zoo), and *the cuckoo* can be heard, which makes them real. The same is true of *the colour blue*, but with respect to *the number three* doubts start to rise. As *three* is rooted in triples of perceptual objects, it has some claim on being real. But it also seems that, for abstract objects, since they are at a larger distance from sensory experience, existence is less of an issue. The deeper the abstraction, the less the ‘need’ for being real.

There is another well-known objection that, for lack of target, has not yet been brought up against image theory, but otherwise doubtlessly would have: the argument from *indeterminacy*. In Van Inwagen’s words:

[I]f the Meinongian is asked, “About your Mr. Pickwick—has he an even number of hairs on his head?,” he will answer (Dickens having

been noncommittal on this matter), “He neither has nor lacks the property of having an even number of hairs on his head; he is therefore what I call an *incomplete object*.” (Van Inwagen, 1977)

Treating Sherlock Holmes in this way (Doyle having been equally neglectful), however, would seem to pull the carpet straight from under the characterization principle. Conan Doyle certainly did not tell a story about an incomplete person—nobody in the stories ever noticed anything strange about Holmes in *that* respect! He is a man of flesh and blood, with (like any other such man) a well-determined number of hairs on his head.

The explanation I would favour is that, here again, the confusion is due to a mixing up of the fiction and the fictional character. As for the latter: non-existent objects are never incomplete, any more than real ones are. What is (always) incomplete, at least for concrete objects, is our knowledge of them. As was argued in this regard in 2.4, transcendentally, perceiving/thinking about *something* P ($\exists x : P(x)$) does *not* mean—in contrast with the realistic picture—that there is some specific P one is perceiving/thinking about. The content of the perception/thought itself is ‘some P ’. Therefore, transcendentally the so-called *selection problem* (after Sainsbury (2009): ‘[H]ow does an author pick out, from among the cloud of nonexistents, the right one to be the fictional character he wants to create?’) does not arise. To increase our knowledge about real people, we have to count hairs. To find out about Holmes, the authority is with Doyle: if he says Holmes has an odd number of hairs, that is the truth; but if Doyle forsakes, *both* extensions of the story are equally acceptable: the one about Holmes with an even number of hairs, and the one about Holmes with an odd number. None of them exists. None of them suffers from incompleteness.

For a final note, it is interesting to see how the familiar paradoxes of the existence-predicate fare under this proposal. Take the ontological argument for the existence of God, first presented by Anselm of Canterbury. Let God be a being of the greatest conceivable perfection. If this being is unreal, then it lacks the quality of being real, and this means that it is *not* of the greatest conceivable perfection—contradiction. Hence, God is real! The usual objection, that being real is not a property like being strong or omniscient, has to be reconsidered. Meinong, after his pupil Ernst Mally, distinguished *nuclear* from *extra-nuclear* properties⁶, existence being extra-nuclear. This distinction has been criticized for being badly defined, but image theory has the resources to give it a new meaning.

Images are templates of properties, searching for a match by scanning reality. Clearly, however, existence itself cannot be among the properties thus scanned. The only meaning that can be given to ‘scanning existence’ is that, *after* the scanning as a whole has taken place, it yields a positive outcome. Hence, the nuclear properties can be identified with the properties encoded in the image:

⁶The English terms are due to Parsons (1980)

the ones that are involved in the scanning. These are also the properties that characterize the object, its ‘being thus’. They appear to determine the way we think about it. As Kant put it:

Thus when I think a thing, through whichever and however many predicates I like (even in its thoroughgoing determination), not the least bit gets added to the thing when I posit in addition that this thing is. (Kant, 1993, A600/B628)

That our brain recasts information about existence as a property of the object is, it must be admitted, somewhat unexpected⁷. But this analysis clearly upholds the most common diagnosis of the flaw in the ontological argument: being real as a property is fundamentally different from being e.g. strong. However perfect the conception of the Supreme Being, being real cannot be part of it.

6.4 Imagination

The fairly broad term *imagination* I shall use here to refer to the interplay of images as they govern our mental life, *independently* from perceptual input. Theoretically, images combining all kinds of properties could arise, yet, although much of this happens in fact, we do not feel there to be what Berto (2018) calls a ‘runabout imagination ticket’. There is an inner structure limiting the ‘running about’ of images. How to account for this observation within image theory?

In this context it is good custom to quote David Hume’s statement to the effect that it is (logical) possibility which sets the boundaries:

’Tis an establish’d maxim in metaphysics, that whatever the mind clearly conceives includes the idea of possible existence, or in other words, that nothing we imagine is absolutely impossible. (Hume, 2003, p. 23)

Established or not, this idea is not widely endorsed today (Priest, 2016; Berto and Jago, 2019). In image theory, although the analysis of conjunctive properties in 6.1 may cast some doubts on the conceivability of a round square, images that scan for intrinsically contradictory properties cannot be excluded—if only for the sake of argument it seems far better not to. If so, then by the Central Hypothesis there is such an object, even though it is non-existent, and necessarily so.

If images can pick their properties from all the conceptual repertoire the cognitive system has, without guidelines whatsoever, then how is it that we do see lots of structure in normal people’s mental life? At least part of the answer is that, as

⁷Arguably, however, there are more putative extranuclear properties that could be vindicated by image theory, e.g. *being possible*, *being complete* (Berto and Plebani, 2015), but also *being abstract*.

we have seen in Chapter 3, maintaining images comes at a cost for the organism, and thinking about round squares (for those not involved in the metaphysics of impossible objects) is a waste of resources. Yet, from this it only follows that one has to use one's imaginative resources sensibly; but what does *that* mean when we are only concerned with inner mentation, with no role for perception? It seems that, if imagination is to be of any practical use for the organism, then, despite the absence of perceptual information, it must be able to generate *knowledge*.

As a matter of fact, the scenario for this to happen is neither deep nor complicated. Image theory—viewing the mental processing of perceptual information as an essentially algorithmic passing over of the output of one image to another—is committed to a picture whereby the informational channels are already present *before* any perception actually occurs. Such a system is excellently equipped to run *simulations*, even in cases where no perception occurs at all. No doubt simulation is an all-pervading feature of our mental functioning. In their integrative review about different types of mental simulation (experience, practice, consumption, and goal achievement), Kappes and Morewedge (2016) conclude:

In each case, mentally simulating an experience evokes similar cognitive, physiological, and/or behavioral consequences as having the corresponding experience in reality (...).

So the structure of the argument must be this. No one appears to be surprised that we can reach sensible insights when there is actual sensory input. *That* this can be done—that the human brain has evolved all the way from unicellular bacteria to the sophisticated computer it is—is of course by no means trivial or obvious, but given this to be the case, and given that it works by and large in the way it does, the ability of the brain to simulate is no big miracle. All that has to be done is to put the machinery in motion *as if there were* perception, while at the same time adding some cognitive disclaimer to the effect that it is not really there⁸.

The type of knowledge that is acquired in this way is *conditional*: *if* this and that is/would be how things are, *then* this and that will/would follow. Conditional beliefs do not always start at the perceptual level; they could as well start 'higher up', as when one thought follows from another. Berto (2018) describes this process using the example of weighing one's chances of jumping over a stream⁹:

Will you succeed if you jump the stream? You don't just blindly try, for that may be dangerous: you may fall in the water, or hurt yourself. So you mentally simulate the scenario: you imagine jumping. You integrate the input importing beliefs concerning the distance, your physical abilities, past performance, etc. You let the story unfold.

⁸The disclaimer appears to be fallible, given the *substitution effects* that Kappes and Morewedge (2016) describe, e.g. that imagined events are attributed an enhanced likelihood.

⁹The example is from Williamson 2016.




Figure 6.1: The introduction rule for implication

In the imagined scenario, you get to the other side. You add the conditional ‘If I jump, I will get to the other side’ to your belief stock.

The addition is a valuable one. It is a piece of, in principle non-empirical, knowledge. The conditionals we entertain in daily life of course enjoy a considerable input from memory and experience (‘physical abilities’ and ‘past performance’), and are therefore inductive, rather than deductive; but even in cases where they do not: even if *no* perception whatsoever is involved (not even a stored memory of past perception) a conditional belief is a real *belief*, a judgment, not mere musing (Williamson, 2016).

This is a point deserving of some emphasis. Theoretically, belief acquisition through simulation parallels the use of the *introduction rule for implication* in the system of natural deduction for propositional logic. The rule provides for the option, at some point of an argument, to insert an *assumption*, say p (often on top of what is already known), then proceed the reasoning until some conclusion q has been reached, and then *discharge* the assumption and conclude $p \rightarrow q$ (Figure 6.1). The ‘temporal’ insertion of p in the argument is just like a simulation: it is a test for *what-if*. Here we find the computational basis for the principles of *consequence* and *hypotheticality*, discussed in 1.4. It is to be expected that this cognitive aspect is an important factor in guiding imagination (cf. also Stalnaker, 1968; Ramsey, 2000).

In the literature the distinction between what is and what is not epistemically relevant is sometimes framed as that between ‘free’ and ‘constrained’ (Kind, 2016), or ‘voluntary’ and ‘involuntary’ (Williamson, 2016) imagining. E.g. Kind (2016):

[I]t is our ability to constrain our imaginings in light of facts about the world that enables us to learn from them.

This appears to suggest that it is, first of all, the *limiting* of imagination that endows it with epistemic relevance. However, in assigning constraining as such this central role I believe that something important is at risk of being overlooked.

Imagination—in those blessed with a reasonable state of mental health—is of course constrained, but this is not what *makes* it epistemically relevant. What makes it relevant (absent perception) to have images of p and q is that there may or may not be an implication running from the one to the other—and this is not something one can choose for! It is the *relation* of ideas (as Hume would call it) rather than the ideas alone. From ‘I jump’ there is an implication to: ‘I will get to the other side,’ not to: ‘I will fly away,’ even though both might be imagined. It is *therefore* that thoughts about flying away are less practical to have than thoughts about landing on the other side.

Hence, my suggestion is that the constraining that is (correctly, it seems) introspected in the case of sound reasoning, rather than being the leading force, has to do more with the weeding out of useless thoughts after cognitive evaluation. Trimming the stream of consciousness is done to save mental resources, not to prevent madness. This should change our appreciation of unconstrained imagination. Kind (2016):

(...) I do not mean here to privilege constrained imagining over unconstrained imagining. Both of them have their place in human life. Our unconstrained imaginings are important for many of the activities that are important to us, and rightly so—from engaging with art and literature to fantasizing or pretending.

This association of unconstrained imagination with epistemically ‘soft’ activities like art and fantasizing does not seem to give this activity its due credits. To my mind, unconstrained imagination (or rather, *relatively* unconstrained imagination: no human activity is entirely at random) is better seen as *searching*. Mental creativity is allowed to move more or less freely within some distance from a starting point, in order to explore the conceptual territory so as to discover the most promising hypotheses. In terms of Figure 6.1 this means testing various p 's to see whether there are implications $p \rightarrow q$ (in most cases only probabilistic, of course) to be found which are worthwhile. This playing with ideas may happen with something of a plan behind it or in a totally haphazard way, but the cognitive validity of the outcome will not depend on the process of its formation.

6.5 Modality

Transcendentalistically, the road from imagination to modality is a short one. Modality is about how things could be, else than how they are. Consonant with the departure from realism I have advocated, I would not suggest to embrace *modal* realism à la David Lewis (1986a) either, but I am nevertheless glad to make use of his method of specifying modal statements in terms of non-actual worlds

or scenarios¹⁰. In our experience of the world, non-actual scenarios are extremely important: everything actual is constantly surrounded by clouds of possible alternatives, many of which are not even known to be non-actual. Furthermore, many—if not most—of the properties that we ascribe to objects have a modal character. An egg is *fragile*, even if it does not actually break—we can almost *see* this when we look at it. This is also the way in which a snake is dangerous, a concrete floor hard (suppose you fall!), and a plan promising.

Lewisian modal realism may be counterintuitive; one thing I take Lewis to be perfectly right about is that in a realistic view on the world it is hard to find a credible *non-realistic* way to describe whatever is non-actual. Alternatives like the various kinds of *ersatzism* (e.g. Heller, 1998; Sider, 2002; Jago, 2012) face a similar problem as was discussed in regard of fictional entities: how to fulfil the *characterization principle*? If the objects in non-actual worlds do not have their properties *in the same way* as those in the actual world, then how do they have them at all? Unabridged realism—where this condition is met in the most immediate of ways—always seems to have the upper hand. This is important, since in the draft ontology non-actual scenarios are presented in exactly the same terms as actual ones, albeit under some modal operator. Furthermore, the argument about intermediate cases, as it was given in 6.3 with respect to fiction, seems to apply here just the same. Possibilities, by their very nature, *could* be actual, so it would be strange if the objects figuring in any mere possibility would have their properties in a different way. But then, how could they *not* be real?

In image theory the solution for this problem follows quite naturally from what has been said in 6.3. There are non-existent objects, and they have properties just like existing ones—what distinguishes them from the latter is only that they are unsupported perceptually. Hence, all worlds but the actual world are, transcendently speaking, non-existing. Yet, this shortcoming does not prevent us from thinking about them and thinking them through: it is in this way that we consider counterfactual situations (as was described in 6.4) and their consequences. Therefore, image theory is able to account for our modal intuitions in a very specific way, keeping the advantages of possible worlds, while avoiding both the extravagance of realism about what only *may/might* have been, and the artificiality of ersatzism. The transcendentalist has the option wholeheartedly to endorse Stalnaker (2003):

[I]t seems to me that modal notions are basic notions, like truth and existence, which can be eliminated only at the cost of distorting them.
(p. 7)

I take it that the draft ontology is quite clear about alternatives to the actual state of affairs. When we say that it could have been raining today, irrespective

¹⁰Considering entire worlds is convenient for logical purposes, but not even the actual world is ever wholly before our mind's eye, let alone the others. But for the said logical purposes, scenarios can easily be interpreted as sets of worlds.

of the possibility of rainy days in alternative universes otherwise just like this one, this comment does not seem to be about such places. What it means is only *that it could have been raining* here and now. There is no deeper level of interpretation.

But there *is* a deeper level of explanation. Modality is primitive on the level of objects, but not on the image-level. Just like objects, according to the Central Assumption, emerge where procedures reach a unified output, modal notions seem to emerge due to the possibility to imagine (to form images of) non-real objects and situations, and draw conclusions from them. In that sense, transcendentalism allows for *actualism* without the need to resort to primitive dispositions, essences, or potentialities (Fine, 1994; Vetter, 2011, 2015). This is not to say that essences, dispositions, or potentialities might not be very useful in illuminating our human appreciation of possibility, necessity, and counterfactuality—that these are strongly tied to actual objects, rather than general laws and regularities seems a most reasonable assumption. But a fact that tends to go missing in this approach is that dispositions do not in any way obviate the need to account for *alternatives* to what there actually is. There must be scenarios different from how things really are, but otherwise just the same.

7.1 On Balance

If all you have is a hammer, everything looks like a nail, or so the most popular wording of A.H. Maslow's aphorism goes. In the preceding chapters lots of hammering has been going on. Image theory is not the worst of hammers: some of the nails sank in quite smoothly. Others took some more force, however, and of some it can be disputed if they were really fixed or just bent. Maybe an occasional *screw* was rammed in. Be that as it may, image theory, as presented here, is an unfinished project. A closer look on the matter will reveal, I believe, apart from less convincing claims, also unexplored territories. This chapter is dedicated to discuss both.

Most theories, besides a positive message, have a negative one also—often the latter is more convincing than the former, and image theory is not necessarily an exception. Realism—the view that objects are present in the world independently of human cognition—has been something of a strawman in the preceding chapters. I believe some forceful reasons have been provided to show that, no matter how charitably one is prepared to look at the matter, the prospects of realism ever convincingly to come to terms with the way the world appears to us are bleak and will remain so. It is not that there *could* be no independent reality that mirrors the draft ontology—how could that be excluded?—but the parallels between the structure of the world and that of our cognitive system are too conspicuous to be a coincidence.

Does it follow, then, that any better ontological theory should be transcendentalist? It seems so. Denying the existence of any independent reality whatsoever (even though this flag has been defended vigorously many times in the history of philosophy) seems outrageous, but denying that our perceptual and conceptual abilities have a strong—indeed, a decisive—impact on the way we experience the world does not look too promising either. The world as it appears to us, if based in reality, is shaped by cognition—that is transcendentalism. The genius

of Kant's idea is that it allows a sensible appreciation of the subjectivity of our view on the world, without letting go of the fundamental insight that the facts themselves stand unmoved by our cerebrations.

Does *image theory* follow then? Surely not: image theory is just one way to cash out transcendentalism. Kant's own transcendental idealism is very different theory in many respects. Furthermore, image theory is not a single theory; it is just a framework in which explanations for various phenomena can be formulated. The version presented above has its own presumptions, viz.:

- The draft ontology is the right target of explanation.
- The concept of object has a transcendental nature.
- The human cognitive system is computational.
- It makes sense to think of images as the functional units of this system.
- The role of procedures is correctly described by the Central Assumption.

Each of these is open to amendment.

Thus, the ideas unfolded in the above are not a wholesale package. What has been said in the above can be read as:

- I An argument against realism
- II A plea for transcendentalism
- III A couple of proposals within the framework of image theory.

The claims are in decreasing strength. It is possible to accept I and II while rejecting all or some of the proposals under III, or to accept only I. And of course one might opt for I and II, but choose a different theory to explicate II—or stop at II and leave it there.

And of course one could reject it all. Then again, some concepts belonging to image theory may have *heuristic* value: they may be useful outside this framework as well. It appears to be a good attitude towards theories in general to treat them as more than just some amount of truth or falsehood. Ideas and concepts belonging to theories—even false theories—may have some worth by giving rise to new ideas, or prove more fortunate in a different context.

7.2 An Explanatory Limitation

In what has been presented in this book, the concept of *image* is central. The concept endows the theory with advantages over realistic ontologies. All of these advantages, however, have been bought for an assumption that at first blush may seem utterly question-begging.

Here is why. Realistic ontologies all appear to have their loose ends when it comes to explaining how predications connect to objects, as we have seen when discussing Bradley's regress (3.1). Universal theories need a relation to connect universals to individuals, which must be unlike other relations (i.e. it must be a 'special' universal or no universal at all). Trope theories need similarity and compresence, which must also be unlike all other relations (i.e. 'special' tropes or no tropes at all). There does not seem to be a way around this: to take the sting out of Bradley's regress the newly arising relations must be less substantial, to be taken less seriously, less real, something like that. These are fundamental problems in every realistic theory, but does image theory fare any better? Here, rather than a couple of weakly justified primitive concepts, there appears to be a *circularity* at the heart of the explanatory clockwork of the theory. Understanding the working of images requires knowledge about how they run their procedures and how these procedures themselves can be described, their structure, subroutines and so on; in sum, their computational structure. This is essentially the type of research done in neuroscience, artificial intelligence, and related disciplines. But of course, when speaking about how images 'scan' their objects, we refer to the interaction due to *their* properties and relations, and if *these* stand unaccounted for, then any resulting explanation seem to rests on shaky ground. Was not all this set up to *explain* properties and relations?

This is the limitation I anticipated in 1.6. One consequence of the transcendentalist approach to ontology as defended here is the fundamental impossibility of analysing predication in such terms as have been purged of the concept itself. Let us say that we can see that Snowy's fur is white. Although we can dispute the objectivity of this condition and invoke an explanation in terms of images for how we experience it visually, denying that Snowy's fur is at least *a certain way* would instantly make it a mystery how it could be scanned according to some procedure, causing the image to produce a positive signal. From the viewpoint of image theory, the facts about the fur—those that are responsible for how it reflects light rays—are also what causes the interaction with the image. Considering this, it might seem as though image theory is viciously circular in the following way: we see that Snowy is white because our image for whiteness produces a positive answer with respect to Snowy—which is caused by Snowy's being white!

But this is not quite true: despite containing circles of this type, explanations by means of images are not generally vacuous. There is no belying that the explanations for the foundational notions of image theory are only delivered by approaching images as just parts of our brain: things alongside tables, chairs, and cellphones; and this reasoning 'at the image level' is done with knowledge about objects in general. It is this bit of ontology which provides the presumptions whereby the ontology of the world is explained. However, the latter is importantly *richer* than the former. To describe images and their interaction with objects, only a small group of very clear-cut concrete objects are needed, with a limited

set of properties and relations. No kinds or attributes are needed, no meaning, reference, or intentionality, and surely no non-existent objects. Nothing much, one could say, that Quine would have disliked; only some idea of possibility and necessity cannot be missed to understand how images give rise to procedures. This (among other things) is why I take modal notions to be principles of the understanding (1.4).

Thus, there is circularity, but, one could say, a *productive* circularity, which is enough for being non-vicious. Clearly, though, whoever comes along and produces an explanation where these circles have been removed, will have done a better job. But there are good reasons to be sceptical about the possibility of such an explanation. If what has been stated in Chapter 1 is true, then it is fundamentally impossible for us to sidestep our understanding of reality in terms of objects.

7.3 Objectivity

Some issues transcend the purely technical aspects of image theory. It might be that (some mature version of) image theory could smoothly explain the entire *structure* of the draft ontology: the inventory of concrete and abstract objects, their relations and inferential connections, the causal links between them and what not. Then there would still be some questions that would not even have begun to be addressed, one of which is *objectivity*.

Image theory is about experience of objects and their properties and relations, but the way these are handled by the theory could seem to make them intolerably *subjective*. If all properties are ‘in the eye of the beholder’, does that imply that the world in fact has no properties, which means that everything is really the same? The shallow answer is: of course not; image theory has a different way of accounting for the structure of the world, but does not deny it. What gives image-based properties their objectivity is that different images, belonging to different beings, can function in exactly the same way, i.e. have the same group of objects as their (emergent) extension—they typically do so. And surely there is no *volition* involved in determining what will be the result of scanning reality. As the description of the procedure for even-segmentedness shows, the working of procedures is fixed.

All this is true, and one might wonder if fixedness of the outcome of a procedure, irrespective of where and by whom it is performed, is not objectivity enough. But there are some pitfalls on the road. Might there not be ‘alien’ intelligence, that apprehends the world in ways radically different from ours? Such intelligence might be thought of as based on a cognitive system that cuts the pie of reality in a completely different way than we humans do. I do not mean *bats* in the role Nagel (1974) had in mind for them. Their experience of the world is bound to be very different from ours, but when a bat detects a fly, we can see one as well: we agree that there is one. Clearly, there are whole domains of objects that the bat

does not know (and who knows what is beyond *our* perceptual reach?), but in the area where we are both perceptually active it does not look as though there is a *massive* mismatch between the objects that exist from our viewpoint and that of the bat. Really alien intelligence would e.g. ‘see’ objects that from our point of view are entirely random gatherings of elementary particles—or worse.

The options for alien intelligence are likely to be limited, however. Notice that, *if* there is such intelligence, it will probably be forever hidden from our sight. For, if we could see, or otherwise detect, such beings, they would consist of normal matter, the same stuff that stones, doorknobs, computers, planets, and stars are made of. This is the type of stuff that, in principle, we can analyse. We could gain insight into the perceptual organs of these beings, and do research into their cognitive systems. But if we would have access to the workings of all this, we could also make use of this knowledge to build devices to detect the objects that these beings can detect, and their intelligence would no longer be alien to us. Surely this argument is sketchy, and arguably presumes a lot of technical sophistication from our part, plus the power to study the intelligent beings or their voluntary consent. But all this only has to be possible *in principle*, and, since it does not essentially transcend the scientific study of some worldly phenomenon, it may not be such a tall order after all. If the argument is successful, it has the important consequence that different cognitive systems that know of each other’s existence cannot ever live in *fundamentally* different worlds¹

It may be reassuring, granted this, to know that we, and also animals, computers, and aliens, all ‘experience’ by and large the same joints of nature—but then again, how does that go beyond mere intersubjectivity? The correspondence could, after all, have been arranged by a more sociable relative of Descartes’ demon. Critics of this answer would probably want to hear how it is the result of things *actually being some way or other*. Let there be some intersubjectivity involved: let it be true that Snowy looks white thanks to the way our eyes and brains work; still the fact that he looks that way to us should be interpretable as the consequence of certain facts about his fur. There is a clear sense in which exactly those facts are what it is for Snowy to be white. To put it differently, must there not be at least *some* class of properties which are *ways of being*, not just ways of appearing?

It is not easy to see what sort of answer would count as satisfactory to address such worries. The objectivity of objects themselves has been handled in the negative so far: the world does not cease to exist when we stop observing it. But will running this narrative with respect to properties lead us *far* beyond the shallow answer? That both these bowls are round must surely express a very deep fact about them, but for us humans it is just that they are both round. One could hope that something of the truth is allowed to shine through by such strictly

¹Albeit on different grounds, Donald Davidson famously makes this claim as well in his discussion of the principle of charity (Davidson, 2013).

insufficient statements—think of a Wittgensteinian ladder. But it is quite a big ladder we need here, and it might not even be sufficient. Among philosophers, after all, irrespective if they like it or not, the idea is familiar that some properties are indeed intersubjective. Surely our sense of what a white dog looks like means nothing to Nagel’s bat or facet-eyed flies—it is too simple to decide that these creatures just do not have the right senses to pick up these features, as if they are there all the time, lying in wait to be detected. The facts about this type of properties can be philosophically framed to anyone’s liking, but it is no *great* exaggeration to say that they are ‘created’ by the perceptual tools of the observer.

The only qualification to be made here—although I do not know how much it will effect—is that some properties seem far *more* objective than others. *Structural* properties have before (in 1.3) been singled out as the most convincing claimants of objectivity. Maybe the whiteness of dogs as we feel it to be will die with the last observer with regular colour vision, but not only is the even-segmentedness of centipedes far less dependent on the specifics of the sensorium; it is also quite unthinkable that this feature does *not* reside ultimately in the reality of what constitutes the centipede. Again, the most illuminating words that can be said appear to be in the negative. A final sceptical note can be added: even though something *really is* the matter with a centipede in virtue of which we call it even-segmented, it is quite doubtful—or so I would say—that there is genuine objectivity behind our observation that two even-segmented centipedes have something *in common*. After all, it is the element of commonality that chiefly resides in the image.

7.4 Conceptual Content

A closely related issue is how to individuate and where to locate *conceptual content*. As for the first part of the question: what is needed is convincing *identity criteria* for images: criteria which, for images *i* and *j*, would be sufficient for their encoding the same property *P*. These cannot be missed, since a single object which, given property *P*, is *the* image, is lacking. Every cognitive agent has their own stock of images, and of some of them it may even have several copies; and even the staunchest nominalist will admit that there must be truth conditions for the statement that different observers observe ‘the same property’. Let the discussion about what this *means* be postponed for the moment; it would be nice to know at least *when* it is the case. Observe that the problem is not with the *unity* of the group of objects which are *P*. Each one of the private images for *P* that each of the cognitive agents has does all by itself a perfect job in unifying this class of objects. Unlike universals and tropes, images enjoy a fully unproblematic existence—the only problem is that there exist *many* of them, which is *too* many. It is inconceivable that all images encode different properties, as this would mean that no two people could ever observe the same property.

So when do different images encode the same property? It would seem that at the very least they should be *extensionally* the same, but ‘extension’, in a transcendentalist ontology, is not the simple and basic concept it is in a realistic one. The intended condition, i.e. that two images should return the same result upon scanning the same object presumes that whatever is ‘the same object’ is given beforehand, which is clearly not so. This is a puzzle lying at the heart of the transcendentalist understanding of objects. An answer to what it means for two procedures to be extensionally the same should start with the ideas developed in 6.1. I shall not elaborate on this point here, for to decide whether images stand for the same property, extensional equivalence is not enough anyway. One of the charming features of image theory is the ability to distinguish coextensive properties like *cordate* and *renate*², since the procedures for determining them are clearly different, which implies that they have different images.

How about the images being *intensionally* the same, which means that both procedures will give the same outcome under every possible circumstance (where ‘circumstance’ must be interpreted in terms of ultimate input pattern)? Probably better, but intensionality is not all there is to explain. If there are genuine cases of *hyperintensionality* (as there seem to be; think about the concepts of equilateral and equiangular triangle) image theory should be able to handle them too. So here is the next step: both images should—within the context of their hardware—run *exactly similar procedures*, which means that the concrete scanning events that they produce should be exactly similar *formally* (for similar input). To give an exact definition of ‘formally’ may not be very easy, but the idea is by no means contrived. There should be a one-to-one correspondence between both procedures: with respect to events in the procedure and with respect to their causal relatedness (cf. Duží et al. (2010), who speak of ‘procedurally isomorphic constructions’). The advantage of this idea is that it seeks no recourse to similarity of *image*. The images may be different: in size, in material (neurons or silicon), in source of energy, and so on, as long as they ‘do’ the same thing during the scanning process. Every step of the one image, under every conceivable circumstance, should have an analogue in the other. Thus *multiple realizability* (as per Putnam, 1967) is guaranteed.

A design like this, to my mind, has fair prospects of being the closest one can get to solving the problem of identity criteria for properties in image theory. Its limitations, though, should not be overlooked. A subtle point, yet one of concern, is the following. In 6.1 an explanation is suggested for the transcendental identity of objects. The explanation is speculative, but it could be falsified if it would turn out that the conditions for people to perceive identity do not coincide with the proposed condition on the image level. This is possible only thanks to the fact that the objects are seen from *one* perspective: *x* strikes *me* as being the same object as *y*. However, the hypothesis just formulated is about ‘trans-perspective

²‘Having a heart’, and ‘having kidneys’, respectively.

identity', and therefore it seems a lot less clear how it might be evaluated. What would have to be the case for it to be true or false? A truly illuminating answer to this question would be more than welcome.

As for the second part of the question at the start of this section: clearly image theory locates conceptual content *in the image*. In that sense, despite its equal insistence on viewing thought as computational, it stands in contrast to most approaches within the computationalist tradition, e.g. Fodor's *representational theory of mind* (Fodor, 1975, 1981, 1994), where the content bearing elements of the cognitive system are defined in terms of how they causally interact with objects. Since the objects are understood realistically, this yields an *externalist* understanding of conceptual content: the concept DOG just is what is nomenclally connected to doghood. Content externalism may not be the most intuitive way to account for the *intentionality* of human concepts, but it seems true that, in a realistic ontology, there are not too many alternatives to secure a credible meaning for the intimate relation between mental content and what it is 'about'. Identifying the one with the other is a radical way to remove a gap that would otherwise look quite unbridgeable. On the other hand, in doing so another gap opens, namely that between mental content and self-knowledge: if mental content is (partly) outside 'the head', then self-knowledge about such content appears to imply knowledge about certain contingent facts 'out there', which seems absurd (Boghossian, 1997; cf., however, Williamson, 1995). Whatever be one's appreciation of the gap between the mind and the world, transcendentalism at least has the advantage that there is no such gap. If the emergence of objects is brought about by mental processes, then genuine intentionality can be combined with an internalist understanding of conceptual content.

Nice though it may seem to have such a clear internal computational locus for conceptual content, it does give rise to a complication that is better not overlooked. Most facts that the representational theory of mind takes care of should come out true in image theory as well. For one thing: the activation of images *should* be caused (among other things) by instances of what they apply to, but as for the latter, Fodor (1998) points out that it would be a grave mistake to limit this causal connection to just one (specific) perceptual type of access:

The merest ripple in dog-infested waters may suffice to cause *dog*-thoughts in the theoretically sophisticated. (p. 78)

To be sure, it may be that all the non-perceptual mechanisms that sustain semantic access to *doghood* depend, 'in the long run' on one's having and exercising perceptual capacities. But not, according to the present view, on one's having any *particular* perceptual capacity (remember Helen Keller). (p. 79)

Clearly, image theory is vulnerable to criticism at this point. Images have been defined mostly in terms of perception ('scanning'), and in specific cases this will come down to the operation of specific perceptual capacities.

Helen Keller's³ image for dogs will differ in important respects from that of a seeing and hearing person, but this problem could be overcome. 'Higher up' in the procedure (where perceptual information is translated into information about structural properties) there will be much similarity after all. The remark about 'the merest ripple in dog-infested waters' is more reason for alarm. We know many things only by *inferring* them from information we can access directly. To know if a box from a store contains peanuts, we normally look at the label to find out, even though, epistemically speaking, it would make more sense to open the box and look inside. The label-image, which is activated first, must have ways of redirecting this information so as to activate the peanuts-image. But what *makes* the latter a peanuts-image is that it can run the procedure to find out about peanuts directly.

To make label-reading and like operations possible, it must be that the more advanced cognitive systems can deploy images in several *indirect* ways. Apart from scanning (direct processing of perceptual information), it must be possible to draw inferences, which comes down to using the output of other images. In the peanut case, there will be an appreciation of the intentionality of the maker of the box: why would someone adorn it with such a picture, if not to communicate the content? For less informed observers a mere associative bond might be sufficient: repeated concurrence of the label and peanuts could have caused this bond to be strengthened in time. Such alternative routes to knowledge cannot be packed in the image itself: they must work because of the embedding of the image in a cognitive system that knows how to pass by the obstacles to direct perceptual knowledge about the world. For the human race the overwhelming majority of known facts takes such bypasses. The world is something we know from being told about.

The practice of simply identifying the concept of the image by the way it directly (non inferentially) reacts to perceptual information, is further complicated by the fact that, for many images, this mode of operation is hardly, if at all, available. How do we come to believe that a golden ring is gold? Perception does play its part, but complex chains of reasoning, taking into account the price and the expertise and trustworthiness of the jeweller, do most of the work. Relying on extensional adequacy will not provide the final solution: if, for various reasons, our ways of recognizing gold are prone to error, does that make our image less an image for gold? It is highly likely that the 'real' image for gold—the one that actually plays that part in humans—is one enabling a procedure that, under reasonable circumstances, delivers a verdict on the being gold of some object, making use of every available source of information. If you see cube-shaped crystallization, be sceptical. If you are colour-blind, ask someone else who is not. If

³Helen Keller was an American woman who became deaf and blind when she was 19 months old, yet, thanks to extensive teaching, finished university, and became an author and an important intellectual voice.

you know an expert, ask that person.

Speaking of experts: here a new aspect of meaning enters the picture, for in image theory there appears to be room for meanings that ‘ain’t in the head’, if not quite in the way Putnam (1975) intended. To borrow Fodor’s (1994) notable example, consider the concept *elm*. Many of us who can distinguish between an oak and a beech will also have enough mastery of this concept to be forgiven for using this word, but there is far more to being an elm than what is in the head of the average citizen passing them by. This is where the *expert* comes in. Her knowledge can be viewed as an external extension to the citizen’s images, since it is, functionally, part of the procedure by which the elmhood of some object can be determined. This is all entirely acceptable and fully in line with the general ideas behind image theory. Moreover, there is every reason also to include *artificial* procedures—by observation and measuring devices—into the ‘extended image’. *Pace* Putnam, however, neither water, nor twin-water *itself* could ever be part of it: what is the target of the image does not itself partake in the procedure deciding whether it is there. Images, be they internal or external to the observer’s brain, are *independent* of the existence and size of their extension. This excludes Putnamian externalism, but leaves plenty of room for concepts that (far) transcend private thought.

Be all this as it may, the sobering fact is nevertheless that inference (the ability to infer as it would be needed in image theory) is not computationally straightforward—far from it! It might seem as though the output of image *i* being the input of image *j* would do the trick, but such an explanation would leave it a mystery how cognition discriminates between valid and invalid inferences. If the brain fails to establish the *right* connections, everything could follow from everything! Image theory stands in need of an account of inference: flexible, and with wide, if possible general, applicability. It should be able to explain why, say, the image for *being a white dog*, when activated, activates the image for *being white*. Appealing to the validity of $p \wedge q \rightarrow p$ will not do: the image for *being a white dog* is *not* the image for *being a dog* and *being white* connected by an AND-gate (cf. 6.1). The conclusion must be that lots of work remains to be done here.

7.5 Indeterminacy

Some dogs are white and some are not, but of many dogs we are not quite sure. Some people are tall and some are not, with many in between. In-between conditions, abounding in real life, are a challenge to realism, since realism suggests that things must be one way or another. Surely there is no problem with being neither tall nor short, but not being tall seems to imply being non-tall. If that is so, then, apart from being tall or being non-tall, there can be no third option: *tertium non datur*.

However, despite the near-inexorability of this principle logically, real life does not seem to support it: wherever there are two ‘complementary’ possibilities, the devil invariably contrives to squeeze a nondescript something in between (Van Rooij, 2011). In many cases this does not bring big logical challenges: whoever is neither so as to be called *tall*, nor so as to be called *non-tall*, can be medium-sized. Dogs can be pragmatically white, but in fact light grey, beige, or slightly spotted. But some in-between conditions are really hard to live with. It is here that I think image theory can do some illuminating work.

Take truth. I take truth because the paradoxes that this concept gives rise to need no more than a few simple steps to bring about. Here the role of procedures as being constitutive of the way objects appear to us as being thus and so (the predications we apply to them) makes it possible to make sense of the so-called ‘gaps’ between truth and falsehood (cf. also Moschovakis et al., 1994). Our pre-theoretic notion of truth goes by the name of *transparent truth*. It is an intuitive rule, if ever there was one:

Truth There is a property *being true*, such that, given any proposition p : p is true iff p .

Thus, if Snowy is white, then *that Snowy is white* is true (and vice versa). In a step-by-step description, the procedure could be thought of as following a Tarski-style recursive definition of truth (Tarski, 1933, 1944). Take the subject. If it is not a proposition, it is not true. If it is, parse the content of the proposition. This will lead to an analysis where the end-points are elementary predications (in first-order logic these are represented by *atomic sentences*). In the example the proposition is that Snowy is white. Since it is not truth-functionally compound, parsing the proposition just returns the same proposition, whereupon the procedure moves to the elementary predication as a sub-procedure. By the latter a test is performed on Snowy’s being white. If the outcome is positive, then it is this output which feeds the main procedure and makes it produce a positive outcome with respect to the truth of the sentence, after which the process terminates.

In this example the sub-procedure is perceptual, but this is accidental. If the proposition is *that what Susan says is true*, then the procedure scanning it will redirect to what Susan says, which is also a proposition. If what Susan says is *that Snowy is white*, then the procedure safely lands on a perceptual sub-procedure at last, and the result can be returned and processed to conclude that it is true that what Susan says is true. Now, in most everyday cases such evaluations, although they may pose lots of practical problems, do not suffer from any *fundamental* impediment. Either the procedure terminates, or it gets stranded somewhere, due to lack of available information. The fundamental problem with a predicate like truth is not that it may encounter circumstances in which it gets stuck and fails to deliver an outcome—no predicate is perfect. The trouble is that, for a couple of well specified cases—Liar sentences and their relatives—the stalemate

is not a matter of bad luck. For these cases we know for a fact that no conclusion can *possibly* be reached. In a Liar sentence:

91. *l*: *l* is false

the search for truth goes on forever. To find out whether *l* is true, the procedure first has to run the sub-procedure testing its content, i.e. *that l is false*. Parsing this content yields *l* again, and the procedure is caught in a loop: it will never return a signal to start the way back to the answer (cf. Gupta and Belnap, 1993).

The paradoxicality of Liar sentences is as baffling as it is, in part because its insolubility is seen through immediately. Let us, however, for a moment imagine that this were not so. Suppose we were *ignorant* about the loop and did not know why we could not find an answer. Many wise people have studied *l*, but no one ever found out whether it is true or false. But this (so the story continues) has not struck anyone as being a problem for the concept of truth *itself*. Of innumerable propositions we do not know if they are true or false—what hubris to think otherwise! To most of us, this solution will look extremely unsatisfactory. To accept such a diagnosis, we just know too much: that there cannot be an outcome is not because of any human limitation, but due to simple facts about the sentence and the concepts it exploits. Ignorance concerns matters where there *is* something to know, even though we do not. I nevertheless believe that, for a final diagnosis of Liar sentences and the concept of truth, comparison with ignorance is more fruitful than is normally thought.

I shall return to that later. A common reaction to the problem with transparent truth is to *adjust the definition* of truth, e.g. by adding a new truth-value in between truth and falsehood. This could be implemented in the following way. For the result of a truth-evaluation there are three options: output *true*, output *false*, and no output. The strategy we will examine here is that whereby ‘no output’ will be regarded as output *neutral*, thereby embracing a three-valued logic (eg. Kleene, 1938). Now *l* will turn out to be *neutral*.

From the viewpoint of image theory, this strategy is problematic, since it involves replacing the original procedure for truth by *another* procedure. As a result of this, what is being tested is no longer *being true*, but *another* property, say *being-true'*. It is simply not possible to ‘adjust’ any procedure in this way and still make it scan for the same property. The only justification for this would be that transparent truth were incorrect as a definition of truth: that the procedure were somehow defective. The ground for believing this to be so, however, cannot be that it gives the wrong outcomes when it gives any. It must be its proneness to paradox. What we expect from *being-true'* as an alternative for *being true* is that its outcomes are equally correct, and that there is always one. *The gap must be plugged*, and since it is obvious that neither *true* nor *false* will do for a plug, something else must. Hence, the principle of bivalence must be given up.

However, the task to find a plug that closes the gap without leakage has proven Sisyphean. It is not that no solution can be given which provides a sound

evaluation for l in terms of truth (*neutral*), but whenever the language has been enriched so as to make this possible, a *new* condition can be formulated which gives rise to new paradoxical sentences, known as ‘revenge Liars’. The recipe for inventing new paradoxes is embarrassingly simple: bring back bivalence. Although there is no objection whatsoever against procedures with more than two possible outcomes (say P , Q , and R), they can always be recast as binary procedures. The said procedure goes proxy for *three* procedures: with outcomes P vs \overline{P} , Q vs \overline{Q} , and R vs \overline{R} . So any version of *being-true*’, with outcomes *true*, *false*, and *neutral*, is at the same time a binary procedure with outcomes *true* and $\overline{\text{true}}$. But now we have the following:

92. l' : l' is $\overline{\text{true}}$

The proposition l' is an example of a revenge Liar: a slightly altered version of the Liar to evade a solution to the original one. If l' is *true* it must be $\overline{\text{true}}$, thus *false* or *neutral*, but if so, then it is $\overline{\text{true}}$, and therefore *true*. But it cannot be both, and down comes the rock!

The method normally chosen to block revenge—to make the rock stay in place—is to constrain the vocabulary, whereby predicates like $\overline{\text{true}}$ are banned from the language, at least from what since Tarski’s work on the semantics of truth has come to be known as the *object language* (a more sophisticated version of this solution is the Tarskian hierarchy of languages). Yet, the lingering feeling of dissatisfaction with solutions of this type is hardly surprising: there is little reason to believe that the (formal) trimming of language in order to prevent mentioning conditions like $\overline{\text{true}}$ has a relevant correspondence to any such bridling in natural language or real life. This feeling is vindicated by image theory: the image for $\overline{\text{true}}$ and like predicates is obtained by only a minimal modification of the procedure for *being-true*’. All that has to be done is to bundle some of the outcomes together by an OR-gate. Mastery of the new concept is really next door to that of the old ones.

For a diagnosis of Liar sentences from the viewpoint of image theory the central observation is that the idea of a procedure for *being-true*’ rests on a questionable assumption, which is that *neutral* can be an outcome in the first place. Let \mathcal{T} be the procedure for *being true* and \mathcal{T}' that for *being-true*’. Then the rules for \mathcal{T}' could be described as follows:

- Whenever \mathcal{T} would return *true*, return *true*.
- Whenever \mathcal{T} would return *false*, return *false*.
- Whenever \mathcal{T} would not halt, return *neutral*.

The implementation of the last rule, however, would require that \mathcal{T}' have a way of telling whether \mathcal{T} halts or not. But deciding whether a procedure halts or not (as a matter of necessity) is a fundamentally different type of ‘scanning’

than the examples we have seen before, which were all, in the end, perceptual tests. Such tests can fail for various practical reasons (bad lightning, failing sense organs, distance), but there is always the *in principle* possibility that it be done properly. With respect to the—predicted, not actual!—behaviour of an algorithm there is no such warrant. In 1936 Alan Turing famously published his proof that *the halting problem* has no solution: there is no general test to determine whether a given algorithm will terminate or not (Turing, 1936).

This may seem irrelevant, since no one doubts that \mathcal{T} will not halt on l . Consider, however, whether \mathcal{T} will halt on l' . Suppose it does not, then, by the rules of \mathcal{T}' , \mathcal{T}' must return *neutral* when run on l' . But this means that l' is *true*, and if that is so, then \mathcal{T} halts on l' after all: contradiction. So let us, alternatively, suppose that \mathcal{T} halts on l' . Given the rules for \mathcal{T}' , this makes l' either *true* or *false*, which, as we have seen, leads to contradiction as well. So does \mathcal{T} halt on l' ? It looks like the problem is with the predicate $\overline{\text{true}}$, and therefore, with \mathcal{T}' . Unlike \mathcal{T} , which may sometimes fail to return an output, but is an entirely feasible procedure, \mathcal{T}' is *impossible to implement*. As a consequence, three-valued logic (in this sense) is not available to image theory. There is no *being-true'*.

What moral should be drawn? If, as image theory states, properties must be associated with cognitive procedures, then ‘gaps’ between being and not being P are just a fact of life. Some procedures do not halt, so a verdict is never reached. That there *must* nevertheless be a fact of the matter with respect to P —that something *has to be the case out there*—is a realist presumption that the transcendentalist need not embrace (a conclusion very much in line with Dummett’s view in realism (Dummett, 1982, 1991)). Most importantly, it is not possible to turn the gap into an alternative outcome. No outcome really is no outcome! If a procedure fails to stop, to interpret *that* as a stop on some meta-level does not work—all we can do is leave it there.

Sometimes there is no outcome—connoisseurs of the Liar and his gang smell danger at such a seemingly reassuring contention. For saying that there is no outcome *also* seems to be drawing a conclusion, and *every* conclusion with respect to Liar-type propositions can be routinely shredded by a suitable version of revenge. Such as:

93. l'' : l'' is either false or there is no outcome.

The quintessential point, however, seems to be that ‘there is no outcome’ is not a statement about the world like ‘Snowy is white’. Statements about the predicted behaviour of procedures are *counterfactual* statements: if the procedure were to run for an indefinite amount of time, then it would or would not stop. Although counterfactuals are often seen as (to paraphrase L.P. Hartley) ‘a foreign country, where they do things differently’, we may as well have to face the possibility that certain counterfactual statements just do not allow for a final evaluation. There is no reality (not even understood transcendently) corresponding to them, just like there is no reality corresponding to a statement like ‘It is unknown whether

Snowy is white' (as opposed to 'Snowy is (not) white'). Here is the similarity with lack of knowledge that I mentioned above. Some statements lack the objectivity to be evaluable in the same way that 'Snowy is white' is. The second disjunct in (93) appears to belong to this group (cf. also Wittgenstein's (2010) discussion of exclamations).

All this looks like an encouragement to view transparent truth as being all truth there is. It does not lead to contradiction: the only way to derive a contradiction from the Liar is to start with the assumption that every proposition is either true or false, i.e. that the procedure for *being true* always returns an outcome—which it evidently does not. How to modify classical logic to handle these cases is an exercise that I leave to the reader.

7.6 Causality

If anyone should feel disappointment that the concept of *causality* has not received due attention in this text, then what is to blame is my failure, not to understand its significance, but to understand the phenomenon as such well enough to provide some account of it terms of image theory. Causality is a massively important element in human life, but also one whose explanation in more fundamental terms is notoriously difficult. All the same it is hardly conceivable that image theory should have nothing to say about causality, for it is in the class of exactly those phenomena that look extremely real, yet show distinctive traces of being human-perspective-laden.

Hence, my modest objective will be to provide some suggestions as to whether (and to what extent) causality should be understood as a principle about how the world really is or as an aspect of how humans cognitively connect to it. For the transcendentalist is not trivially committed to the second. By way of clarification: that there is space and matter, and that there are regularities in the world, I take to be unmistakable aspects of the world as it is in itself, even though we become conscious of them in ways that are specific for our type of intelligence. Beings with a different kind of cognitive make-up might experience these things differently, but in some way or other they would be there for them, too. They are not the phenomena for whose existence *as such* image theory could give an explanation. But that we experience there to be both concrete and abstract objects, and also view the world in terms of propositions and facts, appear to be features of our cognitive system that other types of intelligence might totally lack. It would be interesting to know where causality fits in the picture.

The idea that causality is a projection of human experience is not new. Russell spoke of 'a relic of a bygone age' (Russell et al., 1992, p.193), and some notable reasons can be quoted to see it that way. Although most branches of empirical science are not in the least inclined to question the objective reality of cause and effect, with the coming of modern physics a downright causal account of the

observed phenomena has become harder to maintain. It is maybe too much to say that causality on the micro-level has been proved incompatible with quantum theory, but this could also be due to the inherent vagueness in the concept itself. As Bell's theorem famously implies⁴, if quantum physics is to be understood in terms of the causal interaction between (hidden) variables, then it must be non-local (i.e. it must be allowed to have effect on indefinite distances), but then, due to special relativity, it is no longer assured that causes precede their effects in time. Yet, the latter rule, *temporal asymmetry*, I take to be among the most unnegotiable intuitions about causality⁵. Although all this has very little effect on the goings-on on the visible level, it does cast doubts on the *fundamentality* of causation as a feature of the world as it is 'in itself'⁶.

Of course it is possible to view the intuitive concept as an general purpose rule-of-thumb, based on a 'deeper' fact about the way the world is put together: a scientific principle of causality that may violate these principles (albeit only in the fringe, perhaps). Without taking a stance on these matters, I believe that, for present purposes, it makes most sense to view intuitive causality as the starting point: causality as it is found in the draft ontology. Then the observation appears to be that there are many recurrent patterns of correlation in the world, but only some of them we pick out as being cause and effect.

Which ones, and why? In an acclaimed paper, Michael Dummett argues that the origin of how past and future differ as regards our ability to change the course of events is epistemological:

The difference between past and future lies in this: that we think that, of any past event, it is in principle possible for me to know whether or not it took place independently of my present intentions; whereas, for many types of future event, we should admit that we are never going to be in a position to have such knowledge independently of our intentions. (Dummett, 1964)

Independently of our intentions—this is a core notion. The human mind is more than just a cognitive system: it also has the ability to execute intentions, i.e. the ability to *act*. What Dummett describes is an intriguing interplay between this and the ability to know: there is a sense in which they are mutually exclusive. If matters in some domain are 'fixed', then we cannot act to change them, and if we can, they are not fixed. Now 'to be fixed' does seem to have a connection with being knowable. Dummett argues that if the condition of being fixed in this

⁴Bell's theorem (Bell, 1964) states that no theory about 'hidden variables' displaying local causality can reproduce the predictions of quantum-mechanics (in particular those pertaining to entangled particles).

⁵This concerns *intuitions*: in theoretical physics backward causation is sometimes considered as a possibility (Price, 2001; Argaman, 2010).

⁶There are other observations sustaining this doubt (Chiribella et al., 2013; Brukner, 2014).

sense is lifted from some events in the past, then backward causation becomes less absurd.

If causality is in this way related to our ability to know vs to act (rather than the direction of time), then this would seem to support *manipulability* accounts of causality (Menzies and Price, 1993; Woodward, 2003; Pearl, 2009). As a short characterization of this approach, Woodward (2016) writes:

[V]ery roughly, if C is genuinely a cause of E , then if I can manipulate C in the right way, this should be a way of manipulating or changing E .

Despite philosophical qualms, this account appears very successful in picking out the right regularities as being causal in nature. It also provides a natural explanation for why knowledge about causal relations is important knowledge to have, and also why experimenting is the method of choice to find out about them. Admittedly, it takes some willingness to consider wildly counterfactual conditions, but even here the result is hardly sophistic: indeed, *had* someone changed the course of that old meteor just in time, the dinosaurs might still be around, might they not?

Manipulability theories look anthropocentric—hence the philosophical qualms. But here one could reply that the perception of causal connections is not the perception of things that were not there before: it is to *label* certain patterns in the world as belonging to that class. And of course the transcendentalist has been here before. Saying that causality is part of the human perspective is not saying that Kennedy’s death was *not* really caused by Oswald’s bullets, any more than saying that objects are part of that perspective is saying that there is really no Eiffel Tower. Maybe the gods would see no causes, but they would see Oswald fire and Kennedy fall.

Consequently, what seems to be needed to get any grip on the phenomenon of causality, is a far better understanding of the phenomenon of *action* itself—a phenomenon that itself has been nominated for being observer-dependent (e.g. Dennett, 1989). The exact role of image theory in this project remains to be established. It is not hard to describe bodily movements as being the result of computational goings-on in the brain, but only part of this counts as action—the *intentional* part, and which part is that? My expectation would be that, once we have a good account of action in terms of the human cognitive system, action and causality will turn out to be two sides of one coin.

Bibliography

- Aguilar-Guevara, A. and Zwarts, J. Weak definites and reference to kinds. In *Semantics and Linguistic Theory*, volume 20, pages 179–196, 2010.
- Argaman, N. Bell’s theorem and the causal arrow of time. *American Journal of Physics*, 78(10):1007–1013, 2010.
- Armstrong, D. M. *A Theory of Universals: Volume 2: Universals and Scientific Realism*, volume 2. CUP Archive, 1980.
- Armstrong, D. M. *Universals: An opinionated introduction*. Westview Press, 1989.
- Armstrong, D. M. A World of States of Affairs. *Philosophical Perspectives*, 7: 429–440, January 1993.
- Armstrong, D. M. *A world of states of affairs*. Cambridge Univ Press, 1997.
- Armstrong, D. M., Bacon, J., Campbell, K., and Reinhardt, L. *Ontology, Causality, and Mind*. Cambridge University Press, March 1993.
- Audi, P. Grounding: Toward a theory of the in-virtue-of relation. *The Journal of Philosophy*, 109(12):685–711, 2012.
- Bell, J. S. On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1 (3):195, 1964.
- Benacerraf, P. What numbers could not be. *The Philosophical Review*, 74(1): 47–73, 1965.
- Berto, F. Taming the runaway imagination ticket. *Synthese*, pages 1–15, 2018.
- Berto, F. and Jago, M. *Impossible worlds*. Oxford University Press, 2019.

- Berto, F. and Plebani, M. *Ontology and metaontology: A contemporary guide*. Bloomsbury Publishing, 2015.
- Berto, F. and Tagliabue, J. The world is either digital or analogue. *Synthese*, 191(3):481–497, 2014.
- Black, M. The identity of indiscernibles. *Mind*, 61(242):153–164, 1952.
- Boghossian, P. A. What the Externalist Can Know A Priori. In *Proceedings of the Aristotelian Society*, volume 97, pages 161–176. Oxford University Press Oxford, UK, 1997.
- Boolos, G. To be is to be a value of a variable (or to be some values of some variables). *The Journal of Philosophy*, 81(8):430–449, 1984.
- Bradley, F. H. *Appearance and reality: a metaphysical essay*. Routledge, 2016.
- Brukner, Č. Quantum causality. *Nature Physics*, 10(4):259–263, 2014.
- Campbell, K. The metaphysic of abstract particulars. *Midwest studies in philosophy*, 6(1):477–488, 1981.
- Caplan, B. Creatures of fiction, myth, and imagination. *American Philosophical Quarterly*, 41(4):331–337, 2004.
- Carlson, G. N. and Pelletier, F. J. *The generic book*. University of Chicago Press, 1995.
- Cavina-Pratesi, C., Kentridge, R. W., Heywood, C. A., and Milner, A. D. Separate channels for processing form, texture, and color: evidence from fMRI adaptation and visual object agnosia. *Cerebral cortex*, 20(10):2319–2332, 2010.
- Chalmers, D. J. *How can we construct a science of consciousness?* MIT Press, 2004.
- Chiribella, G., D’Ariano, G. M., Perinotti, P., and Valiron, B. Quantum computations without definite causal structure. *Physical Review A*, 88(2), 2013.
- Correia, F. *Existential dependence and cognate notions*. Philosophia Munich, 2005.
- Davidson, D. On the very idea of a conceptual scheme. *The American Philosophical Association Centennial Series*, pages 209–222, 2013.
- De Waal, F. B. M. Intentional deception in primates. *Evolutionary Anthropology: Issues, News, and Reviews*, 1(3):86–92, 1992.

- Denkel, A. On the compresence of tropes. *Philosophy and Phenomenological Research*, 57(3):599–606, 1997.
- Dennett, D. C. *The Intentional Stance*. MIT Press, 1989.
- Dennett, D. C. Real patterns. *The Journal of Philosophy*, 88(1):27–51, 1991a.
- Dennett, D. C. *Consciousness explained*. Penguin UK, 1991b.
- Dorr, C. *The simplicity of everything*. PhD thesis, Princeton University, 2002.
- Dummett, M. Bringing about the past. *The Philosophical Review*, pages 338–359, 1964.
- Dummett, M. Realism. *Synthese*, pages 55–112, 1982.
- Dummett, M. *The logical basis of metaphysics*. Harvard university press, 1991.
- Duží, M., Jespersen, B., and Materna, P. *Procedural semantics for hyperintensional logic: foundations and applications of transparent intensional logic*, volume 17. Springer Science & Business Media, 2010.
- Evans, R., Hernández-Orallo, J., Welbl, J., Kohli, P., and Sergot, M. Making sense of sensory input. *arXiv preprint arXiv:1910.02227*, 2019.
- Evans, R., Hernández-Orallo, J., Welbl, J., Kohli, P., and Sergot, M. Evaluating the Apperception Engine. *arXiv preprint arXiv:2007.05367*, 2020.
- Fine, K. Essence and modality: The second philosophical perspectives lecture. *Philosophical perspectives*, 8:1–16, 1994.
- Fine, K. The non-identity of a material thing and its matter. *Mind*, 112(446): 195–234, 2003.
- Fodor, J. A. *The language of thought*, volume 5. Harvard university press, 1975.
- Fodor, J. A. *Representations: Philosophical essays on the foundations of cognitive science*. Harvester Press Brighton, 1981.
- Fodor, J. A. *The modularity of mind*. MIT press, 1983.
- Fodor, J. A. *The elm and the expert: Mentalese and its semantics*. MIT press, 1994.
- Fodor, J. A. *Concepts: Where cognitive science went wrong*. Oxford University Press, 1998.
- Frege, G. Über begriff und gegenstand. *Vierteljahrsschrift für wissenschaftliche Philosophie*, pages 1–11, 1892a.

- Frege, G. Über sinn und bedeutung. *Zeitschrift für Philosophie und philosophische Kritik*, 100:25–50, 1892b.
- Frege, G. Der gedanke: Eine logische untersuchung. *Beiträge zur Philosophie des deutschen Idealismus 1 (1918/19)*, pages 58–77, 1918.
- Frege, G. On concept and object. *Mind*, 60(238):168–180, 1951.
- Frege, G. *The foundations of arithmetic: A logico-mathematical enquiry into the concept of number*. Northwestern University Press, 1980.
- French, R. An Argument for the Ontological Innocence of Mereology. *Erkenntnis*, 81(4):683–704, 2016.
- French, S. *The structure of the world: Metaphysics and representation*. Oxford University Press, 2014.
- Frost-Arnold, G. Make ontology easy again. *Metascience*, 25(3):497–500, 2016.
- Geach, P. T. Identity. *The Review of Metaphysics*, 21(1):3–12, September 1967.
- Geach, P. T. *Reference and generality: An examination of some medieval and modern theories*. Cornell University Press, 1980.
- Goodman, N. The new riddle of induction. In *Fact, Fiction and Forecast*, chapter III. Bobbs-Merrill, second edition, 1965.
- Goodman, N. and Quine, W. V. O. Steps toward a constructive nominalism. *The journal of symbolic logic*, 12(4):105–122, 1947.
- Governatori, G., Maher, M. J., Antoniou, G., and Billington, D. Argumentation semantics for defeasible logic. *Journal of Logic and Computation*, 14(5):675–702, 2004.
- Gupta, A. and Belnap, N. *The revision theory of truth*. Mit Press, 1993.
- Hare, B., Call, J., and Tomasello, M. Chimpanzees deceive a human competitor by hiding. *Cognition*, 101(3):495–514, 2006.
- Hawley, K. Mereology, modality and magic. *Australasian Journal of Philosophy*, 88(1):117–133, 2010.
- Heller, M. Property counterparts in ersatz worlds. *The Journal of Philosophy*, 95(6):293–316, 1998.
- Hempel, C. G. Studies in the Logic of Confirmation (I). *Mind*, 54(213):1–26, 1945.

- Herre, H. The ontology of mereological systems: A logical approach. In *Theory and Applications of Ontology: Philosophical Perspectives*, pages 57–82. Springer, 2010.
- Hofweber, T. *Ontology and the Ambitions of Metaphysics*. Oxford University Press, 2016.
- Hofweber, T. Idealism and the harmony of thought and reality. *Mind*, 128(511): 699–734, 2019.
- Hooker, C. A. The relational doctrines of space and time. *The British Journal for the Philosophy of Science*, 22(2):97–130, 1971.
- Hovda, P. What is classical mereology? *Journal of Philosophical Logic*, 38(1): 55–82, 2009.
- Huang, C.-R. and Ahrens, K. Individuals, kinds and events: classifier coercion of nouns. *Language Sciences*, 25(4):353–373, 2003.
- Hume, D. *A treatise of human nature*. Courier Corporation, 2003.
- Jago, M. Constructing worlds. *Synthese*, 189(1):59–74, 2012.
- Johnston, M. and Forbes, G. Is there a Problem about Persistence? *Proceedings of the Aristotelian Society, Supplementary Volumes*, 61:107–155, 1987.
- Jubien, M. *Ontology, Modality, and the Fallacy of Reference*. Cambridge University Press, 1993.
- Kalhat, J. Structural universals and the principle of uniqueness of composition. *Grazer Philosophische Studien*, 76(1):57–77, 2008.
- Kamp, H. Two theories about adjectives. In *Meaning and the Dynamics of Interpretation*, pages 225–261. Brill, 2013.
- Kant, I. *Critique of pure reason*. Everyman, 1993.
- Kappes, H. B. and Morewedge, C. K. Mental simulation as substitute for experience. *Social and Personality Psychology Compass*, 10(7):405–420, 2016.
- Kind, A. Imagining under constraints. *Knowledge through imagination*, pages 145–59, 2016.
- Kleene, S. C. On notation for ordinal numbers. *The Journal of Symbolic Logic*, 3(4):150–155, 1938.
- Kolmogorov, A. N. On tables of random numbers. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 369–376, 1963.

- Kolmogorov, A. N. Three approaches to the quantitative definition of information. *Problems of information transmission*, 1(1):1–7, 1965.
- Kripke, S. A. *Naming and necessity*. Springer, 1972.
- Ladyman, J. and Ross, D. *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press, 2009.
- Laycock, H. Some questions of ontology. *The philosophical review*, 81(1):3–42, 1972.
- Lewis, D. *On the plurality of worlds*, volume 322. Oxford Blackwell, 1986a.
- Lewis, D. Against structural universals. *Australasian Journal of Philosophy*, 64(1):25–46, 1986b.
- Lewis, D. *Philosophical Papers: Volume II*. Oxford University Press, 1987.
- Lewis, D. *Parts of Classes*. Blackwell, Oxford, 1991.
- Lewis, D. Humean supervenience debugged. *Mind*, 103(412):473–490, 1994.
- Locke, J. *An Essay Concerning Human Understanding (First published 1690)*. Philadelphia, PA: The Pennsylvania State University, 1999.
- Maher, M. J., Rock, A., Antoniou, G., Billington, D., and Miller, T. Efficient defeasible reasoning systems. *International Journal on Artificial Intelligence Tools*, 10(04):483–501, 2001.
- Maher, P. Probability captures the logic of scientific confirmation. *Contemporary debates in philosophy of science*, pages 69–93, 2004.
- Marcus, G. F. *The algebraic mind: Integrating connectionism and cognitive science*. MIT press, 2001.
- Maurin, A. Trope theory and the Bradley regress. *Synthese*, 175(3):311–326, 2010.
- McDaniel, K. Against composition as identity. *Analysis*, 68(298):128–133, 2008.
- McDaniel, K. *The fragmentation of being*. Oxford University Press, 2017.
- McDermott, D. and Doyle, J. Non-monotonic logic I. *Artificial intelligence*, 13(1-2):41–72, 1980.
- McKay, T. and MacKay, C. *Plural predication*. Oxford University Press on Demand, 2006.

- Menzies, P. and Price, H. Causation as a secondary quality. *The British Journal for the Philosophy of Science*, 44(2):187–203, 1993.
- Moltmann, F. Properties and kinds of tropes: New linguistic facts and old philosophical insights. *Mind*, 113(449):1–41, 2004.
- Mormann, T. On the mereological structure of complex states of affairs. *Synthese*, 187(2):403–418, 2012.
- Morzycki, M. *Modification*. Cambridge University Press, 2016.
- Moschovakis, Y. N. A logical calculus of meaning and synonymy. *Linguistics and Philosophy*, 29(1):27–89, 2006.
- Moschovakis, Y. N. et al. Sense and denotation as algorithm and value. *Lecture notes in logic*, 2:210–249, 1994.
- Nagel, T. What is it like to be a bat? *The philosophical review*, 83(4):435–450, 1974.
- Newman, A. A metaphysical introduction to a relational theory of space. *The Philosophical Quarterly (1950-)*, 39(155):200–220, 1989.
- Nickel, B. Generics and the ways of normality. *Linguistics and Philosophy*, 31(6):629–648, 2008.
- Nicolas, D. Mass nouns and plural logic. *Linguistics and philosophy*, 31(2):211–244, 2008.
- OED. Online Etymology Dictionary, July 2020. URL <https://www.etymonline.com>.
- O’Leary-Hawthorne, J. The bundle theory of substance and the identity of indiscernibles. *Analysis*, 55(3):191–196, 1995.
- Parsons, T. *Nonexistent objects*. New Haven: Yale University Press, 1980.
- Partee, B. H. Compositionality and coercion in semantics: The dynamics of adjective meaning. *Cognitive foundations of interpretation*, pages 145–161, 2007.
- Pearl, J. *Causality*. Cambridge university press, 2009.
- Poincaré, H. *Science and hypothesis*. Science Press, 1905.
- Pollock, J. L. Defeasible reasoning. *Cognitive science*, 11(4):481–518, 1987.
- Price, H. Backward causation, hidden variables and the meaning of completeness. *Pramana*, 56(2-3):199–209, 2001.

- Priest, G. Thinking the impossible. *Philosophical Studies*, 173(10):2649–2662, 2016.
- Putnam, H. The Meaning of Meaning. In *Mind, Language and Reality*, volume 2 of *Philosophical Papers*, pages 215–271. Cambridge University Press, Cambridge, 1975.
- Putnam, H. Psychological predicates. *Art, mind, and religion*, 1:37–48, 1967.
- Quine, W. V. O. On what there is. *The Review of Metaphysics*, pages 21–38, 1948.
- Quine, W. V. O. Natural kinds. In *Essays in honor of Carl G. Hempel*, pages 5–23. Springer, 1969.
- Quine, W. V. O. *From a logical point of view: 9 logico-philosophical essays*, volume 9. Harvard University Press, 1980.
- Quine, W. V. O. *Philosophy of logic*. Harvard University Press, 1986.
- Quine, W. V. O. *Word and object*. MIT press, 2013.
- Ramachandran, V. S. and Gregory, R. L. Perceptual filling in of artificially induced scotomas in human vision. *Nature*, 350(6320):699–702, 1991.
- Ramsey, F. P. General Propositions and Causality. In Braithwaite, R. B., editor, *The Foundations of Mathematics and other Logical Essays*, pages 237–257. Routledge, 2000.
- Rea, M. C. Supervenience and co-location. *American Philosophical Quarterly*, 34(3):367–375, 1997.
- Rosen, G. Metaphysical dependence: Grounding and reduction. *Modality: Metaphysics, logic, and epistemology*, pages 109–36, 2010.
- Rosen, G. Abstract Objects. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2018 edition, 2018.
- Rosen, G. and Dorr, C. Composition as fiction. In Gale, R., editor, *The Blackwell Guide to Metaphysics*, pages 151–174. Blackwell, 2002.
- Russell, B. *The relation of sense-data to physics*. Nicola Zanichelli, 1914.
- Russell, B. *The philosophy of logical atomism*. Routledge, 2009.
- Russell, B., Slater, J. G., and Frohmann, B. *Logical and Philosophical Papers, 1909-13*, volume 6. Psychology Press, 1992.

- Sainsbury, R. M. *Fiction and fictionalism*. Routledge, 2009.
- Saunders, S. Are quantum particles objects? *Analysis*, 66(1):52–63, 2006.
- Schaffer, J. The individuation of tropes. *Australasian Journal of Philosophy*, 79(2):247–257, 2001.
- Schaffer, J. Two conceptions of sparse properties. *Pacific Philosophical Quarterly*, 85(1):92–102, 2004.
- Schaffer, J. Monism: The priority of the whole. *Philosophical Review*, 119(1):31–76, 2010.
- Schaffer, J. Grounding in the image of causation. *Philosophical studies*, 173(1):49–100, 2016.
- Schiffer, S. R. *The things we mean*. Oxford University Press, 2003.
- Sider, T. The ersatz pluriverse. *The Journal of Philosophy*, 99(6):279–315, 2002.
- Sider, T. Bare particulars. *Philosophical perspectives*, 20:387–397, 2006.
- Sider, T. Parthood. *The Philosophical Review*, 116(1):51–91, 2007.
- Sider, T. *Writing the Book of the World*. OUP Oxford, 2013a.
- Sider, T. Against Parthood. *Oxford studies in metaphysics*, 8:237, 2013b.
- Sommers, F. *The logic of natural language*. Clarendon Press Oxford, 1982.
- Stalnaker, R. C. A theory of conditionals. In Rescher, N., editor, *Studies in logical theory*, pages 98–112. Oxford: Blackwell, 1968.
- Stalnaker, R. C. *Ways a world might be: Metaphysical and anti-metaphysical essays*. Oxford University Press, 2003.
- Strawson, P. F. *Individuals*. Routledge, 2002.
- Swinburne, R. Thisness. *Australasian Journal of Philosophy*, 73(3):389–400, 1995.
- Tarski, A. *Pojecie prawdy w językach nauk dedukcyjnych*. Nakładem Towarzystwa Naukowego Warszawskiego, 1933.
- Tarski, A. The semantic conception of truth: and the foundations of semantics. *Philosophy and phenomenological research*, 4(3):341–376, 1944.
- Thomasson, A. L. Fictionalism versus deflationism. *Mind*, 122(488):1023–1051, 2013.

- Thomasson, A. L. *Ontology made easy*. Oxford University Press, 2014.
- Toledo, A. and Sassoon, G. W. Absolute vs. relative adjectives-variance within vs. between individuals. In *Semantics and linguistic theory*, volume 21, pages 135–154, 2011.
- Turing, A. M. On computable numbers, with an application to the Entscheidungsproblem. *J. of Math*, 58(345-363):5, 1936.
- Turner, J. Ontological pluralism. *The Journal of Philosophy*, 107(1):5–34, 2010.
- Turner, J. Curbing enthusiasm about grounding. 2016.
- Uzquiano, G. Plurals and simples. *The Monist*, 87(3):429–451, 2004.
- Van Cleve, J. Three versions of the bundle theory. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 47(1):95–107, 1985.
- Van Inwagen, P. Creatures of fiction. *American philosophical quarterly*, 14(4): 299–308, 1977.
- Van Inwagen, P. Composition as identity. *Philosophical Perspectives*, 8:207–220, 1994.
- Van Inwagen, P. *Material beings*. Cornell University Press, 1995.
- Van Inwagen, P. Quantification and fictional discourse. In *Empty Names, Fiction, and the Puzzles of Non-Existence*, pages 235–247. A. Everett & T. Hofweber, 2000.
- Van Rooij, R. Vagueness and linguistics. In *Vagueness: A guide*, pages 123–170. Springer, 2011.
- Varzi, A. C. Mereological commitments. *dialectica*, 54(4):283–305, 2000.
- Varzi, A. C. The extensionality of parthood and composition. *The Philosophical Quarterly*, 58(230):108–133, 2008.
- Vetter, B. Recent work: Modality without possible worlds. *Analysis*, 71(4): 742–754, 2011.
- Vetter, B. *Potentiality: From dispositions to modality*. OUP Oxford, 2015.
- Wetzel, L. *Types and tokens*. MIT Press, 2009.
- Wiggins, D. *Sameness and substance renewed*. Cambridge University Press, 2001.

- Williams, D. C. On the elements of being: I. *The review of metaphysics*, pages 3–18, 1953.
- Williams, J. R. G. Illusions of gunk. *Philosophical Perspectives*, 20(1):493–513, 2006.
- Williamson, T. Is knowing a state of mind? *Mind*, 104(415):533–565, 1995.
- Williamson, T. Knowing by imagining. *Knowledge through imagination*, pages 113–23, 2016.
- Wittgenstein, L. *Philosophical investigations*. John Wiley & Sons, 2010.
- Wollheim, R. *Art and its Objects*. Cambridge University Press, 1980.
- Woodward, J. *Making things happen: A theory of causal explanation*. Oxford university press, 2003.
- Woodward, J. Causation and Manipulability. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edition, 2016.
- Worrall, J. Structural realism: The best of both worlds? *Dialectica*, 43(1-2): 99–124, 1989.
- Yi, B.-U. Is mereology ontologically innocent? *Philosophical Studies*, 93(2): 141–160, 1999.
- Zalta, E. *Abstract objects: An introduction to axiomatic metaphysics*, volume 160. Springer Science & Business Media, 2012.
- Zeki, S. The visual image in mind and brain. *Scientific American*, 267(3):68–77, 1992.

Samenvatting

De dingen voor ons: Over wat object-zijn is

Wat is een object? In dit proefschrift wordt dit schijnbaar triviale begrip aan een onderzoek onderworpen. Geïnspireerd door inzichten van Immanuel Kant krijgt het begrip ‘object’ een *transcendentalistische* behandeling, waarvan het kernidee is dat het beschouwen van de gehele werkelijkheid in termen van objecten, die eigenschappen hebben en op bepaalde manieren gerelateerd zijn, een *menselijke* manier van begrijpen is, inherent aan ons cognitieve systeem.

In hoofdstuk 1 en 2 wordt de kenvermogen-afhankelijkheid van objecten verdedigd met behulp van het begrip *geworteld zijn* (‘rootedness’: het gegrond zijn van een object in andere objecten). Betoogd wordt dat de *eerste-orde logica* moet worden gezien als uitdrukking van de belangrijkste principes die de menselijke cognitie beheersen. Er wordt een kort overzicht gegeven van wat ik de *voorlopige ontologie* (‘draft ontology’) noem, ons intuïtieve beeld van de wereld, dat deze principes beogen te beschrijven. Er wordt verdedigd dat de studie van het menselijk cognitief systeem een essentieel onderdeel is van de methodologie om te komen tot een meer systematische ontologie. Meer dan als formele calculus wordt de eerste-orde logica—of liever gezegd, een uitbreiding daarvan die de status van ‘object’ verleent aan een grotere verscheidenheid van symbolen—gebruikt als een geïnterpreteerd symbolisme: een (niet-ideale) weergave van de voorlopige ontologie. Wat betreft de realiteit achter dit symbolisme wordt betoogd dat het menselijke venster op objecten een samenspel is van denken en waarnemen, waarvan de taal getuigt. Gewapend met dit conceptuele kader worden enkele fundamentele begrippen opnieuw bekeken: *identiteit* (zowel in de zin van ‘dit-heid’ als ‘gelijkheid’) en *meervoudigheid*.

In hoofdstuk 3 worden de belangrijkste gereedschappen voor het verklarende werk ontwikkeld. Het hoofdstuk begint met een bespreking van Bradley’s regressie, een berucht probleem in de ontologie. Na bestudering van enkele voorgestelde oplossingen wordt de (realistische) veronderstelling van *objectcompositie* geïdentificeerd als de bron van het misverstand, ten gevolge waarvan een nieuw

en meer specifiek argument voor transcendentalisme beschikbaar komt. De eerste elementen van de transcendentalistische benadering worden uitgewerkt in de vorm van wat *image theorie* zal worden genoemd, een theorie rondom de modules van het menselijk cognitief systeem ('*images*') die *procedures* uitvoeren m.b.t. wat voor onze zintuigen toegankelijk is. Zij stellen eigenschappen vast en bewerkstelligen daarmee het *te voorschijn komen* van de objecten zelf.

Hoofdstuk 4 en 5 zijn gewijd aan de analyse van twee belangrijke manieren waarop nieuwe objecten ontstaan uit bestaande: *mereologische fusie* en *abstractie*. De Lewisiaanse veronderstelling van 'ontologische onschuld' wordt tegen het licht gehouden, wat leidt tot de conclusie dat *mereologisch nihilisme* het vrijwel onherroepelijke gevolg ervan is. Een transcendentalistische benadering van mereologische fusie blijkt echter nieuwe inzichten op te leveren over de werking van procedures: *naburighedsrelaties* ('vicinitieve relations') tussen objecten (relaties waarbij het kennen van de een van de relata het 'aanwijzen' van de andere mogelijk maakt) maken lussen in procedures mogelijk, van belang voor de representatie van meervoudigheden en hun hiërarchische ordening. Deze gedachtegang culmineert in de *Centrale Aanneme van Image Theorie*, een algemene hypothese over de voorwaarden waaronder objecten tevoorschijn komen. *Abstractie*—het ontstaan van objecten zoals *de tijger* en *witheid* uit tijgers en witte dingen—is een fenomeen dat in veel opzichten overeenkomt met mereologische fusie. Door de voorgestelde transcendentalistische interpretatie van deze objecten blijkt algemeenheid de standaardtoestand te zijn en individualiteit de uitzondering, veroorzaakt door de eigenschappen van de ruimte. Het terrein van 'diepere' abstractie en de rol van abstractie in *perfectie* worden kort verkend.

Hoofdstuk 6 is gewijd aan enkele overgebleven kwesties waar image theorie iets over te zeggen heeft, met een centrale rol voor het *zijn* zelf—en het ontbreken ervan, niet-zijn. Eerst wordt de relatie van *identiteit* besproken, een verre van triviale relatie, waarvan het enige moeite kost om die haar juiste plaats in image theorie te geven. Het begrip *geworteldheid*, tot hiertoe weinig meer dan een intuïtie, wordt nader uitgewerkt. Het wordt voorgesteld als alternatief voor 'relatieve identiteit'. De voorwaarden om echt te zijn, onecht, of iets er tussenin, worden behandeld aan de hand van ideeën over *bestaan* en *fictie*. Het hoofdstuk sluit af met enkele gedachten over *verbeelding* en *modaliteit*. Tenslotte worden in hoofdstuk 7 enkele van de sterke en zwakke punten van zowel de transcendentalistische benadering als image theorie besproken, alsook mogelijke verdere toepassingen.

Summary

The Things Before Us: On What it Is to Be an Object

What is an object? In this thesis it is this seemingly trivial concept which is up for investigation. Inspired by some of the insights of Immanuel Kant, the concept of object is given a *transcendentalist* treatment, the central idea of which is that the practice of viewing all of reality in terms of objects, having properties and being related in certain ways, is a *human* mode of understanding, inherent in our cognitive system.

In Chapter 1 and 2 the mind-dependence of objects is defended with the help of the concept of *rootedness* (the being grounded of some object in other objects). It is argued that *first-order logic* should be seen as an expression of the main principles guiding human cognition. A short survey is done of what I call *the draft ontology*, our intuitive picture of the world, that these principles are intended to describe. It is argued that the study of the human cognitive system is an essential element of the methodology to arrive at a more systematic ontology. Rather than being put to work as a formal calculus, first-order logic—or rather, an extension of it that grants ‘objecthood’ to a greater variety of symbols—is used as an interpreted symbolism: a (non-ideal) representation of the draft ontology. As for the reality behind the symbolism, it is argued that the human window on objects is an interplay of thought and perception, attested by language. Armed with this conceptual framework, some fundamental notions are reconsidered: *identity* (both in the sense of ‘thisness’ and ‘sameness’) and *plurality*.

In Chapter 3 the main tools for doing the explanatory work are being unpacked. The chapter starts with a discussion of Bradley’s regress, a notorious issue in ontology. After examining some of the solutions that have been suggested, the (realistic) assumption of *objectual compositionality* is identified as the source of confusion, whereby a new and more specific argument for transcendentalism becomes available. The first elements of the transcendentalist approach are worked out in the form of what will be called *image theory*, a theory centred around the modules of the human cognitive system (*images*) that perform

procedures on what is perceptually available to us. It is these that determine properties and thereby effect the *emergence* of the objects themselves.

Chapter 4 and 5 are dedicated to the analysis of two major ways in which new objects arise out of existing ones: *mereological fusion* and *abstraction*. The Lewisian claim of ontological innocence is evaluated, leading to the conclusion that it has *mereological nihilism* as its almost inevitable consequence. A transcendentalist approach to mereological fusion, however, is shown to bring new insights into the working of procedures: *vicinitive relations* between objects (relations which are such that knowing one of the relata allows ‘picking’ the other) allow there to be loops in procedures, which are important for the representation of pluralities and for their hierarchical arrangement. This line of reasoning culminates in the *Central Assumption of Image Theory*, a general hypothesis about the conditions for objects to emerge. *Abstraction*—the emergence of objects like *the tiger* and *whiteness* out of tigers and white things respectively—is a phenomenon showing many similarities with mereological fusion. By the suggested transcendentalist interpretation of these objects, generality turns out to be the default condition and individuality the exception, brought about by the properties of space. The territory of ‘deeper’ abstractions and the role of abstraction in *perfection* are briefly explored.

Chapter 6 is devoted to some remaining issues that image theory has something to say about, with a central role for *being* itself—and for its absence, non-being. First, the relation of *identity* is discussed, a relation that, far from being trivial, it takes some effort to give its proper place in image theory. The concept of *rootedness*, which so far has been little more than an intuition, is worked out in greater detail. It is proposed as an alternative to ‘relative identity’. The conditions of being real, being unreal, and being halfway will be addressed, guided by ideas about *existence* and *fiction*. The chapter closes with some thoughts on *imagination* and *modality*. Finally, in Chapter 7 some of the virtues and shortcomings of both the transcendentalist approach and image theory are discussed, along with possible further applications.

Titles in the ILLC Dissertation Series:

- ILLC DS-2016-01: **Ivano A. Ciardelli**
Questions in Logic
- ILLC DS-2016-02: **Zoé Christoff**
Dynamic Logics of Networks: Information Flow and the Spread of Opinion
- ILLC DS-2016-03: **Fleur Leonie Bouwer**
What do we need to hear a beat? The influence of attention, musical abilities, and accents on the perception of metrical rhythm
- ILLC DS-2016-04: **Johannes Marti**
Interpreting Linguistic Behavior with Possible World Models
- ILLC DS-2016-05: **Phong Lê**
Learning Vector Representations for Sentences - The Recursive Deep Learning Approach
- ILLC DS-2016-06: **Gideon Maillette de Buy Wenniger**
Aligning the Foundations of Hierarchical Statistical Machine Translation
- ILLC DS-2016-07: **Andreas van Cranenburgh**
Rich Statistical Parsing and Literary Language
- ILLC DS-2016-08: **Florian Speelman**
Position-based Quantum Cryptography and Catalytic Computation
- ILLC DS-2016-09: **Teresa Piovesan**
Quantum entanglement: insights via graph parameters and conic optimization
- ILLC DS-2016-10: **Paula Henk**
Nonstandard Provability for Peano Arithmetic. A Modal Perspective
- ILLC DS-2017-01: **Paolo Galeazzi**
Play Without Regret
- ILLC DS-2017-02: **Riccardo Pinosio**
The Logic of Kant's Temporal Continuum
- ILLC DS-2017-03: **Matthijs Westera**
Exhaustivity and intonation: a unified theory
- ILLC DS-2017-04: **Giovanni Cinà**
Categories for the working modal logician
- ILLC DS-2017-05: **Shane Noah Steinert-Threlkeld**
Communication and Computation: New Questions About Compositionality

- ILLC DS-2017-06: **Peter Hawke**
The Problem of Epistemic Relevance
- ILLC DS-2017-07: **Aybüke Özgün**
Evidence in Epistemic Logic: A Topological Perspective
- ILLC DS-2017-08: **Raquel Garrido Alhama**
Computational Modelling of Artificial Language Learning: Retention, Recognition & Recurrence
- ILLC DS-2017-09: **Miloš Stanojević**
Permutation Forests for Modeling Word Order in Machine Translation
- ILLC DS-2018-01: **Berit Janssen**
Retained or Lost in Transmission? Analyzing and Predicting Stability in Dutch Folk Songs
- ILLC DS-2018-02: **Hugo Huurdeman**
Supporting the Complex Dynamics of the Information Seeking Process
- ILLC DS-2018-03: **Corina Koolen**
Reading beyond the female: The relationship between perception of author gender and literary quality
- ILLC DS-2018-04: **Jelle Bruineberg**
Anticipating Affordances: Intentionality in self-organizing brain-body-environment systems
- ILLC DS-2018-05: **Joachim Daiber**
Typologically Robust Statistical Machine Translation: Understanding and Exploiting Differences and Similarities Between Languages in Machine Translation
- ILLC DS-2018-06: **Thomas Brochhagen**
Signaling under Uncertainty
- ILLC DS-2018-07: **Julian Schlöder**
Assertion and Rejection
- ILLC DS-2018-08: **Srinivasan Arunachalam**
Quantum Algorithms and Learning Theory
- ILLC DS-2018-09: **Hugo de Holanda Cunha Nobrega**
Games for functions: Baire classes, Weihrauch degrees, transfinite computations, and ranks

- ILLC DS-2018-10: **Chenwei Shi**
Reason to Believe
- ILLC DS-2018-11: **Malvin Gattinger**
New Directions in Model Checking Dynamic Epistemic Logic
- ILLC DS-2018-12: **Julia Ilin**
Filtration Revisited: Lattices of Stable Non-Classical Logics
- ILLC DS-2018-13: **Jeroen Zuiddam**
Algebraic complexity, asymptotic spectra and entanglement polytopes
- ILLC DS-2019-01: **Carlos Vaquero**
What Makes A Performer Unique? Idiosyncrasies and commonalities in expressive music performance
- ILLC DS-2019-02: **Jort Bergfeld**
Quantum logics for expressing and proving the correctness of quantum programs
- ILLC DS-2019-03: **Andras Gilyen**
Quantum Singular Value Transformation & Its Algorithmic Applications
- ILLC DS-2019-04: **Lorenzo Galeotti**
The theory of the generalised real numbers and other topics in logic
- ILLC DS-2019-05: **Nadine Theiler**
Taking a unified perspective: Resolutions and highlighting in the semantics of attitudes and particles
- ILLC DS-2019-06: **Peter T.S. van der Gulik**
Considerations in Evolutionary Biochemistry
- ILLC DS-2019-07: **Frederik Mollerstrom Lauridsen**
Cuts and Completions: Algebraic aspects of structural proof theory
- ILLC DS-2020-01: **Mostafa Dehghani**
Learning with Imperfect Supervision for Language Understanding
- ILLC DS-2020-02: **Koen Groenland**
Quantum protocols for few-qubit devices
- ILLC DS-2020-03: **Jouke Witteveen**
Parameterized Analysis of Complexity
- ILLC DS-2020-04: **Joran van Apeldoorn**
A Quantum View on Convex Optimization

- ILLC DS-2020-05: **Tom Bannink**
Quantum and stochastic processes
- ILLC DS-2020-06: **Dieuwke Hupkes**
Hierarchy and interpretability in neural models of language processing
- ILLC DS-2020-07: **Ana Lucia Vargas Sandoval**
On the Path to the Truth: Logical & Computational Aspects of Learning
- ILLC DS-2020-08: **Philip Schulz**
Latent Variable Models for Machine Translation and How to Learn Them
- ILLC DS-2020-09: **Jasmijn Bastings**
A Tale of Two Sequences: Interpretable and Linguistically-Informed Deep Learning for Natural Language Processing
- ILLC DS-2020-10: **Arnold Kochari**
Perceiving and communicating magnitudes: Behavioral and electrophysiological studies
- ILLC DS-2020-11: **Marco Del Tredici**
Linguistic Variation in Online Communities: A Computational Perspective
- ILLC DS-2020-12: **Bastiaan van der Weij**
Experienced listeners: Modeling the influence of long-term musical exposure on rhythm perception
- ILLC DS-2020-13: **Thom van Gessel**
Questions in Context
- ILLC DS-2020-14: **Gianluca Grilletti**
Questions & Quantification: A study of first order inquisitive logic
- ILLC DS-2020-15: **Tom Schoonen**
Tales of Similarity and Imagination. A modest epistemology of possibility
- ILLC DS-2020-16: **Ilaria Canavotto**
Where Responsibility Takes You: Logics of Agency, Counterfactuals and Norms
- ILLC DS-2021-01: **Yfke Dulek**
Delegated and Distributed Quantum Computation