

ABSENCE & NOTHING

STEPHEN
MUMFORD

the philosophy of what there is not

OXFORD

Absence and Nothing

Absence and Nothing

The Philosophy of What There is Not

STEPHEN MUMFORD

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

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

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

© Stephen Mumford 2021

The moral rights of the author have been asserted

First Edition published in 2021

Impression: 1

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

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

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

British Library Cataloguing in Publication Data

Data available

Library of Congress Control Number: 2021911519

ISBN 978-0-19-883153-2

DOI: 10.1093/oso/9780198831532.001.0001

Printed and bound in the UK by
TJ Books Limited

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

Contents

<i>Preface</i>	vii
<i>List of Table and Figures</i>	xi
1. Soft Parmenideanism	1
1.1 Nothing really matters	1
1.2 The first argument	2
1.3 Soft ontological Parmenideanism	3
1.4 Soft methodological Parmenideanism	13
1.5 Is the cat out of the bag?	16
1.6 The way forward	18
2. Negative Properties	19
2.1 What would be a negative property?	19
2.2 Arguments against negative properties	20
2.3 For negative properties	30
2.4 Summation and further work	36
3. Nonentities	38
3.1 Negative particulars	38
3.2 Non-beings	39
3.3 Limits, boundaries, edges, and stops	45
3.4 Privations	48
3.5 Omissions	52
3.6 Negatives norms	53
3.7 Negative epistemic states	54
3.8 Logical and mathematical negatives	57
4. Causation by Absence	64
4.1 Negative causation	64
4.2 More cases	65
4.3 Creation <i>ex nihilo</i>	66
4.4 Why causation by absence spells trouble	69
4.5 Attempts to stop escalation	72
4.6 Other ways of explaining causation by absence	77
4.7 Explanation but not causal explanation	82
4.8 Summary of conclusions	84
5. Mere Possibilities	85
5.1 Possible queues ahead	85
5.2 Reification	87
5.3 Fictionalism	91

5.4	Grounding possibility in what there is	93
5.5	Parmenidean possibility	99
7.	Perception of Absence	101
7.1	Seeing what is not there	101
7.2	Our subject matter	103
7.3	Perceptual theories	107
7.4	Cognitive theories	113
7.5	A Dennettian account	118
7.6	Where this leaves us	125
8.	Empty Reference	128
8.1	Can we talk about nothing?	128
8.2	The axiom of existence	131
8.3	Proxy referents	133
8.4	Aboutness	137
8.5	Fictionalism	138
8.6	Pretence	140
8.7	Reference failure	143
8.8	Aboutness again	145
9.	Negative Truth	147
9.1	The hippo not in the room	147
9.2	Molnar's problematic	150
9.3	Solutions proper?	153
9.4	Attempted solutions that reject at least one of (Mi)–(Miv)	159
9.5	A more recent solution	164
9.6	Impasse?	166
10.	Negation and Denial	168
10.1	Partial success?	168
10.2	The equivalence thesis	169
10.3	The separate functions of assertion and denial	172
10.4	What are you denying?	186
10.5	Can you deny a denial?	189
10.6	Logic and complex denials	191
10.7	Internal and external negations	193
10.8	Conclusion	194
11.	Negative Belief	196
11.1	The final reckoning	196
11.2	Inventories of being and non-being	196
11.3	How to believe that something is not	201
11.4	Types of negative belief	205
11.5	Ontological Parmenideanism	206
11.6	Final words	207

<i>References</i>	209
-------------------	-----

<i>Index</i>	217
--------------	-----

Preface

I am grateful to the Leverhulme Trust in granting me a Major Research Fellowship from 2016 to 2019 specifically to complete this work. Leverhulme remains willing to fund blue-skies research that has no immediate practical benefit. The present book falls firmly into that category.

Problems relating to absence and nothingness have troubled me since 1989 when I read Armstrong's *A Theory of Universals*. He discusses at one point whether there are negative properties, such as being non-red, non-circular, uninitiated, sugar-free, and so on. I immediately shared Armstrong's instinct that such 'features' did not deserve a place in reality. The problem was, however, that when Armstrong provided specific arguments against negative properties, they all seemed question begging. As an unknown PhD student, I wrote a letter to Armstrong (in the days before email) on the other side of the world and explained my worries. To my delight, a few weeks later I received a detailed reply in which Armstrong admitted that his arguments had begged the question and he then offered some further considerations. Negative properties did not seem particularly natural and perhaps should be excluded from our ontology on that basis. I couldn't rest content with this explanation either but had to shelve my concerns and get on with my PhD, which was on an unrelated topic. Sometime later, doctorate secured, I revived my interest in negativity and presented a paper in which I attempted a new argument against negative properties. This contained an elementary error so I shelved it indefinitely. If Armstrong's arguments had failed, mine had failed spectacularly.

Negativity would not go away, however. Over time, I noticed a number of other philosophical problems that concerned what there was not: problems such as causation by absence and truthmakers for negative truth. The latter was a problem I took from George Molnar. In one of the final papers he published, he articulated the problem in what is now widely acknowledged as the definitive formulation. His paper ends on a sorrowful note, however. No one had a solution to this problem and he could find none either. I took this as a challenge and published a couple of small pieces on the subject, with which I was not entirely happy. Still more problems of absence and nothingness arose: problems such as perception of absence, nonentities, holes and shadows, negative facts and the act of denial. In 2005, I first had the idea of collecting these together for a monograph in which there would be a common solution offered. I have a notebook from that date in which I plan out the contents for *A Book About Nothing*, which later became *Nothing Really Matters*. Neither of those was written because I became

immersed in other projects and then unfortunately fell into university management. The idea was to have a single, systematic approach that would dissolve these many related problems and show that nothingness was not a feature of ontological reality but only a feature of our ways of thinking or speaking. Until I examined the detail, however, I could not be sure that it would all work out that way.

The plan was once more shelved as other matters took priority. Only as I reached the end of my period in management did I think once more that I ought to revive my interest in nothing in particular. Having dwelt on absences and nothings for twenty-six years, I was able to make the case to Leverhulme that the book really would not be completed without their valuable support. As Colin Howson wryly observed, they decided I was good for nothing.

I started the detailed work, which took some unexpected turns. Most notably, I recognized the benefit of situating my approach within the Parmenidean tradition. I realized, in doing so, that my concern was with the very first problem seriously discussed in Western philosophical history. That shouldn't really come as a surprise. The topic, after all, is the very division between Being and Non-being. There can be nothing more fundamental than that. Parmenides was largely correct. What exists is everything. What does not exist is nothing. But Parmenides rather overextended this thesis, arguing that we cannot even think about what there is not, and that change is impossible. I was never going to accept that. My sympathies were only softly Parmenidean, I realized, and this could inform an approach to the whole subject matter. Let us test whether we can see through a Parmenidean project, explaining some long-standing philosophical problems, without ever having to invoke an absence or nothingness as part of reality. I will leave the reader in suspense over whether a soft Parmenidean approach is successful in the final reckoning. One further discovery, however, was that we cannot resolve all these issues doing metaphysics alone. An understanding of empty reference, denial, and negation was also necessary. Another possible title, *The Metaphysics of Nothing*, also had to be rejected, then. Only the first half of this book is metaphysics. The second half is epistemology, philosophy of language, and logic.

What have I got to show for over thirty years thinking about nothing? I have certainly learnt a lesson in humility. These philosophical questions are bigger than all of us, and it is very unlikely that a simple response will work. A whole lifetime is barely enough to complete this kind of work. Nevertheless, I hope that the present book offers at least some progress and might allow others to see the issues in a more illuminating light. I'm happy if what I have done enables or encourages others to do better.

Having acknowledged the Leverhulme Trust, I should also thank those who have supported me personally, especially over the last three years when I was working intensely on the book. That includes my family, whom I relocated from

Nottingham to Durham, and all my colleagues in Philosophy at Durham. Being on externally funded leave meant that I was not contributing to the life of the department in all the ways I would have liked. I should thank Oxford University Press, especially philosophy editor Peter Momtchiloff, for once again having faith in my work and making a commitment to seeing it published. Peter's readers then gave me many detailed comments. Appointed readers and referees often get a bad press, and it is almost always a thankless task. It is important for me to point out, then, that the two anonymous readers of the beta-version gave me the most helpful, incisive, and constructive comments I've ever had. In countless places I have followed their suggestions, and it has made for a vastly improved book. Also working on behalf of the press, I must thank my copy editor, Phil Dines, who had to wrestle with, in places, a difficult text, and push me on some deeper-than-you'd-think lexical-philosophical matters, such as whether we should distinguish 'what-is-not' from 'what is not'. Copy editors are among the many unsung heroes who turn an author's almost-ready thoughts into a presentable book.

I would like to thank all those who attended talks I gave on topics related to absence and nothing and for the comments and criticisms received. As this has been over such a long period, there are too many venues to list and names to name, but I hope that, if you're one of those people, this volume does justice to your concerns and you find the answer to your question herein. I am grateful to Liva Rotkale for assisting me with the Latvian emphatic denials used in chapter 10. Some of the material in chapter 10 has appeared in an earlier form as 'Negation and Denial' in the *Cambridge Handbook of Philosophy of Language*, P. Stalmaszczyk (ed.), and appears here with permission of the editor and publisher.

Finally, this book was completed during the pandemic at a time that has been difficult for us all. In my own case, it was an extremely tough year. As I write these words, the crisis is still far from concluded. Friends and family have helped me through, and I will never forget the importance of that. My hope is that some good emerges from this nightmare and that we will re-evaluate the ways we live to produce a better and more caring post-pandemic world.

List of Table and Figures

Table

10.1	Table of assertion and denial	173
------	-------------------------------	-----

Figures

7.1	Kanizsa triangle	104
7.2	Absent dot in box B?	107
10.1	Relative determinacy of assertion and denial	177
10.2	Relative scope of assertion and denial	182

1

Soft Parmenideanism

1.1 Nothing really matters

It matters to us how things are. But how things are not can be equally important. It matters when something you want is not there. In a variety of circumstances, it matters that a drink is sugar-free. A migraine sufferer might want any pudding as long as it is not chocolate. Anyone with an allergy or intolerance will have a similar type of negative preference. It could matter to someone that a particular train does not stop in Wakefield, if they were planning to get off there. Bereavement matters too. When a loved one is no longer around, the feeling of their absence is profound. Some fear their own non-existence more than anything else. In other cases, things not being a certain way can be entirely mundane. Someone might regret not being six-feet tall or not being a better chess player. But some cases of what is not are of great, even cosmic importance. Because the spread of Covid-19 was not stopped, the UK suffered over 130,000 deaths. This omission itself could be explained in terms of an absence: the lack of an effective government strategy. It matters enormously that the planet Earth is not closer to the sun than it is. It is likely that there would be no human life on the planet if it were. Equally, it matters that we have not recently been hit by a giant asteroid: the sort of impact that might have led to extinction of the dinosaurs.

If *nothing* matters so much, we ought to have at least a basic understanding of how it matters. Unfortunately, as soon as we try to understand what-is-not, we are immersed in the deepest of deep philosophical perplexities. If what-is-not is an absence of something, does it mean that absences are parts of reality? How could they matter otherwise? But if they are real, it seems as if we are saying that what-is-not is a part of what-is. That looks downright self-contradictory. Suppose, then, that we straightforwardly reject the reality of absences, voids and nothings, and any kind of entity that looks 'negative'. Can we still give a complete account of the world? For example, in some cases of causation, it looks as if it matters not only what is present but also what is absent. If you had no air to breathe, you would die. What would kill you? The most obvious cause would be a lack or absence of oxygen. It is very difficult to explain the death in other, 'positive' ways: in terms of what was present. Various things were there. But their presences don't entail the absence of oxygen, which it seems would have to be cited as an extra fact about the situation, in addition to the facts of all that was present. Furthermore, it seems that the presences of

all those things are not what explain your death. The absence of oxygen looks essential and indispensable in explaining the effect under consideration.

These questions, we will see, are not trivial. They are easy to state but hard to answer. The distinction between being and non-being must be understood for us to make progress, and perhaps there is no more fundamental distinction in the whole of metaphysics than this. There will be much more we must understand besides.

1.2 The first argument

Nothingness, non-being, was the first philosophical problem in the Western tradition. The oldest such surviving text in which there is sustained philosophical argument is Parmenides' poem *On Nature*, which discusses non-being in the section entitled 'The Way of Truth' (Parmenides *Fragments*). Graham (2010: 4) makes a credible point when he says that the main division in Greek philosophy should not be between pre- and post-Socratics but between pre-Parmenideans and post-Parmenideans. Plato's dialogue *The Sophist* can be understood as a direct response to Parmenides. In one place, the Eleatic visitor, granting a claim that looks anti-Parmenidean, confesses that he might be guilty of patricide (*Sophist*: 241d3–4). Parmenides was from Elea, and the Eleatic visitor is intended, either literally or figuratively, to be Parmenides' son.

Parmenides argued that there is no non-being. Indeed, we should not even speak of it. Gorgias (*On What-is-Not*: 741) soon dissented, claiming that what-is-not exists. If we look at more recent philosophical history, non-being plays a significant role in the philosophies of Hegel, Heidegger, and Sartre. It is not just a view of the Continental tradition, however. Russell (1918: 206–8) admitted that there must exist general facts that are essentially negative. Despite being a professed naturalist, Armstrong (2004: 54–7) followed Russell in accepting higher-order negative facts, which his facts of totality are. Among contemporaries, Schaffer (2004) believes that absences can be causes, which seems to commit to their existence and their powerfulness. Hommen (2013) and Zangwill (2011) accept that there are negative properties, such as being non-red. Barker and Jago (2012) allow negative instantiation as a likely part of reality. And Priest (2005) states that some non-existent things nevertheless are. It seems that there is currently no shortage of analytic philosophers willing to whisper sweet nothings in our ears. Often they will depict themselves as holding a radically unorthodox view. The above list shows, however, that the acceptance of negatives has become mainstream. Take, for example, Sorenson, who describes himself as 'a reluctant apostle of negative metaphysics' (Sorensen 2008: 16), committing to the reality of negative features in the world and our ability to perceive them. Nevertheless, he begins by assuming that 'Philosophers and physicists alike have a strong conviction that reality is positive' and speaks of a natural 'Aversion to negative things'

(Sorensen 2008: 16). On the contrary, a philosophy of ‘negativism’ is currently prevalent. It is virtually conventional wisdom.

Despite this ongoing acceptance of negative entities of various kinds, I will defend the spirit of Parmenides’ position, if not the letter. There is no nothing. Absences and nothingnesses have zero reality.

I am, however, going to adopt a stance that I call soft Parmenideanism. I use this label with a dual meaning. Soft Parmenideanism is my name for a methodology that is Parmenidean in tone but, unlike Parmenides himself, sets out what conditions would have to be met in order for a negative entity to be accepted into our ontology. The hard Parmenidean, presumably, is someone who would never countenance negative entities in any circumstances, which is another way of saying that a rejection of negative entities is their most basic and fundamental commitment. Below, I shall be explaining the soft Parmenidean methodology and justifying its adoption.

Soft Parmenideanism can also be used as a label for a position that is held. Such soft Parmenideanism would be accepted as an ontological thesis. It is soft because it accepts some of what Parmenides said—basically, the rejection of negative existents—but it does not accept all of what Parmenides said. The soft position rejects the more hardline views of Parmenides: such as that motion and change are not real (perhaps the view most associated with the Eleatics) and that it is impossible to think about what is not. As Gallop (1984: 24) shows, it is in any case dubious that these stronger claims followed from Parmenides’ reasonable starting point that what-is-not is not anything. Hard and uncompromising Parmenideanism is certainly entertaining, as Della Rocca (2020) proves, but it is not the view I will be defending here. Taken overall, the present book makes the case for a soft Parmenideanism, both as an ontological thesis and as a philosophical methodology.

We need to understand Parmenides’ original position and why so many have sought to overturn the parts that seem obviously true. Some might think it trivially so. I duly start with a look at Parmenides’ specific claims and his arguments for them. I will then go on to outline an approach of soft methodological Parmenideanism that will be applied in the rest of this book. It follows the spirit of Parmenides’ position, though it leaves open the possibility that the approach might not deliver the hard conclusions for which Parmenides is notorious. It accepts that there are conditions under which some form of negative entity would have to be accepted as real. The rest of this book will then examine whether there are in fact cases in which those conditions are met.

1.3 Soft ontological Parmenideanism

I shall come to my proposed Parmenidean methodology shortly, but first I wish to describe, and to an extent justify, my soft Parmenideanism as an ontological thesis. This is a view concerning what there is and what there isn’t.

We must begin with Parmenides' text, specifically 'The Way of Truth,' which is accepted as the part of the poem in which Parmenides expresses his own philosophy, though he has a character voice the words. I will highlight seven main claims that I think should be taken from Parmenides before explaining why some of them, but not all, might be accepted. Another reader of Parmenides could produce a different list of the Presocratic's most important claims. Della Rocca (2020: 1), for instance, argues that the rejection of all distinctions is a key claim in Parmenides. Parmenides' philosophy is particularly open to interpretation since it is written in hexameter verse and the exact meaning is debatable. I admit also that I have approached Parmenides' text with particular interests and purposes in mind, as the rest of this book evinces. I believe exactly the same can be said of Della Rocca's book, of course.

All those provisos declared, we can say that Parmenides is reasonably attributed with the following views:

Pi. Nothing is not (A)

Pii. There are no degrees of being (A)

Piii. Nothing comes from nothing (A/R)

Piv. Non-being is unknowable (A/R)

Pv. Reality is a single plenum (A/R)

Pvi. Non-being is unthinkable and unnameable (R)

Pvii. Neither motion nor change is possible (R)

As an indicator of what is to come, I have followed each claim with an (A), where I think we should accept it, an (R) where I think we should reject it, or an (A/R) where I think that we can either accept or reject it, depending on how it is interpreted and our other commitments.

Let us take it, then, that a hard ontological Parmenidean is someone who accepts all of Pi to Pvii. Such a position would fully deserve to be called hard. The view would be not just that there is no nothing but also that we cannot even speak or think of what is not. It would follow, of course, that one could not even consistently deny that nothingness is something. If I was a hard ontological Parmenidean, this could be a very short book. None of the discussions that follow, concerning negative properties, causation by absence, negative truth, perception of absence, and so on, could meaningfully be articulated.

The soft Parmenidean, I suggest, is someone who holds to a proper subset of Pi to Pvii, so rejects at least one of them. This immediately suggests the question of how many claims are to be held and allows that there could be degrees of Parmenideanism. Some soft Parmenideans might be harder than others, if they reject just one of Pi to Pvii, for instance, while another rejects five of them. My own soft Parmenideanism, as indicated above, is towards the softer end of the scale, given that I see it as a firm commitment to only two of Pi–Pvii. Furthermore, I commit to the Parmenidean claims that I take to be the most obvious, almost

undeniable, and which cause relatively little difficulty with the rest of our philosophy. But my Parmenideanism is not completely soft. I would be willing to accept Piii, iv, and v if they are interpreted in a certain way or restricted in a specific manner. That could then lead to accepting all of Pi to Pv in some form or other, which would be a fairly robust Parmenideanism. As the above menu indicates, however, I do think that it is a perfectly tenable position to hold only Pi and ii while rejecting Piii, iv, and v. It can be inferred from this that I do not take Pi and ii to entail iii to v.

So that this discussion does not remain forever abstract, I should say something of what I take each of Pi to Pvii to mean, give some textual justification for attributing such views to Parmenides, and provide preliminary reasons for why I think each claim should be accepted or rejected. (The references to Parmenides' text are given by the number of the fragment followed by the line number, all taken from the Gallop 1984 edition. Hence 6.2 refers to fragment 6, line 2)

Pi. Nothing is not

... nothing is not, that is what I bid you consider (6.2)

For never shall this prevail, that things that are not are; (7.1)

Not all of Pi to Pvii carry equal weight, since it is hard to see how anyone who rejects Pi can claim to be any kind of Parmenidean at all, even if they held most or all of the remaining claims. This is Parmenides' most important and significant claim, since it concerns the division between being and non-being. Indeed, I have arranged Pi to Pvii in an order where they become progressively less and less essential to Parmenideanism, as I am conceiving it.

How is Pi to be understood? In its simplest form, which is how I will take it, Parmenides is saying that there is no nothing. There is only what-is. There is only what exists, hence non-being is not something. The point ought, then, to be obvious. Perhaps it is trivial and hardly worth stating. It rejects what Gallop (1984: 24–6) calls the Homeric equivocation (after Homer, *Odysseus* IX.366–7 and 398–412). If you say 'there is nothing in my pocket' you cannot be taken to mean that there is indeed something in your pocket, namely 'nothing', since nothing is not a thing. Similarly, in Homer, to say 'Nobody killed him' does not mean that someone did kill him, namely 'Nobody'. Nor, in Parmenides, when the goddess accuses humans of 'knowing nothing' (6.4), she does not mean that there is an object of knowledge called 'nothing' that we know. Such inferences are also parodied by Lewis Carroll, through his character of the White King, who misunderstands the statement 'I see nobody on the road' (Carroll 1871: 234). We should not, therefore, attempt any kind of positive description of the nature of nothingness, such as we find in Heidegger (1929), for instance.

That the Pi claim seems trivial can be questioned to an extent, however. For one thing, it is open to interpretation exactly what Parmenides meant since the 'is' in 'it is and cannot not be' (2.3) is ambiguous and it is likely that Parmenides

was using it in slightly different ways in different places (Waterfield 2000: 50). The 'is' can be interpreted existentially, as when one says 'the Earth is,' simply meaning that it exists. But 'is' can also be used predicatively, as in 'it is...'. Waterfield also distinguishes a veridical sense of 'is' that might be in play since only if x is F , for some attribute F , can we say that x is real. With only limited and imprecise textual evidence, we cannot rule out any of these meanings of 'is', so we cannot immediately dismiss Π as trivial. For the purposes here, I am accepting an existential interpretation of the claim and thus taking it, as indicated above, as a claim about what there is (everything) and what there is not (nothing).

A second reason to think that Π is not entirely trivial is that Parmenides thinks it worthy of an argument, which he gives at 8.15–16. Gallop (1984: 7) reconstructs it as follows:

- [A] 1. Either 'is' or 'is not' (and not both)
 2. Not 'is not'
 Therefore,
 3. Is.

The disjunction in 1 can also be used to argue from 3 to the conclusion of 2, of course: that not is not. Now it might be significant in the history of philosophy that Parmenides sought to argue for Π rather than merely assert it, since most of what we have from his predecessors consists in assertion or rhetoric rather than arguments (see Waterfield 2000: 49). In many ways, then, this is how serious philosophy truly began. What can be taken from the fact that Parmenides offered an argument is that although the claim in Π might seem obvious and even trivial, Parmenides evidently did not think so since the argument [A] shows that he thought there were assumptions that were even more basic and fundamental than the conclusion expressed in Π . After all, someone could think it possible both that there is and there is not—there is both something and nothing—and this is what premise 1 of [A] denies by making the disjunction exclusive. Of course, once premise 1 is accepted, premise 2 looks a lot like Π itself, so we might not be persuaded by [A] as a reason to accept Π . Perhaps a better argument could be constructed.

There is, however, a third and perhaps the most important reason to doubt that Π is trivial, which is that many philosophers effectively deny it. As we will see in the remainder of the present book, there are reasonable analytic philosophers who think that we have no choice but to accept certain kinds of negative existence. They do not think that they are denying a trivial truth, for they would usually understand that to be confused, but it is almost certain that they are denying Π , under its existential interpretation. Even if one thinks Π is not exactly trivial, one might still think that it is undeniable for other reasons. To put it in familiar terms, one might say that Π is at the centre of our web of belief (Quine and Ullian 1970). A consideration of soft Parmenideanism as a methodology will

reveal, however, that there could be theoretical grounds for denying Pi if it would leave us with an overall more coherent web of beliefs.

There is a final remark on Pi concerning Parmenides apparent necessitarianism. Parmenides characterizes what-is also as what ‘cannot not be’ (2.3) and what-is-not as what ‘needs must not be’ (2.5). Consequently, it would be possible to ascribe to Parmenides a thesis that what is necessarily is and what is not necessarily is not; that is, a necessitarian view of nature, which denies that there are any contingent truths. This could easily be an eighth Parmenidean thesis on the list, but I do not itemize it separately since I think it is best covered under Pv.

Pii. There are no degrees of being

[It] must either be completely or not at all (8.11)

[it] all alike is; (8.22)

... nor is there a way in which what-is could be

More here and less there, since [it] all inviolably is; (8.47–8)

Pii tells us that existence is univocal in respect of degree. Something that exists, exists entirely and absolutely. Nothing can exist to a lesser extent than something else. Existence does not admit of any degrees at all, just as there is no sense of more or less perfect. Hence, the most insignificant speck of dust, if it exists, exists in the same sense and just as fully as the island of Manhattan exists, or the planet Earth, or the whole universe. There is, thus, a stark choice for us when we evaluate what there is. There is no middle ground between existence and non-existence. There is nothing that partially exists or almost exists. One can cut a log in half and the resultant halves exist just as much and in the same sense that the previous whole did.

This thesis, again, might seem very obvious. But there was a time when it was necessary to reassert it. Pii is associated in the twentieth century with John Anderson, for instance, having been spread by his student David Armstrong (see Mumford 2007b: ch. 1). It is seemingly challenged in certain interpretations of Meinong, but, as we shall see (ch. 8, below), this is largely due to Russell’s misrepresentation (such as Russell 1905). Russell’s own statement ‘there is only one kind of being, namely, being simpliciter, and only one kind of existence, namely, existence simpliciter’ (Russell 1903: §427) sounds in accord with Pii. The appearance is superficial, however. Read in context, it is clear that Russell was distinguishing being from existence: both were univocal, but they meant different things. This is not, then, the sort of view a Parmenidean should welcome. We will see that even some contemporary philosophers deny Pii, holding a position named egalitarianism (for example, Zangwill 2011). Pii rules against any equivocating, prevaricating, or obfuscation. Along with Pi, Pii is a Parmenidean thesis that I accept into my soft Parmenideanism. It rules that negative-being cannot just be a different kind of being, or a lesser degree of it. All I can see against Pii is that it might

also be a trivial truth, but again, that is dubious when some have denied it. We would, I maintain, need very good reasons to abandon it, and they had better not be ad hoc.

Piii. Nothing comes from nothing

what-is is ungenerated and imperishable (8.3)

In what way, whence, did [it] grow? Neither from what-is-not shall I allow

You to say or think; for it is not to be said or thought

That [it] is not. And what need could have impelled it to grow

Later or sooner, if it began from nothing? (8.7–10)

It is on the basis of these fragments that Parmenides is attributed with the view that nothing comes from nothing. Parmenides might have been the first to articulate the view, though his formulation is not the best known. Lucretius' 'nothing can be created out of nothing' (*De Rerum Natura*: 31) probably holds that distinction. It is with Piii, however, that matters become more complicated for the soft ontological Parmenidean, since the text is ambiguous between a rejection of causation by absence and another view.

Parmenides argues against the possibility of change, and that includes genesis and perishing (see under Pvii, below), since genesis would involve a change from there being nothing to there being something, and perishing is a change from there being something to there being nothing. Setting aside any such general claim about change and considering in isolation the claim that nothing can come from nothing, then Piii looks a reasonable thesis to hold, certainly if one already accepts Pi. Waterfield explains as follows:

Parmenides denies that anything can come into existence from something that does not exist. Given a state of non-existence, we cannot explain a state of existence since we have no way of moving from one to the other. Since, by definition, what-is-not has no properties, it has no properties that could be taken to explain the generation of what-is. (Waterfield 2000: 52)

Another way of understanding this is that the non-existent has no causal powers, so there is nothing that could give rise to something. We will see (chapter 4), that there are contemporary philosophers who accept the reality of this so-called causation by absence, namely that absences can cause effects. Nevertheless, there is an obvious danger that an acceptance of causation by absence leads to a violation of Pi. The reason is to be found in the Eleatic stranger's reality test in Plato's *Sophist* (247d–e). The stranger says that causal power is the mark of the real: to be real is to bear at least some causal power. If one allows that an absence has a power to produce an effect, then, it entails that the absence is something real. Contrary to Pi, therefore, it is entailed that a nothing is a something. Parmenideans should avoid accepting causation by absence for this reason.

There is a second interpretation of the dictum ‘nothing comes from nothing’, however, that it is possible to reject. If one again sets aside Parmenides’ claim that change is generally impossible, there would be no reason to deny the possibility that there could have been nothing and then there was something, by which I mean spontaneous existence. This is not the same as causation by absence since it is not a case that the previous nothing produces or creates something. It is that suddenly there is something when before it there was nothing. No action of a causal power is required for existence of the something.

To decide which of these two interpretations is correct will depend on how we understand the *comes from* in ‘nothing comes from nothing’. Does it imply causation? If it does, then we should accept Piii, or at least so I will be arguing below. But if it does not imply causation, and one reads its denial as saying only that there could be something where previously there was nothing, then it seems one could meaningfully jettison Piii, under that reading. Nevertheless, as I have indicated, there is a version of Piii that I think a soft ontological Parmenidean ought to accept.

Piv. Non-being is unknowable

You could not know what-is-not . . . nor could you point it out (2.7–8)

Parmenides’ claim Piv is that non-being is unknowable. This again admits of ambiguity. It might be defensible under one reading but then be dubious under another.

If one takes a superficial reading, interpreting the words at face value, then the claim seems false, which is probably why our first judgement will be against Piv. Among the things known, it is tempting to include items such as that London is not the capital of France, that we do not live on Venus, that $2 + 2$ does not make 5, that there are no Arctic penguins, and so on. Sometimes such ‘negative knowledge’ might be important to us, from an evolutionary standpoint, for instance: to know that there is no food at a certain place or that there is no way to jump across a ravine. If Piv amounts to a denial of negative knowledge then, at the very least, we are owed an explanation of how it seems that we can know things of this kind. Perhaps it is possible to explain away the appearances.

Under another reading, however, Piv looks a lot more plausible. To know something, or to perceive it, and thus to point it out, requires the existence of the object of knowledge or perception. It requires that the knower enters into a relation with the thing known and that the perceiver enter into a relation with the thing perceived. This is a factive account, meaning that P can be known or perceived only if P is. Hence, you cannot perceive a pot of gold if it is hallucinated; you can only misperceive it. Nor can you know that you are on the moon if it is not a fact that you are on the moon. You would be mistaken if you thought that you were.

Piv requires further exploration, then. Among the issues we should consider is whether absences can be perceived. Can you see that someone is not in the room?

Can you see darkness? Can you hear silence? These questions will be the subject of ch. 7, below. We need to understand how thoughts can be about the non-existent (ch. 8). And then we also need to understand how we can believe something not to be the case. Even if we cannot know that something is not, it seems that we can still believe it, and this itself requires an explanation (ch. 11, below). The soft ontological Parmenidean is likely to take at least some of Piv seriously, then, or desire an account that explains the appeal of what Parmenides said.

Pv. Reality is a single plenum

... it is all one to me (5.1)

Nor was [it] once, nor will [it] be, since [it] is, now, all together,

One continuous; (6.5–6)

[it] is all full of what-is. (8.24)

... [it] is completed,

From every direction like the bulk of a well-rounded sphere, (8.42–3)

The idea of reality as a single plenum might be acceptable. Monism has made something of a comeback in recent years, such as when properties are understood as a single interconnected whole (Mumford 2004: ch. 10), or there exists only one thing (Horgan and Potrč 2008), or the one whole has priority over the many parts (Schaffer 2010a, 2010b). The hardline Parmenidean Della Rocca (2020) takes strict monism as the key consequence of Parmenideanism and applies it to a wide range of philosophical areas. The problem, from a Parmenidean point of view, is why that single thing has to be filled at every place: full of what-is and completed. The simplest reason, as we will see below (§3.3), is to think that Parmenides could allow no gaps in reality, as that would mean it included non-being in some places. It is not obvious that such an argument survives in Parmenides' fragments, but something like this could have been elsewhere, in what didn't survive (it is quite a philosophical irony that there are parts of Parmenides' poem that no longer exist: see also Della Rocca 2020: ch. 13). We will see, too, that Parmenides denies the possibility of motion and change, but this is not directly connected with the question of whether what-is is a plenum, since motion is possible in a plenum if parts of it move simultaneously. Even if there are no parts in the plenum, motion might still be possible even for an unvariegated whole (as in the homogeneous spinning disc argument, see Armstrong 1980 and Hawley 2001: 73ff). We cannot know whether Parmenides realized this.

It seems that the soft Parmenidean is at liberty either to accept or reject Pv dependent on their other commitments and what additional arguments are brought to bear. As we will see (ch. 3, below), much will depend on how we regard gaps or holes. Are they essentially non-beings? If they are some kind of negative existent, then we can see why reality would have to be filled at every place. If we can make sense of non-being in a non-reified way, however, then

perhaps we can deny the necessity of the plenum. Likewise, if holes or gaps are not essentially negative.

Pvi. Non-being is unthinkable and unnameable

It must be that what is there for speaking and thinking of is; (6.1)

Neither from what-is-not shall I allow

You to say or think; for it is not to be said or thought (8.7–8)

In an extended passage, Parmenides argues that what is available for thinking and speaking cannot not exist (8.8–18). If we can think about it, then it exists; and if it does not exist, then we cannot think of it or name it. It would appear that there are some startling consequences of this view. One cannot say of some *x* that it does not exist, since it would come out as self-contradictory (see Quine 1948: 1). It appears that we cannot coherently say that there is no nothing, even though Parmenides appears to have done exactly that.

It might be, then, that we have to reconsider what we can meaningfully say in relation to any kind of non-being, but it is also clear that Parmenides has hit upon a serious philosophical problem, one known as the problem of empty reference. How can one say that Santa Claus wears a red coat, or that Oliver Twist wanted more porridge, or that centaurs have wings, if neither Santa Claus, nor Oliver Twist, nor centaurs exist? How can we say that there are no mermaids or Arctic penguins? There are various types of reference, concerning truth in fiction and myth, negative singular, and general existence claims, all of which would become problematic if no empty terms are allowed.

Much recent philosophy has been devoted to finding a satisfactory answer to questions like these, which I shall consider below (ch. 8). It is clear, however, that Parmenides' claim is even more wide-ranging than this. It is not just about negative existentials, impossible, mythical, or fictional objects. It applies to all and everything that is not. It rules against any contingent falsehood. Hence, it is supposed to be unthinkable and unsayable that Hillary Clinton is president, since she is not. One would have thought that, even in stating the example, enough has been done to falsify Parmenides' Pvi. The view might not be so easy to dismiss, however. What we need is an account of negative belief (ch. 11) and probably an account of how it is possible to believe falsely, as Plato considered in *The Sophist*.

Notwithstanding that later discussion, it is clear that the soft ontological Parmenidean has no compelling reason to adopt Pvi on the basis of Parmenides' arguments. As Gallop (1984: 24) shows, we cannot validly infer from Pi to the view that what does not exist, necessarily does not exist. The latter is needed to show that what does not exist *cannot* exist (the necessitarianism that I discussed under Pi, above), which in turn, Gallop says, is needed to show that what does not exist cannot be spoken of or thought of. Gallop perhaps assumes here that a merely possible existent can be spoken of and thought of. An argument should

also be given for that, but, it will be shown later, such a view is ultimately acceptable. The consequence of all of this is that in the version of soft Parmenideanism that I defend, Pvi is not included. There is an interpretation of Pvi that would be a more serious matter, however, as we will see below (ch. 8). If the issue is specifically reference, then Parmenides has identified a genuine problem since to refer to *x* implies that *x* exists. When *x* does not exist, we cannot refer to *x*. But if Pvi is not about reference, then what is it about? (Spoiler: it's about about.)

Pvii. Neither motion nor change is possible

...coming-to-be is extinguished and perishing not to be heard of. (8.21)

The above line is the conclusion of an extended argument running from 8.6 to 8.21. Gallop reconstructs it as follows: 'If there is genesis, then *either* the subject arose from what-is-not, *or* it arose from what-is. But it did not arise from what-is-not; nor did it arise from what-is. Therefore there can be no genesis' (Gallop 1984: 15). What-is is therefore ungenerated and imperishable. Parmenides offers further argument. What-is could not arise from what-is-not because:

...how could what-is be in the future; and how could [it] come-to-be?

For if [it] came-to-be, [it] is not, nor [is it] if at some time [it] is going to be.

(8.19–20)

Change involves *what is not* coming to be or *what is* ceasing to be, both of which Parmenides thinks are impossible. Change means that something comes from nothing, which he has already concluded against with Piii. For example, if we try to suppose that a tomato changes from red to green as it ripens, this means that its redness was nothing but then comes to be something, hence comes to be from nothing. As Gallop (1984: 17–18) states, the argument against motion is obscure and uncertain. As at 8.38–41, it might simply be the argument against change in general that is being applied now to position: one position perishes and another is generated.

Pvii is not a thesis that I would welcome into my soft Parmenideanism. One reason, almost too obvious to state, is that the reality of change is entirely apparent to us and can be known both empirically, through the evidence of our senses, and rationally through our ability to think different things at different times. Any a priori argument that rules such changes to be impossible has thus doubtless gone awry or assumes some even more counterintuitive premise, such as that time itself is illusory. The only question, then, is whether we have given a good account of change or whether, with Heraclitus (*Fragments*), we should take its reality as a basic assumption of our philosophy so that we can explain other things in terms of it. In a broadly Aristotelian account, there are powers that exist in potentiality but can come to act. A rich conception of reality, as containing both

potency and act, both being existent, at least has the virtue that it can account for change. The act does not come from nothing: it is explained by what there already is. Change need not be a variety of creation *ex nihilo*, as Parmenides depicts it. Nevertheless, it is clear that Pvii was an important part of Parmenides' philosophy since he devoted a number of lines to it in 'The Way of Truth'. To overturn it is to depart some distance from a hard ontological Parmenidean position. Nevertheless, we should be willing to do exactly that. It is the rejection of nothingness that is attractive about Parmenideanism—since that is reasonable—not the rejection of change, which seems unreasonable. And given that the latter need not follow from the former, there is no reason why we have to accept it.

Here, then, is the soft ontological Parmenideanism that I recommend, and will defend at greater length in the chapters to follow. I accept Pi and Pii: nothing is not and there are no degrees of being. That nothing comes from nothing (Piii) and that nothing is unknowable (Piv) are claims that are to be accepted with certain qualifications and interpretations. Whether being is a single plenum (Pv) is something that can be accepted or rejected irrespective of our other Parmenidean commitments and I shall not go into any further arguments concerning these. (For what it's worth, I have an attraction to some forms of monism and the idea that spacetime is dense, but I will not be arguing for either here. The discussion of holes might bear on the question of the plenum but not directly on monism, as we see in chapter 3.) That non-being is unthinkable and unsayable (Pvi) is something that I reject, as I do the thesis that change and motion are impossible (Pvii).

1.4 Soft methodological Parmenideanism

I now move to a different way of understanding Parmenideanism: one that builds a philosophical methodology aimed at preserving the commitments of Pi and Pii as far as possible while offering principles for evaluating any challenges that they face. This will provide a new way of understanding a distinction between hard and soft Parmenideanism and a way of proceeding with the arguments to come.

Let me call a strong methodological Parmenidean someone who accepts that there is no nothing, that nothing is not something, that everything that exists is 'positive', and that there are no negative entities of any kind, and, furthermore, these theoretical commitments for them are inviolable. The final clause is the crucial one since it then follows that in no circumstances would the strong methodological Parmenidean countenance admission of negative entities into their theory of the world or any part of it. (Note that I prefer to call methodological Parmenideans strong rather than hard.)

It is possible to adopt a softer version of methodological Parmenideanism. Such a position would proceed from the basis that commitment to negative existents is at the very least highly undesirable and ought to be avoided if possible. The

Parmenidean spirit is preserved in this respect. However, the soft methodological Parmenidean permits that there could be some circumstances in which negative existents have to be admitted, where there is no credible alternative to doing so. There would be a clear preference for not accepting negatives, but the view is conciliatory in that it allows that negative entities would have to be admitted if certain conditions were met.

Let us consider why anyone would want to adopt this softer position. First, it is surely too uncompromising in philosophy ever to say that one would never change one's view, no matter what the circumstances. It is reasonable to be open to compromise even with one's fundamental commitments. What is at the centre of the web of beliefs should be revised if an overall more attractive alternative web is found. One's fundamental commitments might be found to entail a contradiction, for example, or generate some other form of *reductio*. Even if a rejection of negative entities is a basic theoretical foundation, then, it should not be immune to revision. The key question is on what basis that revision should occur.

This brings us to a second point, which is that in holding any position, you ought to be able to say in what circumstances you would be prepared to relinquish it. If you were to say that there are no possible circumstances in which you would abandon a certain position, then it looks as if you are saying that you regard it as sacrosanct, unfalsifiable, and immune to rational evaluation; in other words, a dogma.

Suppose one is a Parmenidean, then, in respect of subscribing to Π . The position looks more reasonable if one allows that theoretically there are conditions in which one would admit a negative entity even if one believes that those conditions are not met. This is not to concede that Parmenideanism is a contingent truth. As this is a philosophical thesis, the conditions that concern us are argumentative, regarding the weight of philosophical considerations. Accordingly, I offer the following two conditions for admission of a negative entity:

1. Indispensability: the entity does work that is needed for our best philosophical and/or empirical theory.
2. Irreducibility: the entity is irreducibly negative, its being cannot be explained away in other terms that are entirely positive.

Given the discussion, above, I am suggesting that it would be reasonable to accept a negative existent as real if it meets these two conditions. This indicates a way forward as we consider a host of putative negative entities. The methodology that follows is that we look at the putative negative entities on a case-by-case basis and judge whether there are adequate grounds for them to be admitted into our theory of what there is: to our ontology. They will not be admitted if they are either dispensable or reducible. The soft methodological Parmenidean thus proceeds with an open mind, compared to the strong methodological Parmenidean. If the

best available theory of something—of causation, for example—needs absences to be real, and if they cannot be reduced away in terms of something positive, then we should reluctantly concede that absences are real. Our ontological Parmenideanism has to be relinquished in this area.

Why do we need both conditions? First, indispensability is required because there is the possibility of some irreducibly negative entity that is nevertheless dispensable, if it played no essential role in any of our best theories of anything. For example, suppose that one allowed a notion of negative anti-instantiation, as Barker and Jago (2012) do. It would be reasonable for a soft Parmenidean to argue as follows. Instantiation is a problematic notion anyway, as the Bradley regress shows. Perhaps there is a better theory than the traditional substance–attribute view since it has the problem of accounting for how substances and attributes are united. There might be a better theory (trope theory, factualism) of the world that dispenses with any notion of instantiation, and thus of anti-instantiation. Or, as a variant on this, perhaps we do need instantiation, but anti-instantiation is not an essential part of such a theory of the world since we can account for everything we must without needing to add anti-instantiation. Either of these would be defensible grounds for a soft Parmenidean to reject commitment to any such thing.

This condition derives from the indispensability arguments found in Quine (1976) and Putnam (1979). Theirs are arguments for the reality of mathematical entities from their indispensability to science. The premise is that we ought to be ontologically committed to the entities that are indispensable to our best scientific theories. It is this premise that is of interest since it is one that could be applied generally. Clearly it need not be limited to a defence of mathematical entities. Ought we not to be committed to the entities that are indispensable in any of our best theories, regardless of the subject matter of those theories? The soft methodological Parmenidean is primarily concerned with best philosophical theories and saying we should look to them for our ontological commitments. Unlike Field's (1980) interpretation, I am not proposing a nominalistic or other deflationary use of the argument. Indispensability to best philosophical theory seems as good a basis for ontological commitment that we can get so I am happy to interpret any resulting commitment realistically. Nor am I put off by Maddy's (1992, 1995) charge of anti-naturalistic confirmation holism, especially given that we are dealing with philosophical theories that tend not to be subject to empirical anomalies.

The second condition I apply is irreducibility. There could be indispensable entities in our best theories that on further analysis were reducible to positive entities, or could be replaced by something else that was positive but that played the same role in the theory. For example, we might be tempted to allow the reality of a negative property as the best way to account for negative predications, such as when we say that something is non-red. Suppose, however, that instead of taking non-redness as a single thing, one could reduce it to a disjunction of other colours, with the exception of red. Thus, instead of treating non-redness as a

property, one could say that something has the property of being green or blue or orange or turquoise or... (for reasons given by Gale 1976: 12–16, I don't believe this reduction actually works; but suppose it did). In that case we would have no need to accept the irreducible reality of non-redness. I take this as a separate condition so that the soft Parmenidean has a second way out of admitting negative entities. It might well be that the best all-round theory invokes non-redness. A theory which reduces it to a disjunction of positive properties might not be as good in a number of ways: for instance, it is theoretically messier and it seems to admit open-ended disjunctive properties. But if the negative property of the theory can be reduced to others that are positive, then the soft Parmenidean should accept the way out in order to avoid negative entities.

1.5 Is the cat out of the bag?

I suggest that if both the indispensability and irreducibility tests are met, it would be dogmatic and irrational for us not to accept the reality of a negative existent. A strong methodological Parmenidean wouldn't: but this is where such a position looks unreasonable. Soft methodological Parmenideanism is a position of compromise. Let us not forget, however, that this is only a methodology. The soft Parmenidean proceeds in the hope, and maybe the expectation, that when we go through each of the areas in which negative existents have been posited, that none of those negatives turn out to be both irreducible and indispensable in the best theories of those areas. If that is the end result, then the soft methodological Parmenidean will also have saved a soft ontological Parmenideanism. The compromise, however, is that if a negative entity has consistently resisted attempts at reduction, and no better theory can be found that is without negative entities, then the soft methodological Parmenidean concedes that the negative entity should be admitted.

It might be wondered, however, whether the methodological Parmenidean is compromised in another way too. Suppose that, once all the considerations have been assessed, some kind of negative entity has to be admitted because the best theory of X requires it and there is no reduction available. Then isn't the cat out of the bag? It is an admission that there seems to be at least some negative existence. And, if that is conceded, then it seems that the soft Parmenidean no longer has any grounds to discriminate against further negative entities simply in virtue of them being negative. Of course, any entity must do theoretical work to earn its place within our ontology, so Occamist tests still apply. If such tests are passed, however, it seems that the soft Parmenidean no longer has any principled grounds on which to continue discriminating against negative entities on the basis of their negativity. A charge could be brought, then, that soft methodological Parmenideanism is an unstable position and it should be abandoned once one negative existent is granted.

Nevertheless, the soft Parmenidean should reject this charge. The rationale for methodological Parmenideanism is the ontological Parmenidean thesis P_i , along with P_{ii} . Methodological Parmenideanism is adopted because negative entities are, for various reasons, undesirable in any theory. If the methodological Parmenidean at some points has to admit a negative entity into their ontology, it need not involve them conceding that negative entities are suddenly any less mysterious than we thought they were, or any less undesirable. The expectation would be, instead, that the Parmenidean simply move to a next methodological phase. This is again a compromise, but it would presumably be that while some negative entities are accepted, because there is no better theory that is free of them, if at some future time a new theory is found that is otherwise at least as good as the first, but involves no negative entities, then the new theory is to be adopted. The preference for elimination of negative entities can thus consistently remain. Given this preference, it shows how negative entities should not be multiplied willy-nilly, just because one kind of them has been accepted. The soft methodological Parmenidean can treat them as theoretical anomalies that might receive a solution further down the line where any such list of anomalies is best kept to a minimum.

Now it might be wondered whether soft methodological Parmenideanism is a novel position. Has anyone ever held both that negative entities are undesirable but that they can nevertheless be accepted in some circumstances? Despite the aforementioned tension in the view, it is a position that realists and naturalists in metaphysics have frequently adopted. We could take Russell and Armstrong as examples of realist-naturalists, for instance. It is notable that both reluctantly accepted some variety of negative existent after due consideration. Russell (1918) saw that general facts were indispensable and negative. From a collection of facts such as $Fa, Fb, Fc, \dots Fn$, we cannot derive that $\forall x(Fx)$, since the existence of $Fa, Fb, Fc, \dots Fn$ is consistent with the existence of something else that is non- F . What-there-is does not necessitate that it is all-there-is. Russell conceded that we had to postulate *general facts* in addition to first-order particular facts and that these general facts were irreducibly negative, since they say that there are no more first-order facts than those. Armstrong (1997, 2004) followed this by accepting *totality facts*—higher-order facts that there are no more first-order facts—in order to account for the truthmakers of negative truths. In the cases of both Russell and Armstrong, therefore, we have good grounds to classify them as soft Parmenideans in that they both saw it as desirable to eliminate negative existents but were prepared, reluctantly, to accept them. And among contemporary philosophers, we find that when Schaffer (2004) accepts that absences are causes, and Hommen (2013, 2018) accepts that there are negative properties, they do so with implicit admissions that such things are not desirable in themselves but, rather, negatives that we have to admit because there is no other option. Our best theories demand it. It might even be, then, that the soft methodological Parmenideanism is a commonly held view on the right way to proceed when tackling the difficult questions of negative existence.

1.6 The way forward

We have seen how a distinction can be drawn between ontological and methodological Parmenideanism, with hard and soft varieties of each, but all united by the Parmenidean thesis that nothing is not. I have outlined why soft Parmenideanism is the more reasonable position to hold, both as a set of theses held and as a way forward in tackling the problems created by putative negative existents.

The methodological thesis offers us a way to address the difficult matters that will follow in the rest of this book, and it will be seen that the soft Parmenidean tests will be constantly in the background. When we look at each new issue, we will be considering whether we are really compelled to accept a kind of negative existent or whether we should reject it because there is either a better theory available or some reduction of the negative to the positive is possible. This will allow us to assess whether absences are things, whether holes and edges are negative, whether absences are causes, whether they can be perceived, whether there are negative truthmakers for negative truths, and so on.

At the conclusion of the process, it might turn out that all major putative negative entities can be explained away. The soft ontological Parmenidean could then rest content with that outcome. The soft methodological Parmenidean should not presume that this finding will triumph, however. Even if there is a preference against negative entities, there should not be a presumption that the arguments will be available that rule them out. We proceed with an open mind in at least that respect.

2

Negative Properties

2.1 What would be a negative property?

Many things have a property of being red. A stick could be 10 cm long, a person could have a temperature of 37.1°C, and a stone could weigh 2 kg. The reality of other putative properties is a more controversial matter, including whether there are negative properties. Can something have a property of being non-red? Do we want to allow non-redness as a genuine property, just as real as being red? Nominalists, of course, do not accept the reality of any properties, but there might be reasons even for realists to resist negative properties. An object is blue, for example, and one might think that there is nothing more to us calling it non-red than that it is blue. Since blueness and redness are contraries, then something that is all blue cannot also be red. Perhaps, then, it is possible to do without any negative properties if purely positive properties can explain them all away. Maybe one can draw a distinction here between not having a property, F, and having a negative property of being non-F. The former could be acceptable, while the latter might commit you to the sort of dubious entity against which Parmenides warned.

Our soft Parmenidean project requires that we seek to explain what we can about properties without invoking negative existents. If our best theory of properties leaves us no choice but to accept that there are negative properties, then we would have failed at this hurdle. The hope, instead, is that we can say all that we need to say without invoking irreducibly negative properties. It is hardly worth denying that negative predications are made: for example, when something is said to be not-10-cm long. But whether a negative property is needed to account for the use of negative predication is another matter. Perhaps all the ascriptions of being not-10-cm long can be explained in terms of what there is: what lengths positively exist. Or perhaps the negative predication can be explained away in some other manner that doesn't commit us to allowing a genuine negative property. Such a strategy might be forthcoming.

This discussion arises within the common framework in which both particulars and properties are accepted as elements of the world. Particulars are sometimes said to bear their properties. There are various theories about what properties are: whether they are Platonic transcendent universals, Aristotelian immanent universals, resemblance classes, tropes, and so on. Not much of what follows depends on

these questions of *what* properties are. The concern is with what properties *there are*. Similarly, there are views on the merits of a substance-attribute ontology—another way of saying particulars and properties—and whether there are better alternatives. Perhaps things are not composites of particulars and properties. There is process metaphysics, for example, which sometimes professes to offer a radical alternative to the conventional framework. Nominalists, of course, do not accept that properties are real: there are only particulars. I will not consider this debate, however, as nominalism versus realism is far too big a conversation for here. In any case, for nominalists there are similar questions, such as whether there are negative resemblance classes corresponding to negative predicates, so many of the arguments I will be discussing apply there too.

In approaching our problem with an open mind, applying the methodology detailed in chapter 1, it is worth considering whether an intuitive dislike of negative properties has a justification here or whether it is only a matter of prejudice. We want to know whether there are any direct arguments against the existence of negative properties. It is a further question whether there are irresistible arguments in favour of negative properties, such as whether they are indispensable. Nevertheless, there is a lot to learn if we start with the case against.

2.2 Arguments against negative properties

Armstrong presents four direct arguments against the reality of negative properties, after which he offers his own account of negative predication. It would be useful for the Parmenidean project if we could endorse these arguments. Regrettably, we cannot. I will present each argument in turn, followed by some comments. The importance of doing so resides in what the failure of Armstrong's arguments shows us, which includes, among other things, just how difficult a task it would be to eliminate negatives from our ontology. Hence, we have a demonstration that the problem of negativity is not an easy one.

The arguments are presented by Armstrong while developing a theory of immanent universals, so he speaks sometimes of negative properties and sometimes of negative universals. The distinction between properties and universals can matter (see, for instance, Lewis 1983), but it does not matter here. Whatever properties are—whether they be real universals or something else—it is the possibility of them being negative that is our concern. Accordingly, I will stick to use of the term property, even where the wording of an argument has to be adapted from one originally given in terms of universals. I will then follow the direct arguments with a different kind of consideration: the unnaturalness of negatives.

2.2.1 First argument

The basis of all of Armstrong's arguments is a rejection of what he calls the argument from meaning (Armstrong 1978b: 11). Quinton (1957: 33) articulated the view, which he too rejected, that to every predicate there corresponds a property. If there is a predicate 'P', then there is a corresponding property P, where a property is some real worldly, non-linguistic existent. Now Armstrong accepts (1978b: 19) that there are negative predicates but rejects any direct inference to the existence of corresponding properties. What properties are real is to be decided by a weight of other considerations. But an argument from meaning can be rejected since it is plausible that there can be predicates without corresponding properties, properties for which there are no predicates, multiple predicates for the same property, and multiple properties for which we use the same predicate. The first case is the key one for the current discussion. I take it to be clear that there are predicates without corresponding properties (such as being phlogistonated), and we need to enquire now whether this is the case for negative predicates. There are the predicates 'not-green', 'not-6 feet tall', but are there, or are there not, negative properties corresponding to these?

Employing a shorthand, I shall use $\neg F$ to denote a simple negative property: namely the property of being non-F, where F itself would be the complementary 'positive' property. Non-red and red would be an instance of such a complementary pair. Armstrong's first argument, to paraphrase the original presentation (Armstrong 1978b: 23), is that when a property predicate applies to various particulars, it implies that they are all identical in some respect, which means that they resemble each other. Therefore, when $\neg F$ applies to different particulars, it implies they are identical in some respect, which means that they resemble each other. But many particulars could all be $\neg F$ without resembling.

This argument, like all those that Armstrong presents, has *prima facie* credibility. First of all, it has credibility in the case of the 'acceptable' real properties. All things that are red are identical in some respect: they share the same colour. All things that are 10 cm long are identical in some respect: they share the same length. But if one considers all the things that are non-square, it includes triangles, circles, and trapeziums, as well as the multiplicity of irregular shapes possessed by ordinary objects, such as tables, chairs, and cats. It must also be true that reflections, ideas, events, bodies of gas and water, are non-square. An idea is not square, we can certainly say; indeed, it would be a category mistake to attribute any shape to it at all. There is, as Armstrong sees it, an identity that comes from the sharing of a property. There is a *One that runs through the Many*. His view is that redness is one single 'thing' (a property, he says, is a higher-order particular), yet it is found in many different things: the things that instantiate redness. All this seems to hold of the real properties: those that are

‘positive,’ we might have to say. But it does not hold of negative properties, the argument alleges.

This is what Crivelli (although he puts it in terms of negative kinds) calls the unity problem, which he summarizes so:

Such a presupposition [that there are negative kinds] is dubious. What characteristics are shared by emeralds, numbers, Buckingham Palace, and generosity? Still, they should all fall under non-whiteness, were there any such kind.

(Crivelli 2012: 188)

As Crivelli’s study of Plato goes on to show, however, Armstrong’s argument here is ignoring Plato’s ‘parts of difference’ account from the *Sophist*. According to Crivelli, as seems confirmed in the text of the *Sophist* (255b8–c8), Plato offers the view that there will indeed be something in common to all things that are non-F; namely that they are alike in being different from F. Hence, the circle, triangle, cat, table, and idea have in common that they are all different from squareness. This is a One that runs through those Many.

Now it would be quite natural for a follower of Armstrong’s view to protest that this alleged thing-in-common to all non-square things is not genuinely in-common. But it is hard to see how such a view can possibly be asserted without begging the very question at issue, namely, whether there are negative properties. To make the point bluntly, one can exclude being different from squareness, or being not-square, from being a thing in common to all the non-squares, only if one has already ruled out negative properties. But that is exactly what is at issue. One cannot assume against negative properties in what is purported to be an argument against them.

This is sufficient to defeat Armstrong’s first argument against negative properties. But it also raises a further defect of invoking a resemblance-based argument for adjudication of what is a genuine property. Any resemblance between two or more particulars will inevitably, either immediately or ultimately, be based on the properties of those particulars. In brief, *a*, *b*, and *c* resemble when they have a common property or quality. Aliveness is always aliveness in some respect, and one standard way of understanding ‘respects’ is that they are properties. This is at least the view of any broadly realist metaphysics (we will see, under *Naturalness*, below, that a different view is possible). Thus, the way in which all red things resemble each other is that they are red and it is very hard to see how there could be any other plausible account of what their resemblance consists in (I am, of course, aware that nominalists think they can resolve this matter). It is thus inevitable that when we consider the respect in which all non-F things resemble, it is that they resemble in being non-F. But even if this is a plausible account of resemblance, it cannot be the basis of a criterion by which to separate genuine from spurious properties since the way in which diverse particulars of a kind will

resemble will be precisely in virtue of that property in question. This explains why Armstrong's first argument is bound to be question begging. It challenges the credentials of $\neg F$ to explain resemblance, then rules that it cannot, but only by assuming that $\neg F$ is unacceptable and thus cannot be legitimately a basis of resemblance.

2.2.2 Second argument

Armstrong's second argument originates in McTaggart (1921: §62). If we permit negative properties then everything has the same number of properties, since for every supposed complementary pair of properties F and $\neg F$, every particular must be either F or $\neg F$. For example, if there were three particulars, a , b , and c , and ten complementary pairs of properties consisting of $F, \dots O$ and $\neg F, \dots \neg O$, then a , b , and c will have the same number of properties, namely ten, consisting of one of F or $\neg F$, one of G or $\neg G$, and so on through O or $\neg O$. This conclusion—that all things would have the same number of properties—will hold, of course, no matter how many particulars there are and how many properties.

McTaggart was happy with this conclusion, apparently, but Armstrong instead presents it as a reductio. We cannot know a priori that two particulars have the same number of properties, he insists. An empiricist should say that this question is always an a posteriori matter. Presumably, one should look, examine, and count a thing's properties in order to know how many it has. Only after repeating such an empirical examination could one conclude of two things that they have the same number of properties. An assumption that is responsible for generating this unpalatable conclusion is that there are negative properties that complement each 'positive' property.

That Armstrong considers this a reductio seems to bear some connection, though not explicitly stated, to the so-called Irish principle, which he invokes in several places (e.g. Armstrong 1978b: 11). The principle dictates that if you can know a priori that a particular bears some property H , then it has no such property. For example, it can be known a priori that everything is self-identical: so this tells us that self-identity is not a property. Whether something has a property must be an a posteriori matter in any broadly empiricist approach (although there might be at least some exceptions: one might know a priori that one has the property of being a thinking thing, if Descartes was right). It seems to follow, then, that whether particulars have the same number of properties must be an a posteriori matter.

Armstrong accepts a restriction on McTaggart's argument, however, since it seems to apply only to particulars and properties that are of a matching category. Virtue, for instance, is neither circular nor non-circular. Virtue is not of the right category—objects extended in two dimensions—to be either circular or

non-circular. So here we have a particular, v , for which none of a complementary pair of properties H and $\neg H$, is applicable. Armstrong says, however, that this restriction makes no significant difference: and this seems to be right. For it follows that if a and b are of the same category of existent, then they will have the same number of properties, and this conclusion alone is enough to offend Armstrong's stricture against a priori knowledge of property enumeration.

The second argument can be challenged for much the same reason as the first, however, namely that it is question begging. The way in which it begs the question against negative properties is complicated. In the first place, who is to say that things don't have the same number of properties, namely an infinite number? This is not merely suggesting that a potential infinite number of predicates apply to a thing, which is undoubtedly true. The point can have more ontological seriousness than that for it could just as well be stated in terms of properties. How many properties does the room have in which you are seated, for instance? The answer seems limited only by our finite abilities to know and to describe. Hence, there seems no immediately obvious reason why things must have a finite number of properties, so an argument based on the absurdity of things having the same number of properties would be insecure. The point raises the prospect of perhaps an even greater absurdity, however. One can wonder what sense it makes to enumerate properties at all. Are they the sort of thing that could sensibly be counted in a principled and consistent way? Are red and scarlet one or two properties, for instance? They are non-identical, which makes them look like two, and yet something that is scarlet needs nothing in addition to be also red.

However, let us put aside these misgivings and try to view the argument sympathetically. We might grant that there is at least something in it. An empiricist might say, for example, that a billiard ball must have fewer properties than a battleship, since a battleship is a very complex particular with all sorts of additional properties that something as simple as a billiard ball does not have. Restricting ourselves to non-relational properties, since relational properties might be reducible to non-relational properties anyway (Armstrong 1978b: 78), one could give an almost complete description of a billiard ball in terms of its properties very briefly. It has a particular colour, mass, shape, size, and composition, and perhaps not much else. In contrast, a complete description of all the properties of a battleship would be almost a lifetime's work for someone, so intricate is such a particular.

But even viewed in this sympathetic way, we are still right to resist the proposed reductio. The argument must be question begging, for whether there are negative properties is a bigger, more fundamental, hence prior question to whether particulars have the same number of properties. An answer to the latter question shouldn't determine an answer to the former. If there is to be any priority, it should surely be the other way around; that is, if one is to consider whether particulars have the same or different number of properties, one question to address is whether there are negative properties. McTaggart has given us a reason

why the question of negative properties will affect our enumerative question. If there are negative properties, and a consequence is that particulars have the same number of properties, then so be it. We have an initial grasp on what it is to be a negative property—the one that Plato articulated, for instance—whereas, for reasons given above, there is less of a grasp on how we should properly enumerate properties. There seems less legitimacy, in contrast, in accepting a principle that particulars should not be judged on a priori grounds to have the same number of properties and using this principle to reason that there cannot, therefore, be negative properties. The cleavage between being and non-being is the most basic. It is a more fundamental question of philosophy whether there are negative existents. There must be very solid grounds for the rejection of negative properties. And if there are negative properties, whether particulars have the same or a different number of properties is spoils to the victor, relatively speaking.

Despite that, if the empiricist's concern behind this second argument is that, on the basis of properties, one should be able to provide an account of how some things are more complex than others, then this might still be possible even if we admit negative properties. An account could be outlined in the following way. A more complex thing could have more positive properties than a relatively simple thing has. A simple thing would always have more negative properties than a complex thing. A battleship has more positive properties than a billiard ball has. The battleship has the properties of being staffed, of floating on water, and having gun turrets. The billiard ball would have the complementary negative properties. As suggested above, the billiard ball seems to have relatively few positive properties: certainly relative to a battleship. So simplicity can consist in having relatively few positive properties and complexity can consist in having relatively many. Furthermore, even if the basic McTaggart-Armstrong argument is correct, it is surely an a posteriori matter which of the complementary pair, *F* and non-*F*, a particular has. And it will thus be an a posteriori matter how many positive properties a particular has. Hence, Armstrong's claim that the number of properties a particular has, even if we set aside all misgivings about the legitimacy of such a notion, must be a posteriori known, can be maintained in a form that is only slightly amended if there are negative properties. We simply say that some things have more positive properties than others, which is known a posteriori. That particulars will also have negative properties, and the total number of positive and negative properties together will come out as the same for each particular, seems a less staggering consequence than was first presented. We can conclude, therefore, that the second argument also gives us no compelling reason to reject negative properties.

2.2.3 Third argument

Next, Armstrong (1978b: 25) argues that we should be able to attribute causal powers on the basis of properties. But there don't seem to be any negative causal

powers. Causation by absence, for instance, can be dismissed out of hand, for if you allowed it, you'd have to say that lack of poison is causing you to be alive. An Occamist argument is deployed here: if positive properties are capable of doing all the causal work that is needed, then there is no need to accept negative causal powers besides.

This argument requires a little more attention. The intuition behind it seems perfectly reasonable. There is something very suspect about causation by absence. As we will see (chapter 4), however, causation by absence is not entirely simple to dismiss. Further, providing a *reductio* of causation by absence is not the same as showing how causation can work without it, and it seems no simple matter to demonstrate that we can account for all cases of causation in exclusively positive terms.

The starting point of the third argument is one that is reasonable, however, and I accept it. There is indeed a close connection between the causal powers of things and their properties. The foundation of such a view is commitment to something like the Eleatic reality test (Plato, *Sophist* 247d–e). To be real is to be powerful, and so a property is real only if it adds a causal power to its bearer. For example, that which is spherical must have the power to roll in a straight line down an inclined plane. There is also a developed version of this view supported by Shoemaker (1980 and explored in Mumford 2008). In this, we can understand properties as clusters of causal powers. As some think of it, each property bestows a unique set of causal powers on its bearer. I think the most plausible explanation of this is that the property is the cluster of powers. The powers provide the essence of the property—they determine its identity—hence properties and unique clusters of powers stand in a one-to-one relation. Two properties that differ in their causal powers are thus different properties.

Armstrong's argument, however, moves from this acceptable starting point immediately to a discussion of causation by absence. Notwithstanding the more detailed treatment to follow in chapter 4, below, let us accept for the moment that causation by absence is unattractive and so it would be a disadvantage of any theory of properties if it were committed to it.

We need to back up a little, though. If we retain the view that we should be able to attribute causal powers on the basis of properties, as we should, then let us call the causal powers that we would attribute on the basis of negative properties, were such things to exist, negative causal powers. Now it is not obvious that the only way of understanding such negative causal powers is in terms of causation by absence. Presumably, Armstrong is thinking that if a particular could really have a property such as being non-spherical, then it is its absence of sphericity that would be powerful. This instantly sounds absurd, for how can an absence of anything causally affect something else? What, though, if we understand negative causal powers in a different way? Would it be plausible, for instance, to understand them as powers *not* to do things, since 'positive' causal powers are powers to do things?

The idea of a power not to do something seems coherent. We can think of insolubility, for instance, as a power not to dissolve and resilience as a power to withstand pressure, remaining unchanged. Think also of inertia and immobility. Furthermore, while causation by absence seems problematic, causation of an absence is acceptable. Often, an effect is an absence of change, such as when two opposite powers are operating that balance each other out perfectly and, as a result, nothing happens.

However, the problem we have here is that it is implausible that a negative property must involve a power not to do something. There is a distinction to be made between not having a power to F and having a power to not F, and it is the former that would most plausibly follow from negative properties rather than the latter. The difficulty is that while having the power to not F might plausibly be understood as the having of a power, not having a power to F is not a power. Clearly, it is a case of *not having* a power, contrary to the starting point of Armstrong's argument, which was that having a property involves having a power.

In case anyone is sceptical that there really is a distinction between not having a power to F and having a power to not F, we need only think of the distinction between soluble, insoluble, and non-soluble. Because something is non-soluble, we cannot assume that it is insoluble. We can distinguish a soluble from an insoluble aspirin, for example, but ideas, or numbers, while non-soluble, are not insoluble. Insoluble means that something can be placed in liquid and resist dissolving. Neither an idea nor a number can be placed in liquid and nor does it seem right to say, in any circumstances, that they resist dissolving.

We might, then, have arrived at a better way of supporting Armstrong's third argument than Armstrong himself gave. If the most sensible understanding of 'negative causal power' is one that resolves into no causal power at all, then it suggests that the tie between properties and causal powers will indeed be broken. And it seems reasonable to think that the having of a putative negative property should be allied to not having a causal power rather than allied to having a power not to do something.

2.2.4 Fourth argument

Armstrong has a fourth argument. If we allow that conjunctive properties are real, which he thinks we must, then if we also allow negative properties, then we would have to allow disjunctive properties as well, which are even worse than negative properties. The reasoning is that a complex property that had the form $\neg(\neg F \ \& \ \neg G)$, which is a combination of negative and conjunctive components, would be logically equivalent to the disjunctive property $(F \vee G)$. Armstrong has various arguments elsewhere, which I will not detail, for why there are no

disjunctive properties. But he accepts conjunctive properties since it raises no great difficulty, if one allows F and G to be real properties, that there could also be a property which is the conjunction of them.

Just as his arguments against negative properties are contentious, so are Armstrong's arguments against disjunctive properties. I will not engage with them, however, since I think that this fourth argument against negative properties can be rejected even if we grant as an assumption that disjunctive properties are unacceptable. The key point is that logical equivalence need not mean an ontological equivalence. After all, Armstrong allows that F can be a real property even though it would be logically equivalent to $\neg\neg F$.

Suppose one thought that, despite including negation, $\neg\neg F$ is really just the positive property F , since they are equivalent, so there is nothing wrong with $\neg\neg F$. But then Armstrong's opponent would be at liberty to make exactly the same move and say that if $\neg(\neg F \ \& \ \neg G)$ were to be acceptable (by assumption, in accepting both negative and conjunctive properties for purposes of argument), then the disjunctive property it entails, $(F \vee G)$, must also be acceptable since it is logically equivalent to something acceptable. Whenever we had $(F \vee G)$, we could just treat it as the acceptable (by hypothesis) $\neg(\neg F \ \& \ \neg G)$. If this move were to be questioned, one need only point out that in considering $\neg\neg F$, we were happy to treat it as F , and this shows that an ontological consideration regarding the acceptability of F led us to understand $\neg\neg F$ in that less concerning way. But if we set such a precedent, then we should of course be free to apply the same principle elsewhere. If we do not like $(F \vee G)$, then we seem justified in treating it as something else that is logically equivalent, such as $\neg(\neg F \ \& \ \neg G)$. This fourth argument of Armstrong's, then, does not carry weight.

Having considered Armstrong's direct arguments against negative properties in detail, we find a mixed bag of results. The first two arguments seem to beg the question against negative properties; the third might carry some weight if interpreted in the way that I have offered, though this was not Armstrong's own interpretation; and the fourth argument carries no weight at all.

2.2.5 Naturalness

There is a different approach to negative properties, which is less direct but nevertheless could be powerful. This is found in Quinton (1957), Quine (1969), Armstrong (1978a: 38–41) and Lewis (1983) and says that some properties are natural and some are unnatural, with the option to say that naturalness can come in degrees. We are then at liberty to suggest, as Quinton does, that while the non-blue things form a class, there is no corresponding natural property of being non-blue.

We might well ask what makes a property natural and wonder why negative properties are not natural. The answer is that there is a strain of primitivism going through this view. Quinton (1957: 47), for instance, says the existence of genuine natural classes is just a brute fact, but one that we are able to notice. We can see the distinction, for instance, in the fact that it would take fewer 'specimen instances' to grasp the idea of blueness than it would to grasp the idea of non-blueness (Quinton 1957: 58), which is one of the things that makes non-blueness unnatural. The story would go, presumably, that in seeing a blue book, a blue flag, and a blue shirt, one can quickly see upon what basis these things form a class, and hence it is natural that they do. But when one sees an old shoe, a turquoise blanket, a fire extinguisher, grass, a cardboard box, together with many other things, one still might not grasp that what they have in common is that they are non-blue. They are not a natural class.

One might not be persuaded that this constitutes a conclusive reason for accepting properties *qua* natural classes of things, and therefore a reason against negative properties, but Lewis (1983) puts some flesh on the argument's bones. He is openly primitivist about what counts as natural. There is probably little choice. For reasons given above, one cannot say without circularity that there is a resemblance among natural properties but not among unnatural properties, since one cannot rule out non-F as a way of resembling unless negative properties have already been excluded. Naturalness is being considered as a basis for excluding negative properties so cannot itself be understood by already assuming against them. But what justifies assuming naturalness as a primitive is the work that it can do. It explains laws of nature, for example, which assume natural properties rather than gerrymandered ones, and it explains the content of language and thought. Without the privileging of some properties as natural, or more natural than others, such content would be deeply unstable, as Lewis details.

This kind of argument is common in metaphysics and Lewis (1986a) has employed it elsewhere to argue for the plurality of worlds. We should look not for direct arguments in metaphysics because it is rare that they are available and even rarer that they are conclusive. Rather, we should look at the explanatory power of a metaphysical assumption. In the case at issue, the claim is that there is explanatory power in taking a class of properties to be natural: a class that does not include supposed negative properties. But there remains something slightly unsatisfying about this approach. Explanatory power sounds like an epistemic consideration about which a realist should have reservations. The world is not a slave to our theories. The direction of fit really ought to go the other way. There might, however, be a metametaphysical methodological point in favour of Lewis's approach here. Sometimes, explanatory power might be all we have to go on, in which case it will be rational to base theory choice upon it. Even then, there is a thought that direct arguments would be better. And if there is incoherence in a notion, then a *reductio* argument could yet be brought to bear.

2.3 For negative properties

It is one thing to consider arguments against there being a particular kind of entity. Even if there are few or no good arguments against something, however, that is not automatically an argument in favour of there being such an entity. There might be no evidence against the existence of a certain dinosaur that left no fossilized remains, but that does not make a good argument for the existence of such a creature. We ought to consider, then, whether arguments exist in favour of there being negative properties. Fortunately, such arguments have been offered, which we should now assess.

2.3.1 Negatives as lesser beings?

Negative properties retain some contemporary champions. Zangwill's (2011) support contains important qualifications but certainly admits that negative properties are real. The key qualification is that negative properties are less real than positive properties, which is itself a controversial view concerning degrees of reality. He calls this view *inegalitarianism*. Hence, Zangwill clearly rejects *Parmenides'* Pii. Zangwill is, by his own admission, a reluctant realist about negative properties but thinks that they cannot be denied since they have a 'causal and metaphysical determining role' (Zangwill 2011: 528).

However, it is fair to say that insufficient grounds are offered for realism about negative properties here, since the reasons Zangwill offers do not really count in favour of them. Furthermore, the idea of there being degrees of reality, in which negative properties can cling to a little bit of it, is not convincing. Some of the arguments provided are better understood as grounds for denying negative properties any reality at all.

We should admit that negative properties are real, says Zangwill (2011: 259), because, first, they are among the necessary conditions of positive properties doing their work. For example, rain makes you wet only if you do not have an umbrella (you have the negative property of not holding an umbrella). This is a form of causal determining even though it is weaker than the determining power of positive properties: so we should allow that negative properties are real but not as real as positive ones. Second, there is also a metaphysical determining role for negative properties. This determining role is again weaker than that of positive properties because while negative properties can determine other negative properties, they cannot determine positive properties, but positive properties can determine negative properties. For example, if something has the property of being all red, then it is determined that it has the property of being non-blue, whereas if something is non-blue, it does not determine that it is red, or green, or yellow, or any other positive property. Zangwill calls this the *asymmetrical*

determination thesis. A negative property can only determine another negative property: for example, if something is not coloured, then it is not blue. Determination can seemingly come in degrees, which Zangwill takes as a good basis for saying that something is more or less real. Here, Plato is cited, who takes the forms to be more real than material things since the existence of material things depends on the existence of the forms, but the existence of the forms does not depend on the existence of material things (Zangwill 2011: 530).

This line of thought should not persuade us, however. Suppose we grant that there are degrees of determination, in that some types of things determine more than other types of things, as Zangwill says that positive properties do relative to negative properties. And suppose that we allow that there can be asymmetrical dependence relations, as is undeniable. This would still not establish that negative properties are less real than positive properties. There being degrees of reality might sound reasonable superficially, but it does not withstand scrutiny when we press what it must really mean. We are talking about being or existence, when we say 'real', and as soon as we switch to that language (which Zangwill does only as the exception), the power of inegalitarianism seems to drain away. For instance, suppose it is granted that someone's left little finger asymmetrically depends on that person. I assume that once the finger is detached from the person, it dies and withers away. The finger can exist only if the person does, but the person could exist without the finger, in case of amputation. This does not mean that the finger exists any less than the person if, as a matter of fact, they do both exist. That the finger depends for its existence on the existence of the person still means that they both exist, and they exist absolutely and unequivocally. The person might be more important than the finger, if the finger depends on the person and not vice versa, but having a greater degree of importance is not the same as having a greater degree of existence. We can grant all sorts of asymmetries between things without allowing degrees of existence. The issue seems to be that, like 'unique' or 'perfect', 'exist' does not admit of degrees and nothing that Zangwill has offered gives us a reason to think otherwise. Indeed, what makes the person more important than the finger is the asymmetric determination, but this is an existential determination, the force of which depends upon the fact that, when the finger does exist, it exists as much as anything else. Existence is not a matter of size or importance since a speck of dust can exist just as much as the Eiffel tower. A greater size does not mean a greater existence.

2.3.2 An indispensability argument for negative properties?

Zangwill has not salvaged negative properties by allowing them some lesser degree of reality, therefore, since it is far from clear that he can help himself to any such notion, once we see it for what it is. Of course, the concern here has been the

view that existing to a lesser degree is a way of vindicating what seems disconcerting about negative properties. To that extent, it is Parmenides' Pii against which Zangwill stands. Let us, then, shift the focus back to the simpler claim that there are negative properties, which challenges Pi? It is doubtful, however, that Zangwill has given good reasons to accept that negative properties exist at all, in any degree.

We can consider Zangwill's approach in the light of our account of soft Parmenideanism in chapter 1. Zangwill's offers an indispensability argument: we have to allow for negative properties in order to make sense of some cases of determination. But let us consider the instances of causal and, what Zangwill calls, metaphysical determination separately.

We can start with causal determination, although this matter will be treated in more detail later (chapter 4). The argument here is that there is some causal role for negative properties to play because, although they might lack causal power, there are at least some true causal attributions in which negative properties feature; for instance, you can get wet because (you have the property) of lacking an umbrella or rainwear. Zangwill notes that this seems to make the conditional true that had you been carrying an umbrella (or wearing suitable rainwear), you would have stayed dry. Are these adequate grounds for attributing a causal role to negative properties and thus for allowing that negative properties are real? I am not going to contest the implied inference from having a causal role to being real, since that is merely an application of the Eleatic reality test, which I accept. But I think that considerations of this kind do not make the case for a causal role for negative properties.

We will see later that there are ample complications to consider, but the following intuitive response is basically sound. It is not the absent umbrella and rainwear that makes you wet, it is the rain. The only appeal of a causal role for something negative is the thought that had an umbrella been present, you would have stayed dry. But there is no negative entity, including no negative property, that is responsible for this. There is merely a counterfactual truth. What makes the counterfactual true is the power of the umbrella and rainwear to keep its users dry. It is certainly not that the negative property, of not being properly equipped, makes you wet. There might be some explanatory value in saying that you became wet because you were not properly dressed, but this is plausibly a non-causal explanation the value of which resides only in the counterfactual that, had you dressed properly, you would have stayed dry. It is hard to see any good reason here to invoke negative properties as an essential or indispensable part of the explanation.

The attraction of negative properties having a role in so-called metaphysical determination is different. As already mentioned above, the thought is that a positive property, such as being all red, determines a negative property, such as being non-blue, but the determination does not go the other way since something that is non-blue could be any other colour. Negative properties can determine only other negative properties. However, this sort of consideration is again inconclusive.

It might be that something being entirely red means that it is not blue, but this is insufficient to say that being non-blue is itself a (negative) property. As we saw, if the argument from meaning is rejected, then there is no need to allow a property for every predicate. It can be true that something is not blue without there being a property corresponding specifically to that thing's non-blueness. It could be that the thing having the genuine (positive) property of redness serves as the truthmaker for it being non-blue (see also chapter 9, on negative truth). There is a credible account of how this is so. Being red is a determinate colour: one of a number of determinates at the same level of determinacy, such as green, blue, yellow, and orange. What is distinctive about determinable properties, such as colour, is that there can only be one of the respective determinates at the same level of determinacy instantiated entirely by a particular at a time. So a colour can only be red, or blue, or green, etc., where the disjunction is an exclusive one. It follows that, if something is entirely red, then it is not any of the other determinate colours. We can explain why being entirely red 'determines' that something is not blue, therefore, but at no point does the explanation require that we accept non-blue as a real negative property. Similarly, we could give an account of how being non-coloured 'determines' that something is non-blue in terms of the relation between determinables and determinates without having to allow that non-coloured is a real property.

We can conclude, then, that Zangwill fails to provide an adequate indispensability argument. It is not conclusive that negative properties are required to explain the apparent cases of determination since other, more plausible, explanations of such determination are available, which do not invoke indispensable, irreducibly negative elements.

2.3.3 Affirming the negative

A very different kind of defence of negative properties has come in recent years from David Hommen (2013, 2018). Hommen (2018: 81) argues that negative properties are perfectly conceivable and that there are good reasons to think that they exist. However, this defence of negative properties occurs within a distinctive framework, complete with certain limitations. Nevertheless, the account helps us to understand what is really at stake when we consider the reality of negative properties. Hommen (2018: 82) makes it clear that he is purposefully avoiding an 'onto-epistemological' approach to the problem of negative properties in favour of a conceptual and logical one, since 'the paramount objection to negative properties is that they are conceptually inconsistent' (Hommen 2018: 82).

Now it is not clear who Hommen thinks has made such a charge, since he cites no explicit proponent of the conceptual incoherence of negative properties. Regardless, he does make a case for the conceptual coherence of negative properties

in terms of ‘the class of properties which are denoted by negative predicates’ (2018: 87), where he gives an account of what it is to be a negative predicate. The latter task is quite helpful, since it is far from clear that there is a simple logical or linguistic test of whether a predicate is negative. Is ‘empty’ a negative predicate, for instance? It does not contain the word ‘not’ but still seems to be about what is not: there is nothing in a specific place. Further, saying that someone is unhealthy might mean something ‘positive’: that they have a particular bacterial infection, for instance. Hommen notes that what looks negative and what looks positive could differ across languages, citing the example of the English pair consistent/inconsistent and the German equivalent pair *widerspruchlichfrei/widersprüchlich*, in which the polarities seem reversed. Also, as Quinton (1957: 48, 49) notes, building on a suggestion from Ramsey, there is nothing to stop a language (as long as we ban palindromes) in which a complementary property is indicated simply by the same word spelled backwards, hence ‘neerg’ instead of ‘non-green’.

The lack of a simple criterion of the negative need not compel us to give up entirely on distinguishing positive from negative predication, however, as Zangwill (2011: 531) and Gale (1970: 214) seem to do when they are happy to just assume that we have enough of a pre-theoretic grasp of the positive/negative distinction to proceed. Instead, Hommen argues that what makes a predicate positive or negative is the inferential role it plays (a more sophisticated version of Zangwill’s account). I will return to the positive/negative division, with my own judgement, in §11.3, below.

There are a number of issues worthy of note about Hommen’s general approach, however. First, we can accept the case made that, on the basis of the linguistic evidence, at least of English, we have a willingness to affirm negatives, and this can be understood as an appeal to negative properties. For example,

- (1) This cape is not red,

Hommen says is *not* most plausibly interpreted as saying

- (2) This cape does not have the property of being red.

That is, we do not naturally take (1) as a denial that something has a positive property. Rather, (1) is taken to be asserting:

- (3) This cape has the property of being non-red,

in which a ‘property’ of non-red is affirmed of the cape. Hommen takes this linguistic practice to indicate something significant about ontological commitment: ‘Such claims do seem to commit us to an ontology of negative properties’

(Hommen 2018: 86). I shall take up the issue of denial, and dissent from Hommen's view, in chapter 10.

This brings us, however, to the second noteworthy issue in his view. Does our conceptual scheme and do our linguistic practices determine infallibly, or at least reliably, what there is? Can we read off ontology from our conceptual division of the world? The following looks like an affirmation of such a claim on Hommen's part:

To get an idea of what there might be, and how it might be, it seems helpful to look first at what we say there is, and how we say it is. Such an approach naturally evolves from the idea, idiomatically put, that (the structure of) language mirrors (the structure of) reality (Hommen 2018: 82)

However, this is immediately followed by some necessary reservations:

There is a disclaimer: saying that we may recognize what exists by means of recognizing what is spoken of, and how so, is, of course, not meant as saying that conceptual analysis alone decides what 'really' exists, and how it 'really' is. By no means should we take language for an infallible guide to reality. Quite the opposite: in order to arrive at sustainable ontological conclusions, one has to carefully check how far, so to speak, one can take language at its word. Since the relation between conceptual and ontological facts is not one of strict entailment but rather some contingent kind of mapping, or correspondence, it seems logically possible that the picture of reality we draw through our manners of speaking is occasionally distorted, if not completely deceptive. Admittedly, it would be rather unwarranted to assume that our conceptual representation of the world goes entirely astray. But the hypothesis that conceptualization gets it right all the time seems equally implausible. The very possibility of deviation gives us reason to expect some mismatches between the contents and structure of language and the true contents and structure of the world, if there is such thing.

(Hommen 2018: 83)

These concessions are enough for us to doubt that Hommen provides a robust defence of negative properties. The reasons for questioning the reality of negative properties are unlikely to be that they are conceptually unacceptable or badly defined. Suppose we grant that we can distinguish positive and negative properties conceptually. That doesn't entail that a negative property exists, as Hommen concedes. It is doubtful that the concept of being 'phlogistonated' is incoherent. It might make for a perfectly acceptable predicate whether or not it denotes a real property. The same holds for 'de-phlogistonated'. Whether there is any such property corresponding to the predicate is a matter for further investigation, as Armstrong says when he rejects the argument from meaning. Hommen's defence

does, therefore, raise the issue of what exactly it is that is abhorrent about negative existents of any kind. The problem is unlikely to be that they are logically or conceptually offensive. It is that they are metaphysically offensive because they violate the Parmenidean principle, Pi. Hommen is right that we will have to look elsewhere. Do negative properties play an indispensable role in any causal laws, for instance? Some think so, and we will look further at causation by absence in chapter 4.

Let us grant, then, much of what Hommen says about the conceptual acceptability of negative predications. Yet we are still at liberty to reject negative properties ontologically. Doing so does not mean that we reject the existence of negative predicates. That would be a brave and unnecessary thing to do. Rather, what is at stake is whether there is anything in non-linguistic reality that corresponds to such predications. The soft Parmenidean might then grant that we have a propensity to assert negatives but also have a view that it is ontologically inaccurate or misleading to do so. The assertion of a negative might be perfectly harmless in most contexts. But if we are considering ontology, we are at liberty to say that the linguistic practice does not reflect reality. The reasons Hommen provides will perfectly suffice. From the perspective of any realist framework, we should attempt a conceptual division that fits reality, though without a guarantee that we get it right. Nor should we assume that our forebears who originated languages were expert ontologists, fully aware of the Parmenidean concerns. It is quite possible that they allowed affirmation of negatives because their concerns were not primarily metaphysical. Nor can it be certain that when affirmation of a negative was accepted in ordinary language that it wasn't thought of as a marker of denial, even if it might not look so grammatically.

2.4 Summation and further work

We have now seen various arguments for why there cannot be negative properties but also some arguments for why we should accept them. Our findings might seem disappointingly inconclusive. Arguments against the reality of negative properties are often question begging. We did see that there are some arguments with weight, however, such as that negative properties fail to connect with causal powers. The topic of causation now looks crucial, since an argument in favour of negative properties is that they are indispensable in certain causal claims and explanations. Yet Armstrong has already denied that there is causation by absence, saying it would be absurd to allow that lack of poison makes you healthy. What this shows us is that any final judgement on negative properties will have to wait until we have considered some further issues. In chapter 1, above, we granted that some negative entities might have to be admitted if they are indispensable in some of our best theories. Whether they are necessary for a full account of the

causal laws of our world is thus crucial. As well as indispensability, as a reason for admitting a negative entity, we said that we also had to ensure that the purported entity really was irreducibly negative. Apart from considering causation by absence, then, we should also think whether the negativity of a negative property can be explained away, for instance, in terms of denial rather than affirmative assertion. I will return to that question in chapter 10.

3

Nonentities

3.1 Negative particulars

Having considered negative properties, we ought now to consider putative negative entities of the other major ontological category: particulars. This is not to commit explicitly to a two-category ontology. There might be more than just properties and particulars. But as this is not primarily a work on ontological categories, I will not labour over, for instance, whether events, modes, or tropes are a distinct ontological category. We will find that there are plenty of negative existents to keep us busy, and it is their negativity that constitutes the present set of problems rather than questions about relative relations among categories.

The simplest way to refer to negative particulars is nonentities, though we will find that the term doesn't always fit. Among nonentities we can think of non-existences, fictions, negative facts, absences (many), nothing (only one), emptiness, holes, shadows, silhouettes, darkness, silence, mirages, vacuums and voids, privations, lacks, edges, limits, boundaries, stops and pauses, endings, gaps, negations, denials, nihilism, scepticism, atheism, ignorance, forgettings, failures, omissions, evasions, avoidances, exclusions, prohibitions, the empty set, the empty world, and the number zero. Some of these might permit specific instances. For example, among stops there could be a stop in the traffic or in a piece of music and among holes there could be hollows, tunnels, and cavities. The diversity of the list shows why we need a term as broad as nonentity to collect them.

Our job in this chapter is to consider the problems created by the existence of any of these nonentities and, in accordance with our soft Parmenidean project, whether they can be explained (away) in terms solely of what there is, thereby avoiding the reification of what there isn't. Such is the extent of the list, it will also be useful if we can find a typology of nonentities, and I will offer what I take to be the most pertinent classification before I consider some of the major cases in more detail. I divide the nonentities in the list above into non-beings, limits, privations, omissions, negative epistemic states, normative negatives, and logical/mathematical constructs. The soft Parmenidean project is to deproblematize them all.

3.2 Non-beings

We have heard of a number of non-beings: the golden mountain, Pegasus, the square circle, *Oliver Twist*, Father Christmas, and the highest prime number. To the best of our knowledge, none of these exist. The reasons are various. Some are legendary, mythical, or fictional. Others are impossibilities. Nevertheless, and in defiance of Parmenidean considerations, some are tempted to say that these non-beings have at least some existence. Each is an object of thought, one might say; or an intentional object. Does this give them at least some kind of being? This question was a motivation behind Meinong's (1904) most infamous though also misunderstood thesis. If we take our knowledge of Meinong from Russell (1905), we get the impression that he allowed everything some kind of being (although this was, rather, Russell's 1903 view). Objects that were not real existents, the popular conception goes, nevertheless subsist. They must do so in order to be an object of thought. Items such as the golden mountain thus have some form of being, though subsistence is inferior to existence.

What is a Parmenidean to make of this? Under the popular conception, it suggests a domain of subsistence that is halfway between existence and non-existence. Subsisting objects are almost existents but not existents in the full sense in which the Eiffel Tower or the cup on the table exists. Parmenideans must reject Meinongianism under this interpretation for it is in violation of Pii. There are no degrees of being, and subsistence sounds like existence in an inferior degree. Everything either exists or does not exist, so there is no middle ground for any 'thing' to occupy. Furthermore, there are no nonentities, since things either are or they are not. Nonentities are not entities, therefore: they are nothing.

More recent scholarship (Grossmann 1974, Parsons 1980, Routley 1980, Perszyk 1993, Jacqueline 2015), however, has shown that the issues Meinong addressed and the solution he offered have at least some weight. Subsistence is not really intended as a lesser form of existence. What concerned him was what our thoughts are about when they concern non-existents. And if they are not about anything at all, how are they coherent? How can we be saying or thinking anything when we think about centaurs? This difficulty might be the reason Parmenides warned against naming or thinking about nothing.

If subsistence is a term for an object of thought, then, Meinong's view is not entirely incredible. The golden mountain exists as an object of thought—as an intentional object—it just doesn't exist as a concrete being. If this is how it is, then there would be relatively little trouble for our soft Parmenideanism. If we assume that the claim of existence for an intentional object is clear, and not confused with the claim that the object thought of exists concretely, then no degree or intermediate form of existence is invoked. Thoughts exist and objects of thought exist,

but the objects thought of, in these cases, do not exist. Meinong included other non-concrete beings within subsistence, such as numbers.

Unfortunately, this cannot be the whole story, however. It is implausible as a theory of reference, for a start. When I wonder whether Father Christmas exists, I do not wonder whether my thought of him exists, since that would be self-confirming. To think that Father Christmas does not exist would be self-contradictory. Thoughts can be about things other than our own ideas. If you think that Barack Obama exists, your thought is about Barack Obama. By parity of reasoning, when you think that Father Christmas exists, your thought should be about Father Christmas. We might have a metaphysically satisfying account, if we flatly deny the reality of non-beings on Parmenidean grounds, but we will not have an adequate theory of empty reference. We will need to consider the topic of empty reference in detail below (chapter 8). Leaving that problem aside until later, however, we should still see why treating subsistence as a lesser degree of existence is not an attractive option and not a satisfactory solution to the empty reference problem either. There is a philosophical lesson here: we shouldn't be forced to accept a bad theory merely because we can't offer a better one. That better theory might come in time.

If there is no simple argument from empty reference to the reality of nonentities, are there any other reasons we should accept them? The only other reasons I find are case-by-case pleadings. Holes exist, for instance; and shouldn't we say that holes are a special kind of negative entity? Similarly, there are shadows, which seem to essentially involve an absence of light. Are they negative entities? There is an empty set, an empty world, and the number zero. Aren't these negative entities? Accordingly, we had better look at these candidate negative entities in turn.

3.2.1 Holes

When we think of negative entities, the subject of holes often arises. A hole seems like a type of negative existent since, it might be said: 'A hole is there where something isn't' (the epigraph to Casati and Varzi (1994), from Kurt Tucholsky). We will see that this, nevertheless, does not make holes themselves negative entities and, thus, no immediate threat to Parmenideanism. Following Casati and Varzi (1994) and Kachi (2011), I take holes to be immaterial entities that are parasitic on their 'hosts': the matter that surrounds the hole. The hole will cease to exist if the host ceases to exist. Being immaterial does not mean that the entity is a non-entity, however. It is a part of our reality. We can conceive of something that would qualify as a 'negative' hole, but there is no reason to think that this will be a real entity of any kind, especially once we see how attempts to identify and individuate negative holes break down.

Casati and Varzi (1994) distinguish three kinds of hole: hollows, tunnels, and cavities. Each has different identification conditions. For example, the number of cavities is fixed by the number of surfaces of a host object; that is, where n is the number of surfaces of an object, the number of cavities is $n - 1$ (1994: 53). We have a tunnel when we could loop a string through an object and can count tunnels by how many such loops we can make, where no loop can define more than one tunnel (1994: 47). These conditions tell us something of what it is to be a cavity or a tunnel, but there are more strict formulations (1994: 63–66). The perfect filler of a tunnel will have two free (showing) faces (think of these as visible continuous patches of sand, if the hole is filled with sand). A hollow's perfect filler will have one free face. A cavity's perfect filler will have no free face. Filling a hole does not mean there is no hole: just that there is a filled hole, which could later be emptied.

Holes are dependent on their hosts, but they are non-rigidly dependent, since they can maintain their identity through a change of their hosts; for example, an air bubble (an air-filled cavity) rising through water and into a layer of oil. We see too, with this case, that holes are not essentially at their spatiotemporal locations since it is possible for holes to move. Hence, you can have different holes at the same region of space at different times: think of one bubble moving away from a location and a different bubble drifting into that exact spot.

A key question is whether a hole is a kind of absence and thus the sort of thing that Parmenideans cannot admit. This would be a difficulty, since it can hardly be denied that there are holes. Contrary to Martin (1996) and Kachi (2011), however, I think it is a mistake to say a hole is an absence. A hole might involve an absence, but it is wrong to say that it *is* an absence, and thus I think it wrong to classify holes as negative existents.

Casati and Varzi, Kachi, and others make a case for realism about holes, which are treated as immaterial objects. Because holes are parasitic, however, something material must exist for there to be a hole. There cannot be a hole, for example, in the middle of a region of empty space. Nor, as Coggins (2010: 71) says, could you have a world consisting only of holes. The host of the hole will have a surface that acts as the boundary of the hole, giving it a determinate extent and shape. There is an absence of matter in some cases of holes, namely unfilled ones, but it is discontinuity, boundary, and surface that make the hole, not whether or not the hole is empty. For discontinuity, think of the concave indent of a hollow. For surface, think how a cavity within a rock gives the rock an inner surface as well as an outer surface. Now, why insist that the hole must be a hole in matter? Could there be holes in immaterial things? Would it count as a hole in empty space if there were a material sphere in the middle of it (a 'reverse hole', perhaps)? I think not. It matters that the hole is a gap: something that can be passed through, entered, or filled. Our concern is with holes in things not things in holes.

If you are concerned primarily with the metaphysics of negatives, it is not holes that should be your worry (edges and boundaries might be, but we will come to them shortly). The metaphysical problem of negative existents would be concerned not with the real, actual, 'positive' hole in the doughnut, for example, but with something like, if it was, the non-existent, merely possible hole in the middle of a continuous solid wooden table: a hole that is *not*. This is what a negative hole would be. We have good reason not to admit such nonentities into our ontology. There are differences concerning identity and individuation between real holes and negative, non-existent holes. These differences in identity and individuation conditions could (A) show us a way to distinguish genuine holes from 'negative' ones, (B) add some weight to the case against reification of negative holes, and thereby (C) contribute to the case against negative entities generally. Accordingly, I will mention a number of ways in which holes differ from negative holes.

Casati and Varzi (1994: 129) allow that different holes can occupy the same space at different times. Negative holes are even less constrained. There could be different holes that occupy at least some of the same space at the same time, if negative holes were to be real: they could overlap spatially, which real holes cannot without merging into a single hole. It is not clear that there are even constraints against different holes occupying exactly the same space at the same time, especially given that holes can move. For example, one negative hole, *a*, moves through a solid from left to right between t_1 and t_3 , and another non-existent hole, *b*, moves right to left between t_1 and t_3 where, at t_2 , *a* and *b* would be at exactly the same place. With real, 'positive' holes, there seems at least some additional constraint, since if the above case really occurred then it would be a case of two holes fusing into one and then fissuring.

Holes are not made out of ordinary matter, but they have matter around them: the matter of their hosts. With negative holes, there could be surrounding matter in some cases, namely in un-holed objects. But some non-holes, namely absent holes, could be made out of nothing, since there would be no reason why we could not say there were non-holes in a volume of empty space. It is unclear that non-holes would need hosts, so uncertain would be our grip on their identification conditions. This shows that there is a difference between a possible hole and an absence of a hole, even if this difference is merely counterfactual. One cannot say of an empty region of space that a hole could be there, since there is no host. But one can still say that there is an absence of a hole there. So an absent hole is not merely a possible hole. (There could be the iterated possibility of a hole in an empty space because there could be an object in that space and then there could be a hole in that object.)

Next, we could also have a Ship of Theseus-type case with a real hole, where the matter around a filled hole gets gradually replaced. But then consider how you

could ‘destroy’ a hole by removing all the surrounding matter (taking away all the ship’s planks). Holes need surfaces to form their bounds (Casati and Varzi 1994: 137) and surfaces are actual. The surfaces of absent holes are not actual at all.

Finally, holes can have causal effects, whereas negative holes presumably cannot. You can trip over a hole, and water can flood through a hole. A hole is not merely an absence, though, and these are not simply cases of causation by absence. This is admittedly a more complex issue than it first seems, however. A knock on the non-existent hole, that is really a solid tabletop, will produce a sound. But it is the tabletop making the sound, not the absent hole in it. That is not a case of causation by absence either (see chapter 4 for a treatment of this topic).

We have seen enough to conclude that holes are respectable enough. What is startling about them is that they seem to be immaterial objects. But the key claim is that immaterial does not mean non-existent. Holes have identity and individuation conditions. Those who find nothingness repellent should be against non-existent holes. As well as an immediate aversion to such a notion, however, we see that identity and individuation conditions for such putative existents break down. We hardly have a grip on them at all: are there two negative holes in the table or two thousand? Real holes can be counted; negative ones cannot. This gives us a more considered reason to reject the reality of any such negative entities.

3.2.2 Negative facts

Another significant kind of negative entity that has been proposed is negative facts. Russell (1918: 206–8) thought that we had to accept negative facts to account for general truths. To say that all people in the room are philosophers is to say something like *a* is a philosopher, *b* is a philosopher, *c* is a philosopher, and there are *no more* people in the room than *a*, *b*, and *c*. It was clear too, though, that Russell accepted this conclusion regrettably, and admitted that when he first proposed the idea at Harvard ‘it nearly produced a riot’ (Russell 1918: 187). He thought he had the best theory of quantification, but it included an irreducibly negative component. Russell must have been a soft methodological Parmenidean. So too did Armstrong (1997, 2004) accept totality facts, which *total* a set of first-order facts by permitting no more facts.

Facts are taken—by Armstrong, for instance—as a kind of particular. We have a fact when we have a particular bearing a property, such as when a ball bears the property of redness. If you bring together a particular and a universal, you get another particular: there is a victory of particularity (Armstrong 1978a: 101). Assuming that such an account is sound, this seems to mean that if there are

negative facts, then there are nonentities. Presumably, if there are any negative facts at all, then there will be many of them. The stakes are high.

Recently, negative facts been revived. Barker and Jago (2012) say that one reason there is enmity towards them is that it is not clear what negative facts are (Barker and Jago 2012: 117) and they hope to correct this. They also say that you 'choose your poison and then see what work you can do with it' (Barker and Jago 2012: 120), which I see as a tacit admission that negative facts are undesirable. This is confirmed when they say what a negative fact is. It is not a particular instantiating a negative property, though they do eventually accept negative properties as abstractions from negative facts (Barker and Jago 2012: 123). Instead, they suggest the particular and property are brought into a fact through a kind of nonmereological composition or tie that they call 'anti-instantiation'. For example, the fact that there is not a hippo in the lake consists of the lake anti-instantiating having a hippo within it.

Positive and negative facts do not differ in their constituents, on this view. The lake being frozen and the lake being not frozen have exactly the same constituents. And the instantiation and anti-instantiation that differentiates them is not a constituent: it is the nonmereological way in which the constituents are tied together to become facts.

What is obviously alarming about this account is that it explains one kind of negative with another. Anti-instantiation, whatever it is, certainly seems like some kind of negative. Barker and Jago (2012: 127) admit that anti-instantiation is opaque but claim progress nevertheless: progress in linking negative facts to nonmereological composition. But until the difference between instantiation and anti-instantiation is explained, then we have nothing but a word: a word that is supposed to solve the long-standing problem of explaining negative facts. There is then a claim (Barker and Jago 2012: 128) that, despite having said next to nothing about anti-instantiation, its adoption is justified by the work it can do. But how can we be sure that it really can do the work required of it if we don't know what it is? Anti-instantiation looks like a *deus ex machina*. Some of the work that anti-instantiation is supposed to do is to explain causation by absence, entities such as holes, truthmakers of negative truths and perception of absence, since negative facts are supposed to be perceivable. We have already encountered holes, though, so let us see if we can explain some of those other phenomena too (see chapters 4, 7, and 9) without invoking a negative *deus ex machina*.

Suppose we are unconvinced, then, and remain sceptical about negative facts. We are of course at liberty to do so, but then we should not ignore the problems that led philosophers like Russell and Armstrong to invoke negative facts as solutions. In Armstrong's case, that was the problem of negative truth, and we will also need to say something about universal quantification.

3.3 Limits, boundaries, edges, and stops

It might seem fanciful to suggest that there are negative particulars of some kind. But there is a consideration that even regular entities drag non-being with them into reality. The argument is that every thing has a boundary, edge, or limit. But a boundary is formed only at the point where what the thing is meets what the thing is not. To be a boundary requires that there be some place or surface beyond which the thing is not. And we have a something only if there is some such boundary. To be, therefore, requires that there also be non-being, hence we have no choice but to reify non-being. Stops, endings, and pauses would be like the boundaries of events or processes rather than of objects.

This argument has some notoriety. We find in Plato's *Sophist* (257c5–d13), for example, the idea that there must be negative kinds since there can only be a positive kind if there is something that is not of that kind (see also Crivelli 2012: 205ff). Consider the kind horses, for instance. To be constituted as a kind, there must be something that is not horses. The negative kind non-horses creates a boundary around horses, and only in having such a boundary is horses a coherent, demarcated kind. Think of a child learning the concept of a table. It is one thing to be able to successfully identify tables, but, the argument goes, you have not fully grasped the concept of a table until you can also identify things that are not tables. If you call a chair, or any four-legged thing, a table, then you have still not grasped the meaning of table. Plato's presentation of the argument gives this point more ontological seriousness since it is not just about identifying kinds. Plato is saying that without the corresponding negative kind, the 'positive' kind will not have any being at all since it will have no boundary that distinguishes it from everything else. Something is only constituted as a kind if its non-being is real. Clearly, a similar argument could be mounted for properties generally. Perhaps something is red only if non-red exists in contrast.

Spinoza is attributed with the most famous version of this argument. The attribution comes from Hegel (1812: 113) as 'omnis determinatio est negatio': all determination is negation. As Stern (2016) shows, however, Spinoza never put it exactly this way. What Spinoza did write was:

he who says that he apprehends a figure, thereby means to indicate simply this, that he apprehends a determinate thing and the manner of its determination. This determination therefore does not pertain to the thing in regard to its being; on the contrary, it is its non-being. So since figure is nothing but determination, and determination is negation, figure can be nothing other than negation, as has been said. (Spinoza 1674: 240, translation from Melamed 2012)

If we take a square figure, for example, its squareness is only determined by it having the right boundaries, and these boundaries are made by there being certain places beyond which the square is not. The boundary is the point at which the square's being meets its non-being. Only with its non-being does the square have any being. These old arguments are not gone. Armstrong used them relatively recently, such as in the following:

It is true, of course, that totality facts themselves partake of the nature of negation. They are 'no more' facts. They set a boundary, beyond which nothing, or nothing relevant, exists. All determination is negation, as Spinoza said.

(Armstrong 1997: 200)

How seriously should we take this argument? It seems undeniable that there are edges, boundaries, and limits to things. It might also have to be admitted that they would not be constituted as things if they did not have some such boundaries. Does that then mean that we have to allow non-being as real in order to have any being at all? Or has Armstrong conceded too much by saying that boundaries must be negative?

The Parmenidean might yet have a way of escaping the worst implications of this conclusion, although Parmenides himself seems to have accepted that the argument had force. What is needed is a non-existential use of the quantifier phrase 'there is...' since we want to be able to grant some cases where there is nothing but without reifying this nothing as something that is. As mentioned in chapter 1, above, when I say, 'there is nothing in my pocket', clearly I am not saying that there is something in my pocket; namely, nothing. I use the phrase without existential commitment to the subject. I do not accept, therefore, that simply to be is to be the value of a variable (Quine 1948) since I am willing also to talk about what is not. An explanation of empty reference will then be required, and we will consider the detail of that later (chapter 8).

The soft Parmenidean need not deny non-being. In saying that there is non-being, however, they are not using an existential 'is', which is to say that they are not reifying nothingness. Similarly, we can allow that there is an absence of water in some soil or an absence of books in the house. But this is not committing to absences. You cannot say how many absences of books there are in the house—and there is no entity without (numeric) identity (Quine 1969: 23). To say that X is absent is to say that there is no X. To say that there is nothing beyond a boundary is similar. It is not to say that there is something, namely nothing; so, despite the surface grammar, it would be more accurately rendered as $\neg\exists x(Fx)$.

We saw how Parmenides himself did not allow us to talk of or name nothingness (Pvi) and we can now see how this might be his reason for stipulating that reality is a single plenum. It is hard to say this conclusively, with only brief evidence, but his concern could have been that if he allowed that there is empty

space, or some part of reality that is unfilled, then he has to allow that something is not: there is nothing in the gap, there is non-being. If one then accepts these quantifier phrases as having existential import, then one has committed to them as part of what is. Parmenides' 'solution', then, could have been to deny that there is any empty space. Everything is filled with being. There is no non-being. The soft Parmenidean, however, reading statements of non-being as non-existentially committing, can admit that some things are not but without reifying them. They remain a part of non-being rather than being.

This allows us to resist the worst implications of Spinoza's principle. A table will have a surface: a boundary beyond which the table is not. The kind horses, too, will have a boundary beyond which there are non-horses. We can say this without reifying non-tables, non-horses, or any other negative entities. It might be that such utterances are equivalent to denials rather than assertions, though we will defer that for later too (chapter 10).

It is most likely that Parmenides overreacted to the problem. He was not alone, however, since we can say the same of Armstrong on limits. He admits totality facts into his ontology but somewhat reluctantly since he regarded them to be irreducibly negative (Armstrong 1997: 240). However, in one of his very last comments on the topic, there is something of a revelation:

I accepted and still accept that there has to be such a [totality] state of affairs. But in the past it still seemed to me that this was an addition to the ordinary states of affairs, so I had to talk fast to try to prevent an infinite regress arising. But a cutting-off of all states of affairs is no addition. 'No more' is not something more! The cost is, a cost I suggest must just be paid, that negation in the shape of 'no more' must be admitted into our ontology. Limit is real. It is an ontological feature. (Armstrong 2010: 80)

This considered verdict of Armstrong has some merit. The fact that there are no more states of affairs than all the first-order states of affairs provides the limit on those states of affairs but is no addition of being. It is not something to be added to the world's inventory. It is a fact that there are no more things and so it cannot be a thing itself or its addition would entail that the things it limits are not the limit. This is a way of denying Parmenides' position, then. We can say that there is a limit beyond which there is nothing but still without reifying nothing. The price, as Armstrong sees it, is that we have to admit these negative, 'no more', facts, and we have already seen how negative facts are problematic. Again, however, we can also see that if we had accounts of empty reference and denial, it might be possible to say what we want to say without invoking negative facts as worldly existents. Instead of asserting that there are no more facts, and worrying about the truthmaker for the assertion, perhaps one can deny that there are more facts, where the denial is appropriate precisely when there is nothing at all more.

Similarly for universal quantification, to assert that something is all is to deny that there is more than it.

3.4 Privations

A privation is understood as a deficiency, where something could, should, or ought to have a certain power or capacity, which it does not (Aristotle, *Categories* 12a28–33). This differs from simply not having a power or capacity. For example, we can say of an idea, or of the number 6, that it is not soluble, since it does not have a capacity to dissolve in liquid. But it would not be right to say that these things are insoluble since they are not the sorts of things to which solubility could apply. It would not simply be false but also a category mistake to think of them as soluble, for which reason calling them insoluble is also inappropriate. However, it does make sense to say of a rock that it is insoluble since rocks belong to a category of things—solid physical objects or substances—to which solubility could sensibly be applied. Indeed, there are many solid physical objects or substances that are soluble, such as sugar cubes. We might judge, therefore, that it is appropriate to say insoluble, in the case of rocks, because it indicates what we might call a meta-level privation or meta-privation. It is not that the rock is lacking something *qua* rock, when it is insoluble, since rocks in general have no tendency to dissolve. But at a higher level of abstraction, a rock is a kind of thing—a solid physical object—where some things of that higher kind are soluble.

The distinction between internal and external negation can be used to separate privations from non-privations, where it is desirable to do so. Hence, we might say, using external negation:

- (1) It is not the case that a rock can see

But we would use internal negation when we say:

- (2) The man is unseeing.

A person typically can see but some have a privation of this capacity, which the internal negation indicates. The rock has no (first-order) privation of this ability since rocks do not ever have a capacity to see.

The notion of a privation is not strictly defined but nor is it useless. Indeed, with the above account we can find the answer to some putative counterexamples. Horn (2001: 7) objects that on Aristotle's account, we cannot say that a human baby is toothless since babies are not supposed to have teeth. Yet toothless the baby surely is. '-Less', as a suffix, seems ambiguous between internal and external negation, though. On the above account, it is permissible to externally negate, in

this case: the baby is toothless since it's not the case that she has teeth. Whether or not it is right to internally negate here (the baby is 'untoothed?'), this is enough for a toothless baby.

What should we say of privations? Are they things? Shortly, we will look at a specific case of something that seems to be real but also a privation: a shadow. What makes shadows seem thing-like is that you can see their distinct boundaries. But can you see things like toothlessness or insolubility? Are they any kind of entity? Pre-empting a topic that will also figure later (chapter 10), note that in both (1) and (2) we deny that something can see, even if they fail to do so in different ways. In pointing to a privation, then, we are saying that something is not: that something does not have a certain power. But not having a 'positive' power does not entail having some kind of negative power. This response of treating 'not' as a denial, does not explain the difference between (1) and (2). We are not, however, yet offering a reductive analysis of 'not'. We are merely trying to show that attributions of privations are not attributions of negative causal powers.

Let us move, then, to the more challenging case of shadows.

3.4.1 Shadows

The simplest thing to say of a shadow is that it is an absence of light, which makes it look exactly like a negative entity. There had been relatively little philosophical scrutiny of shadows until Casati's (2004) book-length study. Shadows were more the concern of astronomers and painters and Leonardo began by describing a shadow as an absence of light when he felt compelled to consider the phenomenon (see Casati 2004: 166). Sorensen (2008: 17) maintains that shadows can be seen, which I will not here deny. Our concern now is whether they are real negative existents.

Some reflection reveals that a shadow cannot be only an absence of light. Consider a sealed, windowless room in which there is no light source. There is no light in the room at all, but there are no shadows either. A shadow requires a source of light that is partially blocked: so there is at least some positive element to a shadow, too. The following case supports such a view. Suppose we turn on a flashlight but then seal it inside a box from which no light escapes. This light of the flashlight is wholly blocked and, intuitively, it casts no shadow outside of the box. The point is that its light is *entirely* blocked so it casts no outlined shape upon anything. Assuming that there is nothing else in the box, then there is no shadow in there either. One side of the box's interior is illuminated and one side is not, we expect, but it seems that the dark side of the box is not in shadow. The bulb does not cast a shadow down its own handle since there is a difference between casting a shadow when a light source is partially blocked and when no light is being shone in a direction in the first place. What we need for a shadow is a light being

shone in some direction which then gets partially blocked such that there is an area without illumination that is bordered in at least some place by an illuminated area. A shadow seems to be something but with an essentially negative component nonetheless. At least the area within the shadow is marked off from that around it by an absence of light.

The term shadow is ambiguous between at least three things: the original shadow, the shadow cone, and the cast shadow (Casati 2004: 169). For a cast shadow, the light must fall on a surface, bordering an illuminated area when the light is blocked. Suppose that the universe is of infinite extent and there is an outermost planet that is illuminated by a star. For at least some of the time of its orbit, it has a shadow cone that extends outwards into empty space. On the planet's dark side is what we can call the original shadow. But during this time, it casts no shadow since the partially blocked light falls on no surface. The shadow cone will end at a point. In some circumstances, where a light source is smaller than its partial blocker, an inverted shadow cone could in theory extend outwards infinitely, never casting on anything.

Thinking of the dark side of a planet as a shadow, what we call the original shadow, raises the question of whether all darkness is shadow (see Sorensen 2008: 263). Night-time on Earth is when the sun's light is partially blocked by the body of the Earth itself, though it still illuminates the moon. But it is wrong to conflate shadow and darkness. God said first of all 'let there be light'. Had he not done so, the universe would be one of ubiquitous darkness: but no shadows, I maintain, since there would have been no light to partially block. This is why we have the conceptual apparatus to distinguish darkness and shadow in more mundane cases. The inside of a box is dark as long as there is no light source within it, but we are reluctant, I think rightly, to say the box contains shadow. Even though there is a light source outside the box, which the walls of the box block, the light is entirely, rather than partially blocked. Open the box a bit, however, and you could let a shadow in.

Shadows might be real but there remains a question of whether and in what way they are negative existents. We saw in the case of holes that there were good reasons to take them as real—and thus positive—entities even though they are immaterial. Similarly, shadows have some features that suggest that they are not made of matter. As Casati (2004: 28) explains, they do not seem to have any thickness at all, hence a case can be made for seeing them as the only two-dimensional concrete existents (whereas 2-D shapes are abstract). Nor can we reduce shadows to the material blockers that create them since they are at different places and the shadow typically has a different set of properties from the blocker. For example, a disc suspended on a string below a light source can rotate, but it is implausible to claim that the shadow rotates. The shadow can move, if the disc moves from side to side on the string, but it seems incapable of rotation. Sorensen (2008: 86) thinks otherwise, saying that the shadow is also rotating but indiscernibly so. This seems hard to maintain, however. It is not simply that something about the

shadow is changing but we cannot see it. Rather, nothing about the shadow is changing, since information about the object—that it spins—is not preserved in the shadow it casts.

There is also a puzzle documented by Todes and Daniels (1975) in which shadows that are overdetermined seem to have no obvious cause. Suppose we have a large light source, such as the sun, and towards the end of the day a stone statue, call it S^A , standing between the sun and the wall, casts a shadow on a wall. Because the light source is large, the shadow will be slightly smaller than the statue (if it were a small light source, the shadow would be bigger than the statue). There seems little doubt that the statue creates the shadow; we might call it the shadow's cause (arguably a case of causation of absence—see chapter 4). Now suppose we have a slightly smaller version of the statue S^A , which we call S^B , and position it between S^A and the wall. It is possible to position S^B such that its shadow is exactly the same shape and size as that which S^A was casting. We would then be in a position where we could remove either S^A or S^B (but not both) and the cast shadow on the wall will be unchanged. In this rare but possible scenario, one can argue that neither S^A nor S^B causes the shadow. First, S^B cannot be causing the shadow because it is not blocking out any light. It stands entirely in the shadow of S^A so no light falls on it at all, and something can cause a shadow only by blocking light. But S^A is not causing the shadow either, since the shadow that it creates falls entirely and only on S^B . Neither S^A nor S^B creates the shadow on the wall, therefore.

Now if shadows were material existents, they would need causes. Beings are created, unless they arise spontaneously (and it is implausible that shadows arise spontaneously). If the shady part of shadow were just an absence of light, then it seems we might actually have found a nothing (cast shadow) that really does come from nothing (absence of light), though this is undoubtedly not what Parmenides intended when he made that famous claim. Instead, it might be best to say that the statue, in blocking light, explains the presence of a shadow without causing it. The problem is not that an effect cannot be a non-occurrence, non-change, or absence of something, since such cases seem permissible. The problem seems to be the immateriality of shadows. They are real because they can persist and move but, by being immaterial, there are some things we cannot say about their persistence that we can otherwise say about material things. For example, consider a room in which a chair stands towards the edge, casting a shadow against the wall (again from Casati). You can turn the light off and be no longer able to see the chair but then turn it back on again and be confident that the chair you saw before the light went off is the same chair as the one you see when the light returns. But now consider the chair's shadow. Is the shadow after the light has come back numerically the same one as the shadow before the light went off? Here, we seem less certain. At least, it seems indeterminate whether or not it is the same shadow. It appears the same, but that might just be a type identity rather than numeric identity. I suggest that this indeterminacy regarding numerical

identity arises because shadows are immaterial. In the case of a chair, its numerical identity seems grounded in it being composed of the same matter before as it is now. No such grounding exists for the shadow since it is not composed of anything.

3.5 Omissions

We need to explain how omissions are significant without being ontologically serious. A straight denial that there are omissions would not get us far since it is so at odds with a common way of understanding. It should also be conceded, however, that there seems to be less pressure to admit omissions as things. The task, then, is to reconcile these two intuitions, which seem to pull us in opposite directions.

There are two types of omission that it is worthwhile separating. The first is in human action or, as should be said, inaction. The second is when it is more generally the case that something doesn't happen, outside of the sphere of action. Omissions are included within this chapter because these can be understood as negative events, and while there can be event types, the omissions with which I am concerned are event tokens.

Event tokens would be located individuals, even though those locations could be extended regions. A tyre puncturing on a car can be dated to a time, even if the process of puncturing takes a little while: say, a tenth of a second. If the car is moving, then the spatial location of the puncturing might also be extended over the region through which the tyre moved in that time. Other events have much bigger spatiotemporal extensions. If World War II can be considered an event—and I don't see why it shouldn't, even if it was a complex one that included many other events as parts—then it had a temporal extension of around six years and occurred in a spatial region that included much, if not all, of the world.

Negative events, if there are any such, might also have spatiotemporal locations but it is not clear that all of them do. Suppose I complete a car journey from Durham to Sheffield and arrive without incident and no puncture. Where did the non-puncture occur? It would not have been at quite so small a region as a real (positive) puncture would have occurred, if ever it did. Did the non-puncture occur at every part of my two-hour journey, then? Did it occur at every metre and every second? Or were there several non-punctures, each of which was precisely located? If there are non-puncturing events, surely there was a near infinite number of them occurring (or non-occurring) along the way. Negative events would be curious existents in this way and other ways too. If they were real, then it would not just be non-puncturings that occurred on my journey. I was also not car-jacked, not hit by an elk, nor a nuclear bomb, I didn't write a novel while travelling, no one spoke to me about Simone de Beauvoir, and so on. It is clear that, if there are non-occurrences, then there would be far more of them than occurrences: the world would seem overpopulated with non-events.

There could be attempts to reduce the number of non-occurrences with some sort of relevance condition, but it is hard to see that this could be done in an ontologically serious way. A puncture is the sort of thing that might have happened on a car journey, so it is a relief when it doesn't. In contrast, being hit by a nuclear bomb is such a rare occurrence anyway that it would seem bizarre to be relieved that it hasn't happened. This misses the point of concern, however. It might be perfectly rational to be relieved when unfortunate occurrences that were likely to happen didn't happen, and irrational to feel such relief over unlikely occurrences not occurring, but our question is whether such non-occurrences are real. The matter is an ontological one and it is unsatisfactory, therefore, to bring in pragmatic factors as a basis for deciding what exists. If we have a situation where, if one non-occurrence exists, then a near infinity of them must, it seems a good motivation for saying that none of them are real.

Omissions in actions might seem to provide at least some motivation for accepting non-occurrences as real, however. After all, we hold people responsible for their failures. Sometimes you have nothing to do. But suppose you decline to help a drowning man, not even throwing him a life buoy that is easily to hand, or you fail to raise an alarm when you see a fire and quietly make your own departure instead. Is not your omission responsible for any deaths that ensue? It seems that there are cases like these where you indeed would have responsibility. Again, however, responsibility does not require that your omission is something that exists, or that it is a cause of the deaths. You might be responsible for harm that you cause but being a cause of something is not the only way to be responsible for it (Moore 2009: 304). A sound basis for responsibility is also that you failed to act where, if you had acted, some unwanted event would not have occurred. The man drowning is your fault because you could easily have saved him and were negligent in not doing so. It is not that you have caused something to happen, then. It is that you have failed to cause something else: the man's rescue. As someone once said, with great power comes great responsibility (Mumford and Anjum 2013). Given that omissions need not be causes, and arguably cannot be causes, then, respecting the Eleatic principle, they need not be reified in our ontology. An acceptable theory of responsibility does not need them. The alarming consequences that would follow from any such reification need not be swallowed, then. The soft Parmenidean can accept events as real without accepting non-events.

3.6 Negatives norms

The subject of omissions in action has already seen us stray onto some normative matters. One can be responsible for the failure to do something *if* it is something that you should have done. Now it is one thing to fail to do something that you should. A prohibition, however, is a norm that something ought not to be done.

Bans and exclusions will also fall under this head and possibly exceptions too. A positive norm is one such as that you must submit your tax return by 31st January. A prohibition can be considered a negative norm, such as that you must not operate heavy machinery while under the influence of alcohol, take goods from a store without paying, or swim slowly in the fast lane. Perhaps there are some subtle distinctions to be drawn between *should not*, *must not*, and *ought not* (Anjum, Lie and Mumford 2012) but we need not consider them here since our concern is solely with their negative aspects and whether they would present a challenge to a soft Parmenideanism.

What is characteristic of a negative norm is that one does not need to do anything in order to satisfy it whereas to violate the norm one would have to do something. This contrasts with positive norms that require something to be done, often conditional upon some other circumstances, while doing nothing in those circumstances would violate the norm. If the norm is to make your tax return by 31st January, then an action satisfies it and inaction violates it. With a negative norm, such as that you must not drink and drive, you satisfy the norm just by doing nothing, such as if you are at home asleep, but break the norm by performing the action that is prohibited. There will be some cases where the matter is slightly more complicated. For example, the prohibition against swimming slowly in the fast lane of the pool is satisfied by doing nothing at all, but it is also satisfied by swimming fast in the fast lane. Only an action breaks the prohibition, though the action could also be considered as a failure to swim fast. It is clear that a failure to swim fast is not quite an omission in the earlier sense, however. It is not nothing at all. One can swim slowly only by swimming, hence one is doing something even if it is not doing something in the way that one should.

In the previous section, we saw that omissions do not pose any great difficulty for the soft Parmenidean. Only actions need to be accepted into one's ontology as real and while omissions might have explanatory value in regards responsibility, this is only because someone could have done something that they should have. The omission itself need not be admitted as real, as the negative correspondent to an action. Once that is accepted, then prohibitions also pose no serious ontological challenge to Parmenideanism since, as we have seen, they can be satisfied by an omission. If the above analysis is correct, then we can avail ourselves of the same response when it comes to violations of positive norms, since those violations will also consist in omissions.

3.7 Negative epistemic states

There are a number of negative epistemic states. The case of believing something to not be the case is the main one we should consider since we will see that it is crucial to understanding negative predications and so-called negative truths. A

connection will be drawn with denial, which one could say is another negative epistemic state: declining to affirm or withholding of belief. Such will be the importance of negative belief and denial, however, and how they relate to negations, that they will receive separate treatment in chapters of their own (chapters 10 and 11). There are a number of other negative epistemic states that we should note here, however, even if they do not require as detailed an examination.

Forgetting is a case that we can explain briefly since we have just accounted for omissions. A forgetting is a kind of omission in the cases where there is an opportunity to recall something but a failure to do so. In such instances, we can treat forgetting as nothing at all, as we said with omissions. Matters are slightly more complicated than that, however, since there seems to be both a dispositional and occurrent sense of forgetting.

The dispositional sense of forgetting is where some memory, experience or piece of information that was once known is lost. This forgetting is dispositional in the sense that once it is lost it is never again recalled and there is not even a disposition to act upon it. For example, suppose that someone feared clowns as a small child but now has no knowledge of doing so. If asked whether they have ever feared clowns, there is no disposition to answer affirmatively. In the case of forgotten knowledge of facts, such as what is the capital of Nepal, the knowledge might be retrievable through relearning that same fact. Experiences that are forgotten are usually irretrievable, though perhaps some processes such as hypnosis might be able to uncover them again.

Now let us contrast this with occurrent forgetting, though this is admittedly a misnomer since the forgetting means that there is no occurrence. Why the name is apt, however, is that this is a case where someone has the knowledge in question but when an occasion arises where it would be apt to use that information, the information is not recalled. In this case, therefore, there can be a belief, and thus a disposition to behave, but the disposition fails to manifest itself. A paradigm instance of this is visiting the supermarket with a shopping list in your head that, at some point on the way there, includes avocados. You arrive armed with the disposition to buy avocados, but as you are walking around, filling your basket with other items you wanted, you walk right past the avocados and buy none. Forgettings are thus not actions, nor occurrences, nor events; they are failures to act, non-occurrences, and non-events, hence we do not include them in our Parmenidean inventory of what there is.

Let us consider some other cases: specifically, negative epistemic attitudes. I am including scepticism, atheism, and nihilism here. Scepticism, I take to be a disposition not to believe something, which would be an omission rather than an action. There is a distinction between believing that not-*p* and not believing *p*, which is relevant. The sceptic presumably—especially a Pyrrhonian sceptic—wishes to avoid any commitment to belief and thus, we assume, has a preference for not believing over believing that not. This is also a basis for separating

religious agnosticism from atheism. The agnostic does not believe that God exists but neither do they believe that God does not exist. We can indeed say, however, that they do not believe that God exists. The atheist does have the ‘negative belief’ that there is no God. The scare quote is necessary here since the atheist evidently holds some belief, but it is a belief concerning what is not the case. We shall revisit the topic of how to believe that-not in chapter 11.

Nihilism is supposedly a belief in nothing. It is hard to see how anyone could literally believe nothing—let us call this global nihilism—though that is likely to depend on what you take beliefs to be. I take beliefs to be intimately connected with action and dispositions to behave in particular. I do not see that someone with no beliefs at all could act and they would most likely meet their end forthwith. Instead, though, one might take beliefs to be countable mental entities. This makes beliefs sound a bit like having marbles in a bag. One could say that someone has a thousand of them, or a million, though this does not seem a plausible model of the mind. A nihilist would be someone with an empty bag and no marbles. On either of these conceptions, then, a global nihilism does not look a possible position for anyone to be in. Local nihilisms, however, are perfectly coherent, where this is based on a view that there are no facts within a particular domain. Nietzsche (1883) wanted to counter nihilism in the ethical sphere, for example, so presumably he took it to be a position that someone could actually hold. If we realize that God is dead, he feared, then we might have no reason to hold any ethical beliefs, thinking that they are groundless. Presumably, nihilism is different from mere disinterestedness. Someone might have no beliefs at all about the typical diet in Mongolia, but this is not because they think there is no fact of the matter to know. It is more likely that they have never had an inclination to know such a thing. Nihilism is grounded, rather, on the view that there would be no facts to be found even if one looked.

The case of disinterestedness takes us to another negative epistemic state; namely, ignorance. There are many things of which we do not know, where there is no particular reason that we should know or no practical possibility that we could know. There are many historical events of which I am ignorant, plus factual matters such as the number of stars in the universe. On some matters, I’ve opted to remain ignorant. For example, I have a copy of Simon Singh’s (1997) book on Andrew Wiles’s solution to Fermat’s last theorem. I believe that it’s a very good book, and from reading it I could gain a fair understanding of the theorem and its solution. But I’ve never got around to reading it. My ignorance is, to that extent, voluntary, but it is still not a case of nihilism as I have good reason to think that there would be something there worth believing. Not all ignorance is voluntary, nor even a fault, however. We could say that a very young infant is ignorant of almost everything; but faultlessly, of course. We would not say that the child is a nihilist, even if she believes nothing.

With the warning that more is to come on negative beliefs and denials, we can provisionally state that of the negative epistemic states encountered here, there is nothing to cause the soft Parmenidean concern. In most cases, we can consider these as special cases of mental omissions. Nihilists, ignorants, and sceptics might deliberately choose not to form beliefs, just as in action theory someone might choose not to save a drowning man. In most cases forgetting is accidental rather than a choice, although in some cases it could be deliberate: for example, if you want to forget a trauma or a former love. Ignorance too can be accidental, although, again, someone might deliberately choose not to learn about something. We might also consider believing falsely, which could also be classed as a negative epistemic state. Apart from cases of self-deception, believing falsely will usually not be deliberate. From a Parmenidean point of view, I do not see that there is a great difficulty in the idea of choosing not to know something, just as one might choose not to act. Blame can attach to such omissions, but I do not see that a metaphysical difficulty is created by them.

3.8 Logical and mathematical negatives

There are various seeming negatives to which we appeal in logic and mathematics. It may be that some of these are indispensable, for example, much that we do mathematically would become impossible if it had no number zero (Black 2000). Some of the main negatives are considered below.

3.8.1 Zero

Westerners were late developers when it came to the adoption of zero. Barrow (2000: 42) might be right to suggest that a loathing of nothingness persisted in Western thought that was not found in other ancient cultures. Perhaps, then, this is a negative consequence of Parmenideanism: it blinds its adherents to the indispensability of a number for nothing. While there was a Babylonian zero and a Mayan zero, the West remained unaware of it. The zero of the West came from India, where it is clear that nothingness was far more philosophically acceptable than in ancient Greece. A number that stands for it was not anathema, where it would have been to Parmenides. Indian philosophy accepted various forms and aspects of nothingness (see Barrow 2000: 43) and thus had no prohibition against talking about nothing. The bindu, or dot, would sometimes be drawn as a circle, and represented zero as the number for when there was nothing. Along with the other Arab numerals, zero arrived in Europe via Spain a bit before 1000AD.

Why would resistance to zero seem reasonable? Here is one line of thought. While there might be one table, three sheep, or seventy-five trees, nothing can have the number zero. The table, sheep and trees are existences, which are countable. No one, however, can really own zero sheep. That just means that they have no sheep. They do not own some quantity of sheep whose number is zero. In contrast, something can be one, and some collection of things can be more than one. But no-thing cannot be counted. This metaphysical reality was reflected in Roman numerals: a number system without a zero.

The reasons why zero is needed are instrumental rather than metaphysical. The first is so that we can do arithmetic. Zero is needed because it is what you get when a number is subtracted from itself. In Roman numerals, there is no number that can represent such a case: for instance, VII + III – X = ?

Second, once a positional number system is adopted, it is very inconvenient if there is no zero. There were, for example, hieroglyphic numerals that were pictures standing for thousands, millions, tens, single units, and so on. It did not matter in which order these appeared. For instance, 100,014 was represented by a fish (for 100,000), a heel bone (for 10) and four rods (one each). No zero was needed to represent this number and the position of the hieroglyph did not matter since a fish, heel, and four rods would amount to the same irrespective of the order in which they appeared. Such numbers would become cumbersome, however, and often could be read only after performing some preliminary calculations. A positional system means that the numeral's place within the number gives it a different value. On the right, the numeral stands for units. If it is second from right, it stands for tens, third right hundreds, and so on, as we all know. Position matters, which is why we can distinguish the number 32 from 23 and know that the former is larger. Unless there is a zero, however, we have nothing that we can put in the places where there is a missing value. If there are three tens but no units, then without zero we can only write 3. If we just try leaving a space after the three (3), then it is not easy to see that it is there (I assume the extra space can be discerned only with some uncertainty on this printed page). Similarly, the empty value can appear within a number, and if we had to write 302 as 3 2 then, again, it could be hard to distinguish from 32 or from 3,002. Zero removes the possibility of error in a positional number system.

Third, zero has instrumental value in the sciences, even if that means we must adopt the fiction of it being something (more on fictionalism in chapters 5 and 8, below). As Balashov (1999) argues, science invokes determinable properties that can have zero-value determinates. There can be zero degrees temperature on the Kelvin scale, for instance, and zero value spin, charge, or mass.

Does it really matter, then, if we have a number for what is not? Let us recall the seven Parmenidean claims of chapter 1 again. Pvi was that non-being is

unthinkable and unnameable. With zero, we are not only giving nothing a name, we also see that it is useful to have it, if not indispensable. In the practice of mathematics, it is important to think about nothing, and there are certain things that we will have difficulty thinking about unless we do so, for example, that there are three hundreds, two single units, and no tens (in 302). Nought is functionally useful, then. We also saw in chapter 1, though, that the soft Parmenidean need not accept all seven of Parmenides' claims and, indeed, Pvi should be rejected under at least some interpretations. If we can supply a theory of empty reference, that is, an account of how it is possible to speak and think about nothing, then we can accept zero in mathematics without having to grant that it has troublesome metaphysical consequences. Zero can remain nothing. Neither naming it, referring to it, nor using it, makes it a something.

3.8.2 Negative numbers

The matter must be different with negative numbers, however. These are numbers that in some sense are negative, in that they are opposites of the regular (positive) numbers, but they cannot simply be nothing, in the way that zero is nothing. The argument for this is simple since, if -1 and -2 were both nothing at all, then they would be no different. If nothing is nothing at all then it cannot have any distinguishing features that would make one nothing different from another nothing (and this is an argument for why there can be only one nothing). There must, then, be at least some difference between -1 and -2 , and this difference must be accounted for by something rather than nothing.

Reflection reveals that this quandary is not difficult to resolve, however. It is not that there is one negative sheep or two negative sheep corresponding to -1 and -2 , in a case where I owe someone sheep. Rather obviously, it is regular 'positive' sheep that I owe. Similarly, if my bank balance is $-\text{£}1,000$, there is no negative money involved, but I do owe the bank some real 'positive' money. Measures that extend either side of zero are clearly useful, and it is important to distinguish one side of zero from the other. We should accept, however, that both sides are equally real and concern what is. A temperature of -10°C , for example, is just as much of a temperature as $+10^{\circ}\text{C}$, with 0 providing a reference point. (I have friends in the Arctic who sometimes say 10°C when they mean -10 . Since freezing is the default there, it is only when it is plus-degrees that they bother to stipulate.) Likewise, my savings could grow by $-\text{£}10$ or shrink by $-\text{£}15$ although these cases reveal a preference we have towards stating the positive. Unless there is good reason otherwise, we prefer to report these two possibilities as a shrinking by $\text{£}10$ and a growth of $\text{£}15$.

3.8.3 The empty set

The empty or null set has no members. And as the identity of sets is determined by their members, then there is only one empty set. Sets or, as some say, classes, are curious ontological entities in the first place. Abraham Lincoln is different from the singleton set that contains only Abraham Lincoln and one might wonder what the set of X adds to X . Is it like putting X into a bag or container? Sets are not concrete entities, though, even if their members might be, so there is no container in any concrete sense. What existence, then, should we attribute to the empty set: the set with no members? Is it anything at all? If not, then there is no danger to Parmenideanism here. If the empty set has some form of existence, however, there could be a problem.

Why is the empty set even an issue? One reason is that it is an indispensable part of some accounts of set theory and number theory, notably in Frege (1884, secs 74–77) and Russell (1903). In set theory, we have some cases where there are no members in the intersection of two sets, so we have to say that the intersection is empty (hence, that $A \cap B = \emptyset$). In number theory, we can define the number 0 to be the empty set \emptyset , and then define the number 1 as the set that contains \emptyset as a member, that is $\{\emptyset\}$. 2 is the set that contains 1 and 0 as members, that is $\{\emptyset, \{\emptyset\}\}$, and so on for all the numbers. We have already encountered indispensability arguments in chapter 1. By parity of reasoning, should we say that since the empty set is indispensable to (perhaps) our best account of numbers, then we have to accept the empty set as real? I will put aside the question of whether the best account of numbers really does invoke the empty set since that would take us far into the philosophy of mathematics unnecessarily.

Notwithstanding such an argument, however, there are reasons to avoid allotting the empty set or null-class any existence in the ontological sense. This is enshrined by Russell and Whitehead who state: 'to say a class exists is equivalent to saying that the class is not equal to the null-class' (Russell and Whitehead 1910: 239, and proven at prop. 24.54). Russell (1903: 74) explained elsewhere that 'a class which has no terms fails to be anything at all: what is merely or solely a collection of terms cannot subsist when all the terms are removed.' In fairness to Frege (1895: 89), he had said something similar. We will see that the same concerns arise when we consider the empty world so, to avoid repetition, I shall momentarily defer that discussion. Russell's (1903) view seems broadly correct, however. He allows that there can be a null-class concept even if the null-class does not exist. When we look at empty reference, in chapter 8, we will see how this is so. Having the concept of a null-class is enough for us to construct a number series even if there is nothing in reality that corresponds to the class-concept. This supports a claim that the number series is constructed from nothing (Barrow 2000: 166). However, as Russell points out (1903: 75), what the numbers would be constructed from, in this case, would be

the concept of the null-class, and a concept is something. Hence, the number series does not constitute a counterexample to Parmenides' principle that nothing comes from nothing.

3.8.4 The empty world

The notion of an empty world has come to prominence since Lewis's (1986a) modal realism introduced concrete other worlds into metaphysics whose acceptance was justified on the basis of their overall explanatory value. One such function of concrete worlds is to explain modal statements, effectively standing as their truthmakers. What, then, of the statement 'there could have been nothing'? One answer, within a modal realist framework, is to say that there is an empty world, and because of it, it is true that there could have been nothing.

For now, I will set aside the question of whether statements of possibility are true in virtue of concrete other worlds, since this shall be considered in chapter 5. I shall also defer the discussion of whether it is true that there could have been nothing, which will also be addressed there. The concern here is with whether there can be—or perhaps must be—an empty world (chapter 5 explains why these questions can be separated).

Baldwin (1996) advanced an argument for the possibility of an empty world, which has been developed by Rodriguez-Pereyra (1997), known as the subtraction argument. In simple terms, the argument is that each particular thing that exists could have not existed, from which we can conclude that there could have been nothing. For every world that contains a finite number of objects, we can subtract something from it so that there is a world that contains slightly less. We could continue this process until we have a world that contains just one thing, which we call w_{\min} . If we then subtract that one remaining thing, then that leaves us with world w_{null} : the world that is empty. To accept that there is an empty world is the position named metaphysical nihilism.

Our world is manifestly not an empty world. If there even could be an empty world, however, then, given some interpretations, its possibility would violate Parmenideanism. The existence of the empty world would indicate, according to common ways of thinking, that there being no nothing is at best a contingent truth, which the necessitarian strand in Parmenides denies. As some would say, there is a possible world where there is nothing. That there *could* be a world of nothing is significant in itself. But could there be?

The subtraction argument can be resisted. For a start, it is not a kind of argument that holds generally, as Armstrong (1989: 24) points out. He invoked an analogous case concerning diminishing members of an army, but the following non-military analogue is even clearer. Suppose I have a collection of pebbles that I have found on beaches. Each individual pebble is dispensable in the sense that

I could give one away and still have a pebble collection. But that does not mean that I could give all of them away and still have a collection. Effectively, subtraction of one existent from a world might still leave a world but only until the very last subtraction. One cannot move from w_{\min} to w_{null} because W_{null} does not have enough to still constitute a world. A Parmenidean could reject metaphysical nihilism on the basis that there are certain minimum conditions that have to be satisfied in order for something to count as a world and the empty world, since it contains nothing, does not satisfy those requirements. Therefore, there is no empty world.

We should consider, then, what it would take for something to be a world, and then judge whether something could be empty and still qualify as such a thing. First, though, we should be clear what it would be for any such world to be empty. As Coggins (2010: 3) makes clear, the disputed issue is whether there can be a world that contains no concrete existents. I accept that this is what matters since, if there were a world with no concrete existents, it could still be argued that it contained abstract entities, such as numbers and propositions. It might be debated whether there are abstract entities, but if there are, then there seems no reason why any world would exclude them since they would be eternal and necessary existents, like numbers. How concrete is defined is also a question of debate, but it does not matter for my present purposes what exactly this is. One could assume that to be concrete is to be causally relevant to at least one other thing, either as cause or effect, or one could assume that to be concrete is to be spatiotemporally located.

What of worlds? I take it that worlds must contain at least some concrete stuff (this is an advanced modal 'must,' see Divers 1999). We see this if we consider what a world is. One view is that a world is like a bag around everything, such that if the bag were emptied of its contents, it would just leave the empty bag. Coggins (2010: 27) calls this the container view. When it comes to worlds, however, it is implausible that there is anything corresponding to the bag that contains its concrete contents. We might speak of a world 'containing' this or that, but the term seems to be just a way of talking. The universe might have limits beyond which there is nothing, but those limits are not constituted by anything and, even if they were, it would pose a problem for metaphysical nihilism since that would then be something. One assumes that any such existent container would have to be removed in order to get us to w_{null} . The other conception of worlds is what Coggins (2010: 28) calls the compositional view, which is that worlds are composed just by the sum of their contents. On this conception, then, if all the contents are removed then there is no longer a world. Similarly, a collection of pebbles can have any number of pebbles equal to or greater than 1. Below that, it ceases to be a collection: it is nothing at all. Likewise, if we remove every concrete existent from a world, then we have nothing at all. Lewis (1986a: 73) himself held a compositional view of worlds so, for the reason just given, did not accept that there

was an empty world ('A world is not like a bottle that might hold no beer. The world *is* the totality of things it contains').

Efird and Stoneham (2005) think they can offer a way out for the compositional view. The empty world can be constituted in a suitably compositional way but without putting something into it. Such a world can be constituted solely by the 'null individual', which is defined as the individual you get when you subtract anything from itself. This might be a legitimate mereological-theoretic move, but it is not one that we can take with any ontological seriousness. The problem of an unfilled empty world is hardly any better for being 'filled' with another negative entity since any concerns we have of about the credentials of w_{null} will be the same over the null individual (and this worry applies also to the conceptions of null individual I have seen other than Efird and Stoneham's). A world that contains only the null individual does not, after all, contain anything remotely concrete, since at best the null individual is only an abstract entity. The empty world does not, therefore, get filled by anything. On a compositional view, then, w_{null} still fails to qualify as a world. The reasoning above, as advertised, applies also to the case of the null-class or empty set and thus gives us a reason for thinking that it has no existence either.

It remains to be seen what should be said about the statement that there could have been nothing. That discussion will resume in chapter 5. An interim conclusion, however, can be made that even if there are concrete other worlds, of the kind that David Lewis thought were real, the empty world is not in any circumstances one of them. Parmenideanism can repel the apparent challenge from metaphysical nihilism.

There are many topics where nothingness, absence, or negativity come into the equation. In this chapter we have explained away a number of them in a manner that ought to be acceptable to the soft Parmenidean. However, other issues have arisen that warrant a more detailed examination, such as how empty reference is possible, and we will address them in future chapters.

4

Causation by Absence

4.1 Negative causation

Returning to your sole-occupancy home after a two-week vacation, you find that all your houseplants are dried out and dead. They seemed in good health when you left. An intruder could have poisoned them, but let us assume that no foul play has occurred. There are at least two things that we might want to say, (1) and (2):

(1) Lack of water killed the plants.

As a reason for their death, (1) seems a permissible claim, at least *prima facie*. Certainly there was a lack of water. No one was watering the plants while you were away, and plants need regular water in order to live. It also seems quite plausible, then, that:

(2) The plants died because they had no water.

Some might say that there is no real difference between (1) and (2). They both ascribe responsibility for the plant deaths to an absence of water. But (1) is problematic, I will argue, while (2) is considerably less so. It is a mistake, then, to think that (2) is merely a linguistic variant on (1). The first statement makes a causal claim and thereby a metaphysical commitment that absences can be causes. In contrast, (2) only concedes that absences can be explanatory. Now given that causation by absence raises many difficulties, and that it's possible to explain some phenomena non-causally, then the best response to such cases ought to be that absences are not causes but that they can provide explanans for certain explananda. Absences can play a role in non-causal explanation without raising the same difficulties they would if either they were genuine, *bona fide*, causes or contributors towards causal explanations. This is the argument of the present chapter.

Some call this the problem of negative causation but two different types of case are included under that head. The first is where a putative cause is a negative entity, such as a non-occurrence or an absence of something. The second kind of case is where an effect is a non-occurrence or absence. Dowe (2001) and Schaffer (2004), for instance, pick out these two types. The latter, causation *of* absence, is far less problematic and I will disregard it here. An effect can indeed be that

nothing happens, but it raises no serious philosophical problem beyond that of understanding omissions, which were discussed in chapter 3. It is causation *by* absence that is the real problem and sets a serious challenge even for soft Parmenideanism. There are some contemporary philosophers who believe that causation by absence is nevertheless real (Schaffer 2000, 2004, Hunt 2005, Haldane 2011, Barker and Jago 2012, Clarke 2018). Their claims must be taken seriously.

4.2 More cases

Before entering into the details, it will be best to motivate the problem somewhat more. First, it should be recognized that the case with which we began is not a particularly rare or special one. There are alleged causes by absence all over the place. You could fail an exam due to lack of revision, a fire can spread because a door was not kept shut (Beebe 2004: 291), a train could stop due to a mechanism that creates an absence when a communication chord is pulled (Lewis 2004: 286), a spring can launch forth and impact on a target due to the removal of something that was holding it back (Schaffer 2000: 287), human beings can die from lack of water but also from lack of oxygen or lack of food, you can be sad because you are missing a loved one, a mechanism can fail if a part drops out.

Sometimes it is not obvious that an absence is in play, but further reflection suggests that it is. People died at the guillotine during the French Revolution, and it looked pretty obvious that what killed them was the falling of the blade on their necks. But this works only as a form of execution because they lose their heads. The absence of a head causes death because there is nothing holding the blood in the body and controlling vital bodily functions such as breathing. If there is causation by absence, then this is such a case. Similarly, Schaffer (2004: 200) argues that every human action essentially involves a withdrawal of tropomyosin: its removal being necessary for muscle contraction. If we deny that causation by absence is real, then there will be a number of apparent cases of causation that are not genuinely causal since at least one step in the process requires something's absence. Lewis points this out citing the two types of communication chord used on old trains. One type (using vacuum brakes) involved only presences of actual components doing their job (the addition of air pressure creating a force). The other type (air brakes) worked through creating absences (by removing existing air pressure). Those who pulled the communication chords did not know which type they were pulling. Nevertheless, the trains stopped when either chord was pulled. Can one really say that pulling their chords caused one train to stop but not the other? If we prohibit causation by absence, then that is the implication.

A couple of special cases of apparent causation by absence are worthy of mention. First, there are shadows, which we already encountered in chapter 3. Shadows

are immaterial and essentially involve an absence of light that is partially blocked. If absences can be causes, then we would have an example of something immaterial that was nevertheless causally powerful. Plants can die from lack of water, but they can also die from lack of light. Gardeners know this. You could even kill a person by depriving them of sunlight. If it is accepted that causation by absence is real, and that shadows involve an absence (of light), then it seems that we have to allow immaterial causes.

But that is not the most dramatic case. C. B. Martin (1996) provided the example of the deadly void. A void is different from a mere vacuum. The latter can contain photons since light passes through it. A void contains, or is, nothing at all: a region of complete nothingness. Martin asks the question whether, if you see a void heading towards you, you should get out of the way (Martin 1996: 62). To do so suggests that you accept absences can be causes since, presumably, you evade the void because you think it would do you harm. Lewis concurs:

The void is deadly. If you were cast into a void, it would cause you to die in just a few minutes. It would suck the air from your lungs. It would boil your blood. It would drain the warmth from your body. And it would inflate enclosures in your body until they burst. (Lewis 2004: 277)

It is clear that the void is nothing at all. That is how it is defined. But it seems equally clear that you would die if you entered into it. Doesn't that mean that the void would kill you: that it would cause you to die, as Lewis says? But that is to concede that an absence—in this case, an absence of everything—can be a cause. If you don't accept that absences can be causes, you need an explanation of why you would jump out of the void's path.

4.3 Creation *ex nihilo*

A third special case of causation by absence deserves a slightly longer mention since it is often considered the great riddle of existence. Parmenides said that nothing comes from nothing but there is an important instance that challenges this principle; namely, creation, or the origin of the universe.

There are two main views of creation. One is theological: God created the universe. The other is naturalistic: the theory of the Big Bang. Dissatisfaction remains with both views. I offer a different take on the question, which might allow retention of the Parmenidean principle. We should not think of the universe as being caused to be, by nothingness. Absence has no causal power so there is no 'creation' by it. We have principled reasons to say that if there is something, when previously there had been not, then the something cannot have been caused by the nothing. The spontaneous appearance of something, however, would be

permitted. Spontaneous existence remains a mystery also, but could nevertheless be permitted in a qualified defence of Parmenides.

Attempts to 'explain' creation *ex nihilo*, from the very old to the very new, seem always to run aground. Hesiod, for example, posed and answered the question:

Tell me this, you Muses who have your home on Olympus,
from the beginning, and tell which of them came first into being.
First of all came the Chasm; and then
the wide-bosomed Earth, seat of all...

(Hesiod, *Theogony* 113–16, in Barnes 1987: 55–6)

The story produced an obvious retort:

the Chasm is said to have been the first god to be born.

How could that be? He had nothing to come from and nowhere to go if he was the first.

Then didn't anything come first?—No, nor anything second, by Zeus,
of the things we're now talking about: they existed always.

(Diogenes Laertius, *Lives of the Eminent Philosophers* III, 10,
this translation from Barnes 1987: 56–7)

A Chasm, preceding creation, clearly would not provide us with genuine creation *ex nihilo*. There would be no *nihilo*. The question then just shifts to one of from where came the Chasm. This seems an obvious fudge. Yet a comparison is to be found even in contemporary naturalistic accounts. Davies (2013: 46) argues that questions such as 'what caused the big bang to happen?' are 'based on an entirely false picture of the big bang'. There is no 'before' the Big Bang because the universe came into existence with time instead of in time. And 'the question "What happened before the big bang?" is meaningless' (Davies 2013: 48), just as there is nothing north of the North Pole (Hawking and Mlodinow 2010: 136). The following sounds right:

cosmologists now invite us to contemplate the origin of the universe as having no prior cause in the normal sense, not because it has an abnormal or supernatural prior cause, but because there simply is no prior epoch in a which a preceding causative agent—natural or supernatural—can operate. (Davies 2013: 49)

So far, so good. But Davies goes on to spoil it by speculating on precisely that matter. He considers, in the next paragraph, 'why should time and space have suddenly "switched on"?' and suggests that 'spontaneous origination of time and space is a natural consequence of quantum mechanics'.

How, though, could creation be a *consequence* of anything? How is quantum mechanics prior to the Big Bang, in any sense, either temporally or existentially? Even acceptance of the simultaneity of cause and effect in this case would not avoid the problem. If the Big Bang occurred simultaneously with its cause, it only shifts the problem, since we would still want to know where the productive cause came from. If there really is nothing at all, which is presumably what we would need for a genuine case of creation *ex nihilo*, how would the laws of quantum mechanics *already* hold? Surely they cannot, in which case the following is unsatisfactory:

Quantum mechanics is the branch of physics that applies to atoms and sub-atomic particles, and it is characterised by Werner Heisenberg's uncertainty principle, according to which sudden and unpredictable fluctuations occur in all observable quantities. Quantum fluctuations are not caused by anything—they are genuinely spontaneous and intrinsic to nature at its deepest level....The key step for cosmogenesis is to apply this same idea not just to matter, but to space and time as well. (Davies 2013: 49–50)

Is it any wonder that Parmenides thought that nothing comes from nothing? Notably, however, he also offered an argument against creation or generation, which Gallop (1984: 15) reconstructs. What is did not come from what is not. But it did not come from what is either. Therefore, there was no generation. In argument form:

1. If p (genesis), then either q or r
2. Not-q (genesis could not have come from nothing: for i. what is not is not something real and ii. what would have made it generate at the point it did, rather than sooner or later?)
3. Not-r (genesis could not have come from something: for there is no other thing existing apart from what there is)

Therefore not-p.

Suppose some such argument is accepted. Russell, writing on Leibniz, seems to accept it, for example: 'Again, if time were real, the world might have been created sooner, and no sufficient reason could appear for creating it at one time rather than another' (Russell 1900: 120, evoking Parmenides' premise 2(ii)). What, then, can we say of the existence of the universe? One alternative is indicated by Diogenes Laertius in the quotation above: what there is existed always. But this is inconsistent with the scientific theory of the Big Bang. As I indicated above, it is possible instead to resist creation as an instance of causation by absence, and offer qualified support to the Parmenidean principle, if one allows the first existence to be a spontaneous matter. Davies wants to do so, as the above shows, but cannot

resist claiming that spontaneous existence springs from a law. If we have genuine spontaneity then existence did not *come from* nothing, in the sense of being originated, created, or generated *by* it. It came from nothing only in the sense that there was nothing and then there was something. But it was not that the nothing *caused* the something. That is what Parmenides prohibits, and we can give him that. The price of doing so is that we accept spontaneity. If there is such a thing, it violates the principle of sufficient reason, however, which is invoked in Russell's Leibniz, above. Della Rocca (2020) places the principle of sufficient reason at the centre of his interpretation of Parmenides. If that is right, it gives us another reason why we might back away from hard Parmenideanism.

This discussion is by way of explaining that there need be no specific problem of creation *ex nihilo* over and above the problem of causation by absence. Concern about creation *ex nihilo* is understandable, and a book of this nature could not have ignored it. If I am right, however, it is causation by absence that is the real philosophical conundrum here, to which we should return directly.

4.4 Why causation by absence spells trouble

It is worth noting that nothing about such cases seems to hang on what specific view of causation one accepts. I have not said anything about what I take causation to be, starting with a merely intuitive grasp of it. It could be that, as we go along, we see that some theories of causation find causation by absence easier to handle than others. Indeed, there are a number of arguments to the effect that causation by absence, being so widespread and accepted by common sense, can be used as a test of a theory of causation (e.g. Schaffer 2004: 203). If your preferred theory of causation cannot handle it, says Schaffer, then you had better look for another theory. Clearly that is a consideration if causation by absence really does have to be accepted as a real and genuine case of causation, but that contention is precisely what is under examination here. Only if we conclude that causation by absence has to be accepted will we then start looking at which theories of causation are fine and which are in difficulty. As I will show, however, if we reach that stage then we are already knee-deep in metaphysical sinking sand.

What exactly should causation by absence be taken to mean? Thus far, we have been proceeding only with examples. Remaining as neutral as possible, we can start with the simplest account, (A), which says that causation by absence should be understood as:

(A) An absence of X can cause G

I am deliberately building less theory into (A) than Beebe (2004: 294) does into her initial definition. I am not saying that the absence has to be specifically a

non-occurrence of an event, for instance, since that presupposes that causes are usually events, which is controversial. Nor am I saying, at this point, that the presence of X would have prevented G, since that is more part of the theory of what it is to be an absence-that-causes rather than part of the definition of a cause that is an absence. The case for such an additional thesis must yet be made. In most of our examples, above, (A) would be satisfied in virtue of some specific absence: absence of water, light, a closed door, a block on a spring, a missing person, and so on, allegedly producing death, a fire to spread, impact on a target, and sadness. In the case of a void, (A) is satisfied in a general way. A void would be an absence of everything, which putatively causes deadly destruction. Finally, (A) does not, and need not, assert that any absence has, as a matter of fact, caused an effect. The concern is not with such contingencies. The possibility of causation by absence is what is really at stake since it would be damaging enough to Parmenideanism.

There are a number of reasons why causation by absence should trouble us. The first is a two-step argument that would blow a hole in the Parmenidean project. The argument goes as follows:

- I. Absences are causes
 - II. The Eleatic Stranger principle is true
- Therefore, Absences are real.

The Eleatic Stranger principle (Plato, *Sophist* 247d–e) says that being causally powerful is the mark of reality. If something is capable of having effects, then we should accept it as real. There might be other marks of reality, but this one seems fairly compelling. If it is the reality test that we accept, however, and we grant that some absences can be causes, then we have thereby reified absences. They are causally powerful, thus real. If we accept the two premises, even a soft Parmenideanism looks to be dealt a huge blow, possibly fatal since it is Parmenides' Pi that is threatened. We could search for a different reality test (rejecting premise II)—one that didn't deliver the reification of absences. But the soft Parmenidean, as I have defined her, is more likely to take the conclusion as a reason to reject I, that absences are causes.

A second reason why causation by absence should trouble us is the simple lack of a credible understanding of how an absence is supposed to produce an effect. It is one thing to get agreement on a statement like (1), with which we began, but it is another thing to say how an absence of water, or of anything else, is capable of producing an effect. It could well be that once an explanation of how absences do causal work is attempted, then we end up with something that is not so clearly intuitive after all. For instance, someone might just explain the case in (1) by saying that had the plants been watered, they would probably have lived, which is undoubtedly true and could well be what people have in mind when they say

things like (1) but does not contain a factual causal claim. Perhaps it contains a counterfactual one, but it is not explicitly attributing causal power to an absence.

Suppose, though, that we meet someone who maintains that (1) is literally true and their explanation sticks to the story that absent water causally produced the death of a plant. Absent water—and absent anything—is literally nothing at all, as long as we are not equivocating in some way. It has to be nothing for this to be a genuine case of causation by absence. The cause cannot be an absence in name only: a disguised something. This is, therefore, a nothing that we are told produces something, in contradiction of the Parmenidean principle that nothing comes from nothing. How would the nothing initiate a new causal chain? What action would the absent water exercise on the plant or its soil? By absent water, I do not mean some real water that is merely elsewhere, although that would be bad enough in explaining the plant's death, since there is no action at a distance. Rather, I must mean not any actual, real water that has being. I must mean some non-existent water in the plant soil. Or perhaps I mean that the soil is such that there is no water in it. For the explanation to work, we must be talking of a specific absence. A number of other things will also be absent from the soil, but they would have had no relevance to the plant's death. The overwhelming temptation is to say that it is specifically an absence *of water* that did the killing. What model of causation provides any kind of empirically accessible and plausible theory of such a causal process leading from non-X, or X's absence, to G? With Gibb (2013: 200), then, we could just insist that absences have no causal power.

Of course, there are theories of causation that give little weight to considerations such as these. They do not accept that each case of causation involves a traceable process. Causation might be only a matter of there being counterfactuals that hold between events that are made true by what is the case at other possible worlds that are similar to ours (Lewis 1986a). Some might use it as an argument against more robust theories of causation that they cannot handle causation by absence, as Schaffer (2004) does against Dowe (2001). But this is an unstable dialectical position for both sides. It seems all too easy for each to dig stubbornly in. Schaffer can, and does, say that causation by absence is a perfectly coherent notion that satisfies all the core platitudes concerning causation, therefore we need to adopt a theory of causation that embraces it, while Dowe can say that causation by absence is so ontologically troublesome that a theory of causation that allows it thereby betrays its weakness. It will be clear that my sympathies lie more with Dowe on this. As a causal realist (Mumford and Anjum 2011), I do accept a notion of causal production that seems lacking from supposed causation by absence. But I also see that someone who starts from a Humean or Lewisian account, where causation is a matter of events falling into a pattern, with counterfactual dependence relations between some of them, finds absences as causes less worrisome.

But there is a third reason why causation by absence is problematic and it is apparent that all sides of the discussion accept that this is a problem—one that demands at least some explanation from those who accept that absences are causes. This is the problem that if we permit absences to be causes, then we cannot stop a vast and implausible escalation of causes for each and every effect.

The escalation problem is well known, and is discussed by Dowe (2001), Menzies (2004), Beebe (2004), Lavelle, Botterill, and Lock (2013), and Vaassen (2019: ch. 13) among others. It is a negative variant of the causal selection problem. Out of all the possible causes of an effect, how do we identify, in an objective enough way, which were the true causes? My failure to give the plants water apparently caused their deaths, I am urged to accept. But then it is also true that had my neighbour visited my house while I was away, and watered them, then the plants would have lived. So how is it specifically my failure to water the plants that caused their deaths rather than my neighbour's failure? We see how this issue quickly explodes. Had my uncle watered the plants, they would have lived, so why didn't his omission kill the plants? Why didn't Barack Obama kill my plants, or the Queen of England, or Elvis Presley, or the Belgian prime minister, or Julius Caesar, since if any of them had watered the plants, they would have lived? There is no point protesting that Elvis couldn't have watered the plants since he is dead because, even so, it remains true that *had he* watered the plants, they would have lived.

Here is another example of causal escalation, adapted from Armstrong (1999). Suppose I am walking down the street and rubble falls from a high roof above me. Various heavy stones and slates fall all around, but none hits me. I might say that I'm lucky to be alive and, specifically, that I lived through it because none of the rubble hit me. But if my continued living is due to what didn't happen to me, then it seems just as true that I lived because a girder didn't fall on my head, because I wasn't struck by lightning, because a hippo didn't crush me, and so on. If absences can be causes, then I am alive because of the absence of all the possible things that could have killed me, which are practically infinite in number. This one state of affairs—my being alive—has an infinite number of causes. The point would generalize entirely. Everything that happens happens only because of the absence of all the things that could have stopped its occurrence. The charge is that this problem of escalation shows that acceptance of causation by absence, even if it seems reasonable and intuitive initially, should be resisted at all costs, since it has consequences that play havoc with our ideas of what-causes-what and of the notion of causation in general.

4.5 Attempts to stop escalation

As the escalation problem is well known, it is not surprising that there have been a number of responses. I will explain below why they are not satisfactory, but I do

so as a way of motivating what I take to be the best response to the alleged possibility of causation by absence.

4.5.1 Appropriateness

Let us consider the pair of statements below in relation to the previous example. Walking down the street in Durham, parts of a roof fall off into the street, narrowly missing me. A hit would have been fatal. We can consider the following:

- (3) I lived because falling masonry didn't hit me.
- (4) I lived because a hippo didn't attack me.

Lewis (2000: 196) suggests that both (3) and (4) are true. They will be true because there would be true counterfactuals, according to Lewis's usual theory of counterfactuals, relating to the associated causal claim. In the closest possible worlds in which I am hit by masonry, I die; and in the closest possible worlds in which I am attacked by a hippo, I also die. However, it would be inappropriate to assert (4), in these circumstances, even though it is true. There are good reasons why it is inappropriate. Durham does not have hippos walking the streets. There never was any danger of one attacking me. In contrast, masonry does fall from a roof, in the example we are entertaining. I had a narrow escape from it. Perhaps one could say that the possibility of me being hit was far closer than the possibility of a hippo attacking me.

However, I follow Vaassen (2019: 186) in this case: (4) is not true, contrary to Lewis's view. If I say that I lived *because* of X, then it is not just supposed that X is the case but also that it is connected in some way to my survival. But I did not live *because* no hippo attacked me even though I concede that I would have died if, in some far-fetched scenario, one had attacked. By contrast, it is right to say that I lived because masonry missed me. We judge that most absences are irrelevant to the effect in question. A stone falling on my head is the right absence to explain why I lived. The absent hippo is not. Hence (3) is true while (4) is not merely inappropriate but flatly false.

4.5.2 Abnormality

Hart and Honore (1985: 38) suggest something similar, but in a way that emphasizes the ontology of causation rather than the pragmatics of what it is appropriate or inappropriate to assert. They identify normality and abnormality as the key issue. It could be said that it is specifically my failure to water my plants that causes them to die rather than the Queen of England's failure since it is normal

that I water the plants while it would be abnormal for her to do so. You cannot blame an effect on something's absence when it was not something that happens normally, or so Hart and Honore claim.

However, as Beebee (2004: 295) argues, following Stapleton (1994: 122), there are plenty of cases in which such a distinction is of no use. In the falling masonry case, for example, it is not that masonry normally falls on the street in Durham. It might be a very rare or even unique event. And it is not, which is needed for (3), that masonry usually falls on the street in an exact pattern that leaves a safe spot in the middle: the spot where I happened to be standing at that moment. I survived because I got very lucky. I was unlucky to be walking by at the moment the masonry fell but then lucky that it fell in such a way that I narrowly escaped. Normality is not much use, then. In the example, it is virtually a miracle that I am unharmed. Very abnormal and highly improbable things happen all the time. Just consider the specific six lottery balls that the machine spews out in the weekly draw: each time, a massively unlikely event occurs. We cannot distinguish genuine from spurious (negative) causes on the basis of which are normal and which are abnormal, then.

4.5.3 Contextual salience

Schaffer (2005) argues that spurious cases of causation by absences can be ruled false in a principled way, if we look to the relevant contrast classes. This claim appears in the context of Schaffer's view of causal contrastivism, which says that the proper understanding of a causal connection is not as a two-place relation—of A causes B—but as a four-place relation that includes contrasts for each of the cause and effect. A causal claim is to be understood as having the general form: A rather than F causes B rather than G. Instead of stating simply that my failure to water my plants caused them to die, I should say that my failure to water my plants, rather than watering them, caused them to die, rather than live. I am considering this view only as a theory of causation and not of merely causal explanation (hence I will not contradict the use Lavelle, Botterill, and Lock (2013) make of contrast in explanation, since I will defend a similar account later). This works as a genuine (negative) causal claim because the contrasts it employs are contextually salient. I could have watered my plants and doing so would have had a bearing on whether they lived or died.

In comparison, if we assert that the Queen of England's failure to water my plants, rather than water them, caused the plants to die rather than live, then this is to be rejected. Schaffer (2005: 332) offers the following explanation: 'The queen's watering my flowers is not easily swallowed as a *relevant alternative*.' The Queen was never going to water my plants, whereas it is reasonable to suppose that I would. After all, I am the one who has watered them every day until I went on holiday.

Vaassen (2019: 189) has argued, however, that even if the contrast case seems a relevant alternative, the spurious cases still come out as false. The point here is not whether the spurious cases are false, since I agree with Schaffer, and many other parties to the debate, that they are indeed false. What the defender of causation by absence is looking for is a plausible and principled explanation of why the spurious cases are false but which can still preserve the intuition that there are some genuine cases of causation by absence.

Take the falling masonry case again. This time let us compare it with the absence of a car running me over on the street as a cause of me living. In the case, we still say that I am living because of the falling masonry's failure to hit it me, rather than a passing vehicle failing to hit me, even though people in Britain are far more likely to be killed by vehicles than falling masonry. So if we consider the contrastive account of the latter, we get it that 'A vehicle's failure to hit me rather than hit me caused me to live rather than die.' This seems contextually salient, as Vaassen says, but the contrast—a vehicle hitting me—is 'easily swallowed' as a relevant alternative to it not hitting me. The contrast is quite reasonable for a British street. Yet, we still think that the case is false; that is, I lived through the day because the falling masonry failed to hit me rather than a passing vehicle failing to do so. Schaffer's apparatus has not helped us, therefore, to distinguish spurious from genuine cases of causation by absence and consequently, according to my argument, he has failed to free causation by absence from the problem of escalation.

4.5.4 Stability

Having cited his views of other accounts positively, I come now to Vaassen's own account, which derives from an idea in Woodward (2006). Woodward argues that causal connections are stable, that is, they hold true through a number of variances in other conditions. For example, water dissolves sugar whether it is in a dish or a glass, whether it happens to be a Thursday or a Friday, whether the water is 20°C or 30°C, and so on. A dependence relation will count as stable if it continues to hold against a wide variety of natural changes in background conditions. Vaassen (2019: 191–4) utilizes this insight to separate genuine negative causal claims from spurious ones (Vaassen prefers to use the term insensitivity instead of stability). Let us follow some standard cases and assume that (5) is genuine while (6) is spurious.

- (5) My failure to water my plants caused their deaths.
- (6) The Belgian prime minister's failure to water my plants caused their deaths.

The task is to explain why (5) comes out as more stable than (6) since if I fail to water my plants then it would require far-fetched changes to the background

conditions for my plants to survive. One such change would be that the Queen of England breaks into my house and waters them. Another is that my roof starts leaking and the drips fall exactly into the plant pot. Given the most natural assumptions, then it remains stable that if I fail to water the plants, they will die. In the case of (6), however, if the Belgian prime minister fails to water my plants, the changes to background conditions that would see them live are not as implausible: for example, that I water the plants instead.

However, as Vaassen (2019: 194–8) admits, it is not clear-cut what count as natural changes in background conditions, nor what count as slight and extreme changes to those conditions. He accepts that what counts as a natural change will depend on the conversational context. For example, the presence of oxygen will be a very natural change in many of the conversational contexts we know but it would be an unnatural change if we were talking about events in outer space. The stability of a causal relation then becomes a context-dependent matter, and the same has to be said, therefore, for any distinction we make between genuine and spurious negative causal claims. This finding seems reasonable. We know, after all, that there is some context in which we would indeed judge that it is the Belgian prime minister's failure to water my plants that killed them: where we had a prior agreement, for example.

One might argue that this example is misleading. What killed the plant is absence of water. Who was responsible for the water's absence is another matter and it is no wonder that practical, normative and relevance considerations come into that judgement (see Beebe 2004). It is common to argue that when we look at causation ontologically, whether A caused B cannot be determined in any degree by such pragmatic matters. We ought to accept that view. Causation is one thing and responsibility is another. Clearly it is my *fault* that the plants died, since they were mine and I should have arranged for a neighbour to visit and water them for me. Had I made such an arrangement, to which the neighbour willingly agreed, but failed to fulfil their promise, then it would be their fault. It is hard to imagine any likely situation in which the fault was the Belgian prime minister's, since we do not know each other. Similarly, in Dowe's *reductio* of causation by absence, among the causes of Kennedy's death were the absences of any human being who ever did, does, or will live, who could have been stood between the assassin's bullet and Kennedy. Anyone could have taken the bullet for JFK, but only his bodyguards could be blamed for not doing so.

As we are looking for a metaphysical solution to the problem of escalation, pragmatic considerations are of limited help only. Normative and contextual accounts have the consequence that a negative causal claim is only spurious, or genuine, relative to some context. This does not banish the escalation of causes, considered ontologically, then. It might, instead, explain and vindicate the escalation. It is a possible view, after all, that the escalation is quite natural and ought to be accepted. We need not take that path, however. The escalation problem is

not the only one that causation by absence faces and if we can explain away the latter then we might not have to swallow the bitter pill of the former.

4.6 Other ways of explaining causation by absence

I am going to move on from the problem of escalation now, accepting that it remains unresolved, and consider other responses that are attempts to explain, or explain away, causation by absence. We will see that each of them offers us something useful that will be incorporated into a satisfactory final account.

4.6.1 All causes are positive

One response, that will attract Parmenideans, is to insist that all causes are positive. The response is, then, to deny that there is causation by absence. But this is not just a flat denial. It is also a claim that what appear to be cases of causation by absence can be explained away in perfectly naturalistic, 'positive' ways. Lewis gives us a model of how this would work in practice. After initially calling the void deadly, he then shows how it might not really be doing anything at all:

When the void sucks away the air, it does not exert an attractive force on the air. It is not like a magnet sucking up iron filings. Rather, the air molecules collide and exert repulsive forces on one another; these forces constitute a pressure that, if unresisted, causes the air to expand and disperse... Likewise, when the void boils the blood there is no flow of energy from the void into the blood. It isn't like a stove boiling a kettle of water. The blood is already warm enough to boil, if its vapour is unresisted... [W]hen the void drains your warmth, what happens is that your thermal energy, left to itself, tends to dissipate; and the void provides no influx of energy to replace the departing heat. (Lewis 2004: 277)

The total absence that is the void, then, is not doing anything. It is not making anything happen. It is the things cast into the void that cause their own disintegration. They have forces within them that in more standard situations are resisted and counterbalanced. It is those forces that do any causing that occurs. Although Lewis does not use these cases to support such a conclusion, this understanding could allow us to resist the pressure to accept causation by absence. The withdrawal of a counterbalancing factor might be the occasion for an effect, but that does not mean that it is the cause of it. The cause of it will be the present causal factors that remain when the other is withdrawn (as in Mumford and Anjum 2011: 147).

Is it plausible that we can generalize this as a solution? Consider whether the cause of a room's cooling is that a heater has been turned off. This is *prima facie*

causation by absence. Isn't it better, though, to say that the room temperature was being maintained, in spite of ongoing heat dissipation, by the action of an agent that was more than compensating? Turning off the heater withdraws that action, but it is not the heater's inaction that causes the room to cool. The remaining and still operating causal factors do that. Similarly, the absence of water produces no action on the plant. What is still present—a plant going through a process of photosynthesis that requires water—does all the causing there is. Usually, the plant sucks water out of the soil so that it can live and grow. In the example, there is a point when all the water is used up and does not get replaced, meaning that a required process is unable to continue.

This response, as it stands, is not entirely satisfactory, however. It is a good start and, I maintain, correct as far as it goes, but it is not the whole truth. A specific absence is essential in the explanation of why the plant died, the room cooled, or someone exploded inside a void. The absence seems to be indispensable to the explanation because, even if what is present does the causing, what is present does not entail what is absent, and thus cannot deliver the absence that the explanation requires.

For example, we could have a full description of the plant, listing all that there is about it. We can include an itinerary of the plant's environment in case it is relevant. We can give all the details about the soil in which the plant sits (although we should resist adding that the soil is dry since this might be a way of smuggling in an absence of water to the description when we are trying to explain the plant's death only in terms of what is present). The problem is that there could be everything that is present plus something else, such as water in the plant soil. What is there clearly does not exclude also the presence of water. Indeed, before leaving for the holiday, everything was just as it still is plus there was water in the soil. So if we can have everything that is present, let us call this total presence *P*, plus water besides, *w*, then it cannot be that *P* excludes *w*, which means that *P* cannot entail $\neg w$. The existence of *P*, therefore, cannot alone explain why the plant died because it does not include among it the absence of water—since *P* is only what is present, not what is absent—and nor does it entail what is absent. Without $\neg w$, we do not have an adequate explanation of why the plant died.

Of course, at this point, one could relinquish one's Parmenideanism. Perhaps the soft Parmenidean should say that this is one point at which the project has failed and that we have negative entities—absences—that are an indispensable part of our best theory of causation. But that is to give up too easily. What we can do instead is look at other responses to the problem as we might find some helpful suggestions there.

4.6.2 Possibilia are causal

Bernstein (2015) has suggested that we can account for causation by absence if we allow that possibilia are causal, a view also supported by Hommen (2016).

Bernstein suggests that a failure to water a plant is a watering of a plant at another possible world and that this possibility is a positive entity, as explained within the modal realist theory of Lewis.

This is a difficult view to swallow. In the first place, if one is prepared to step outside the Lewisian orthodoxy, it can be questioned whether mere possibilities are real existents in any case. The question of their reality is at least contentious, we can say, but we will defer that detailed discussion to chapter 5, below. Suppose we granted that possibilities were real, however, then there would still be huge issues to resolve. Why should something not happening in our world require something happening in another world? Put the other way round, how can an event at another world stop something happening in ours? Or to put the point in a way reminiscent of the so-called Humphrey objection (Kripke 1980: 45), how can anything that happens in another world be relevant—in this case causally relevant—to what happens in our world? Finally, in the case that concerns us, how can someone watering a plant in another possible world causally contribute in the slightest to the death of a plant in this world?

Nevertheless, I do not want to say that Bernstein's suggestion is without value. What it shows is the importance to our explanations not just of what happens but also of what doesn't happen but could happen. This is especially important in the case of causation since our judgements, as Hume (1748: 56) says, are not just about one thing following another. There is also a component of a causal judgement concerning what would have happened had things been different: namely, if the cause had not been, would the effect have been either?

4.6.3 Quasi-causation and counterfactuals

The next suggestion, from which we shall take some assistance, is Dowe's (2001: 217) view that causation by absence is not genuine causation but that it involves what he calls quasi-causation. The basis for this is that although it is implausible that omissions can be causes, for reasons already documented, they do support similar counterfactuals to the ones that are true in genuine cases of causation. For example, suppose that a plant is overwatered and dies. Part of the reason we think of this as a causal situation is that we feel inclined to accept the counterfactual that had the plant not had so much water, then it would have lived. It seems acceptable to move from (i) C1 causes E1 to (ii) if \neg C1, then \neg E1. Some say that (ii) is a more or less complete account of why we accept (i). It is not my view, but that debate is largely orthogonal to our current concern.

Now in the putative cases of causation by absence, similar counterfactuals hold. Where the plant does not have enough water, we want to say that if it had received enough water, then it would have lived. We can think of this as a case of moving from (i*) \neg C2 causes E2 to (ii*) if C2, then \neg E2. But Dowe thinks that \neg C2 is incapable of being a cause. It is an absence: it concerns what is not. Dowe

believes that causation involves the transfer of a conserved quantity from cause to effect, and you certainly cannot have a conserved quantity originating from an absence. So (i*) has to be rejected. There is no such causation. Nevertheless, we can still accept the counterfactual in (ii*), which is what Dowe takes as grounds for calling the case quasi-causal.

I will make some use of Dowe's general framework but will not take from it a notion of quasi-causation since it is potentially misleading. 'Quasi' suggests that there is almost or nearly causation in these cases and, if it is problematic that absences can be causes, then it is also problematic how they can be nearly-causes too. Parmenides warned us not to allow degrees of being. To permit that absences are almost causes is to allow them almost-existence. We might easily lose our way if we go down that road, and I'm sure it cannot be what Dowe had in mind. I am going to insist that absences are no types or degree of cause at all.

What I will take from Dowe's account, however, and building on Bernstein's connection of absences with other possibilities, is the counterfactual framework, for this could show how the absence of something can be explanatory. In turn, this opens the door for an account in which absences can figure in non-causal explanations but which are then mistaken for causes. Simply put, we might think of absence of water causing a death, but that is because of our theoretical background belief, which is inductively well grounded, that if the plant had been watered, it would have lived, exemplifying the structure of (ii*). The truth of such a counterfactual can be invoked in the explanation of E by the absence of C. Had C been present, E would not have happened.

Some might be troubled by such a move. It is hardly as if counterfactual conditionals are themselves an uncontroversial topic in philosophy. We may just be jumping out of the causation-by-absence frying pan into the counterfactual fire. Some think that the best way to understand counterfactuals is within the Lewisian framework of possible worlds, about which Lewis and his followers are realists, and concrete other worlds do not seem to be the sort of existents that will please Parmenideans, as I will consider in chapter 5.

Nevertheless, counterfactuals are indeed useful to our primary topic and we should not be put off them altogether just because they are often understood a certain way. A naturalistic, this-worldly account of the truth of certain counterfactuals is readily available for the case of causation by absence. What makes it true that if the plant had been watered, then it would have lived? The intuitive response is not to look at some other possible world where the plant was actually watered but instead to look at the causal power of water to hydrate plants. This power to hydrate is real and natural and plants, plus other organisms, have evolved to use it. As noted above, there is a good inductive warrant for saying that, had the dead plant been watered, then it would have lived, and this is what provides some explanatory power for invoking its absence when a plant has died.

4.6.4 Causal relevance?

As I have said why a notion of quasi-causality is not helpful, it is worth raising a similar issue. Martin (1996: 64) claims that absences and voids can be causally relevant without being causally operative. How so? I find Martin a little unclear on this, but the position seems to be that although a void or absence is not a thing, it is *how a spatiotemporal region is*. An absence cannot act—it has no causal powers—but it can nevertheless have reality in that although it is not a thing or entity, it is a way some region of the world is. Hence, there is no such *thing* as the absence of water in my plant pot, but my plant pot *is* nevertheless such that there is no water in it, and this is causally relevant to the plant's death.

Again, I take it that there is something useful about Martin's way of describing the case but also something not so useful. What is useful is the capacity to talk about absences while rejecting their credentials to be entities or existents. In chapter 10, I will say some more about how we are able to talk about absences without committing to their existence (spoiler: I can *deny* that there is water in the pot instead of *asserting* that there is an absence of it). So invoking an absence does not automatically take absences into the realm of being. That would be problematic because we would attribute being to something that is a non-being, putting ourselves at risk of contradiction. But we can deny of a region that there is something in it.

As I have indicated, this absence can be essential to the best explanation of a certain occurrence, as it cannot be reduced to the presence of other things. What is misleading, however, is to still maintain, as Martin does, that the relevance of such absences is a *causal* relevance. As with quasi-causation, this seems to concede that the absences have at least some significance or importance or role to play in what is caused, which is a dangerous theoretical move.

I accept that there are some things that are causally relevant without being causes, so it is not that the distinction is without use. For example, it is causally relevant that two bodies are a certain distance apart if you want to explain their relative motions. They gravitationally attract. Their distance will be causally relevant to their degree of mutual attraction and, in turn, their motions. But the distance does not cause those motions. The motions are caused by the bodies acting on each other in virtue of their respective gravitational masses. It does not seem right to invoke absences in a similar way to this, however, and attribute causal relevance to them (again, I am not entirely sure what Martin had in mind). Distance plays a role in determining the degree of attractive force between two bodies, as described in the gravitation law. The role of distance is not to cause, but it does non-causally determine in what way causation occurs. Similarly, it is causally relevant to sugar's dissolving in coffee that the sugar is put in the coffee. Being placed in the coffee does not cause the sugar to dissolve. The coffee causes the sugar to dissolve. Being put into the coffee explains only how the sugar got into a position to be causally acted upon in the first place.

We should be reluctant to follow Martin's lead in extending this notion of causal relevance to absences, however. If these are genuine absences, and not just presences by another name, then they really are nothing at all. And how can a nothing have any influence, causal or non-causal, on the rest of the world? Absences cannot make things happen; and they cannot influence the way in which things happen either. It might matter that there is an absence, but we can account for what happens in terms of what is present, as explained in 4.5.1, above. That leaves explanation by absence, but explanation by absence is a different matter from causation by absence, as I will now detail.

4.7 Explanation but not causal explanation

The position with which I settle bears some similarities to Beebe's (2004, see also Varzi 2007), but with some major differences too. Beebe is right, I agree, to separate causation and explanation. Absences simply cannot be causes. They can be explanations, however, in a way that I outlined in 4.5.3, using Dowe's counterfactual framework. I dissent from Beebe, however, where she accepts that absences can provide causal explanations. They cannot, for reasons I will elaborate. I can thus sum up my position, and solution, as follows. Absences cannot be causes. They can be explanations. But they cannot be causal explanations. The taking of absences to be causes mistakes their epistemic, explanatory role for a real metaphysical one, the danger of doing so being that it thereby commits us to the reification of absences.

Beebe accepts that absences are not causes but is willing to allow that an event happened *because* of something being absent. I am happy to say this too. I am accepting statements of the kind given by (2), at the start of this chapter, but not statements of kind (1). A mistake is to take it that (1) and (2) are just two different ways of saying the same thing. But there are *because*s that are not causal *because*s, so one cannot infer (1) from (2). There are good grounds for asserting claims of the same type as (2)—explanatory claims—but there are not good grounds for asserting claims of type (1), since absences cannot be causes.

Absences can provide causal explanations, Beebe nevertheless argues, but she is here invoking David Lewis's notion of causal explanation where 'to explain an event is to provide information about its causal history: information that need not be restricted merely to citing the event's causes' (Beebe 2004: 293). It can be inferred from this that a causal history need not simply be a history of causes, according to Beebe.

It is right that there is more to a history of some event or thing than its history of causes. To say that something persisted between t_1 and t_2 , for example, says nothing about what caused it to be or what caused any of its changes between t_1 and t_2 . For a history to be a causal history, however, there does have to be some

additional constraint. The most straightforward way in which something can be a part of the causal history of something else is to be among its causes, or the causes of its changes. Beebe and Lewis might be right that there are other ways to be part of the causal history, and I would have thought that being of causal relevance, as described above, was one other such way. But there has to be some limitations, if the distinction between a history and a causal history is to be made. For reasons given above, I cannot see how absences, if they really are nothing at all, can be a part of something's causal history since absences cannot be either causes or causally relevant. Nor would I be happy including negative facts in a causal history, as an intended less innocuous alternative to absences. For reasons given in chapter 3, negative facts are likely to be just as offensive to Parmenideanism. Until such time as a credible account is found for how absences, or other negative elements, play any causal role at all, then we should resist the idea that they are part of the causal history.

So we should reject Beebe's view that absences explain because they causally explain. That does not matter too much, however, since we have already seen how absences can be explanatory without being in any way or degree causal. Causing is done by presences, but there are a number of true counterfactuals concerning how things would have been different if those things that are absent were instead present. Furthermore, we can explain why some effects occurred at a particular time where the presence of something ceased at that time. This is where we can say that an absence is the occasion of some effect rather than its cause.

It might be thought that a matter has been resolved at this point. There was a problem of escalation if absences could be causes, where an effect would have a potentially infinite number of causes, and how could we select only a subset of those causes as real without appealing to illegitimate normative or pragmatic considerations? However, don't we have exactly the same problem if we say that instead of absences being causes, there are just true counterfactuals, for there would be at least one true counterfactual for each putative absence that was a cause? How, then, do we distinguish counterfactuals that matter from those that don't, given that, again, there will be an infinite number of true counterfactuals whenever we have an occurrence?

We see that this is no longer a difficulty, however. In a case where the absence of water seemed to cause a death, it is obvious which counterfactual is relevant and is in mind when we consider the case. The counterfactual is, to use Schaffer's term, a contrast case to the absence. The counterfactual to consider is the one where that which was absent is instead present. Second, however, given that there is no ontological claim of causation involved anymore, in the theory I advocate, then it is no disadvantage at all if pragmatic and normative considerations select the counterfactual, and thus the specific absence, that is invoked in explanation. For example, in the case where it is my failure to water the plants that killed them, rather than the Queen of England's failure to water the plants, there are true

counterfactuals corresponding to both. It is certainly true that if the Queen had watered the plants, then they would have lived. But given that in neither case is it allowed that failure to water the plants killed them—that is, caused them to die—then a normative reason why one counterfactual is more relevant, stable, or normal than the other is entirely acceptable. This does not commit us to a normative conception of causation itself (e.g. McGrath 2005).

4.8 Summary of conclusions

Causation by absence has produced much consternation and confusion among philosophers. Soft Parmenideanism has given us a reason to resist flippant acceptance of absences as causes, however, since they might be appealing *prima facie* on intuitive grounds, but their acceptance would wreak metaphysical havoc.

We also found, however, that those who have gone before have made partially useful contributions which, I have argued, can be pieced together to form a satisfactory overall picture. There is no quasi-causation but there are relevant counterfactual truths. Absences can be explanatorily indispensable without being real. Explanations can be useful and informative without invoking causation. And we can stick with the good Parmenidean principle that only present things—what exists—can be causes. The principle has resisted the challenge from causation by absence. We can thus take any standard example and apply the account. Absence makes the heart grow fonder, we are told, and certainly I can be sad because my beloved is away. Their absence does not cause my sadness, but it is an ineliminable part of the best explanation of my sadness. The explanation is non-causal but rests on a variety of other considerations that pragmatically select a counterfactual conditional as the most relevant: had that person been present, I would most likely not be sad.

We are able to conclude, therefore, that there need be no threat to the Parmenidean project from causation by absence. A number of philosophers have accepted its reality, but I have argued that there is no compelling reason to do so. In particular, causation by absence is not an irreducibly negative, indispensable component of our best theory of the world.

Mere Possibilities

5.1 Possible queues ahead

On British roads there are warning signs that display messages. One such message warns of congestion with the words ‘Possible queues ahead’. Is this alarming? It was suggested to me that it shouldn’t be (by Anthony O’Hear, in conversation). A possible queue cannot delay you; only an actual one can. A queue that is merely possible should really have no effect at all, since it is not anything. You should in general be concerned only with what is actual. This example introduces the topic of the present chapter: the serious issue of whether possibilities have being.

An important clarification is needed immediately. It is accepted in standard modal logic that p is possible if p is actual. Some possibilities are thus also, technically, actual realities. This is right since it would be strange to say that something actual was not possible. Where a possibility is also an actuality, then there is no problem in allowing that it is real. The Eiffel Tower is a possible object that is also an actual object, and it clearly has being, existence, or reality. The concern of this chapter is with ‘mere’ possibilities, which are defined as those that are possible but not actual. There are many such possibilities. Rosa could have gone to London today, but she didn’t. Hillary could have won the presidential election, but she didn’t. A hippo could be in the room right now, but it isn’t. These are mere possibilities.

Mere possibilities seem to have no being. The merely possible hippo in the room is nothing at all, we must surely say. Rosa didn’t go to London today so the possibility of her being there is not something. Had she gone, the state of affairs of her being in London would have been actualized and would then have had existence. As long as it remains a mere possibility, it has none. Indeed, one might say that a mere possibility becomes an actuality precisely by coming into existence.

A Parmenidean who wants to say this about the mere possibilities faces a challenge, however, since a flat denial of the reality of mere possibilities leaves something unexplained. The problem is that some possibilities seem genuine or real and some do not. It is a genuine possibility for Rosa to visit London, but it is not a genuine possibility for her to jump to the moon under her own power. If, in Parmenidean fashion, we wish to treat mere possibilities as nothing at all, denying them any being, then how do we distinguish between those possibilities that are real and those that are not? What is meant by ‘real’ or ‘genuine’ here? We will see that this leads to the question of whether we should seriously consider

reifying the possibilities that we think of as genuine. Perhaps the possibility of Rosa visiting London does have some kind of existence, while the possibility of her jumping to the moon does not. It is fairly standard to deal with such a difference by saying that there are types of possibility. A visit to London is *naturally* possible, which jumping to the moon is not, though the latter is at least *logically* possible. There could be a number of different domains of possibility that we invoke. We can talk of what is morally possible, for instance, or legally possible or metaphysically possible. I am not going to make too much of these standard distinctions, not because they are unimportant but because I want to get directly to the issue of whether a mere possibility has being in any domain at all.

It might be that the way of posing the problem has already led us down the wrong path. By speaking of 'a possibility', one is already treating possibility as thing-like (and perhaps such thought has infected consideration of absences elsewhere too). One is using a count noun, as one can speak of a possibility, two possibilities, or several. This makes possibilities sound like individuals. A different locution is to say, 'it is possible that...', which instead suggests that possibility is more like a property or attribute. It might be a property of certain statements or, more seriously, abstract states of affairs. One could also think of possibility as a modification or mode of a proposition—a sentential modifier—which is where the 'modal' in modal logic originates. It could then be that real though mere possibilities only seem like an issue for Parmenideans because they have been conceived as if they are candidate existents, which they are not. Indeed, something like this view will be defended later. Even if that is the case, however, the question of the being of what is merely possible would still need to be asked. After all, we would still be entitled to ask what the being or existence was of the property of being possible. Why would some things have this property and not others? And what would bear the property, if it was only possible? We can say 'it's merely possible that *p*', which suggests that *p* has no actuality, so how can it instantiate a property of possibility? Similarly, if we take possibility as a mode or modification of a statement, it is again reasonable to ask why some but not other statements are so modified. The statements exist, certainly, whether they are modified or not, but there is a question of whether it is true or correct to modify such a statement and if so why.

The goal, as with other chapters, is to explain the phenomenon without reifying nothingness. Here, there are two ways we can go. We could, as with previous topics of negative properties and absences, decide that mere possibilities really are nothing at all. Alternatively, we could find that they do indeed have some being, so they are not, after all, nothing. Either response would be acceptable to the Parmenidean. What we want to avoid is a position in which mere possibilities are taken to be both nothing at all but also real existents. From the foregoing, however, we already have seen that both these views are tempting individually: mere possibilities seem like they are nothing at all and yet some mere possibilities

are real and some are not. Nevertheless, we should resist holding this combination of views.

In the remainder of this chapter, I consider two accounts that effectively reify mere possibilities. After finding that neither is entirely satisfactory, I then consider non-reification; that is, a way of showing that mere possibilities are not real. I conclude in a way that should be satisfactory to the Parmenidean. Mere possibilities really are nothing at all, as we first thought, but some of them are grounded in what is actual, and it is these that we consider to be genuine possibilities. Without making possibilities actual themselves, therefore, we can see why it is nevertheless reasonable to distinguish the genuine from the non-genuine possibilities.

5.2 Reification

I will consider two accounts of mere possibilities that have some popularity and can reasonably be interpreted as reifications. Mere possibilities are real things in the sense that they have some being or actuality. The first view is modal realism and the second is a version of the powers view.

5.2.1 Modal realism

The modal realism I am talking about is the full-blown realism developed in detail by David Lewis (1973, 1986a), from which I am excluding ersatz readings of possible worlds (I will turn to them later). The view under consideration reifies possibilities in the sense that what is merely possible for our world is an actuality at some other concrete world where there is a plurality of such worlds, each spatiotemporally disconnected from the others, but each just as real as any other. Hence, the mere possibility of Rosa being in London today is an actuality at some other world where Rosa does visit London on this day. What is nothing at all in our world is elsewhere something. According to Lewis's view, there is a world for every possibility, hence every mere possibility of our world will be actual somewhere in just the same sense in which things are actual in our world. 'Actual' works as an indexical in this view (Lewis 1986a: 92ff). By saying that something is actual, we just mean that it is true within the world that we occupy. Hence, as I approach a sign that says there are possible queues ahead, what I worry about is that I might be one of the unlucky people to inhabit a world in which there really is a queue. What I am hoping is that I live in a world in which there are no such queues and it is some other poor soul who will have to wait in line. Someone is going to be delayed: I just hope it isn't me.

The advantage of this view is that there is no qualitative difference between the actual and the merely possible. They are both equally real and qualitatively

indistinguishable in theory (not that anyone could be in a position to actually compare them). Specifically, it is not as if the latter has some shady quasi-existence, or subsistence, and needs to somehow acquire an additional property of actuality in order to become real. Nor are there different realms: non-existence and existence, which the possibilities must traverse. There is no such travelling in modal realism. Each world just is what it is: a patchwork or mosaic of unconnected events—just one little thing and then another (Lewis 1986b: ix)—and there is no sense *within* a world of what is possible or necessary. Each world is effectively a Hume world: a world in which things are as Hume (1739) describes in the *Treatise*, conjoined but never connected. Lewis's achievement was to show how a theory of modality (1986a) and a theory of counterfactuals (1973) could be produced within a Humean framework: if the Hume worlds were plural instead of there being just one of them.

Whether this should be considered a realism about modality is moot, since there are no intra-world modal facts. But modal realism is an apt name for our present purposes since, as it suggests, the theory is realist about the possibilities: they are just as real as the actualities since all are actual in some world. Indeed, it is essential to Lewis's theory that there is at least one world where each possibility is actual, thus there will be uncountably many worlds (Lewis 1986a: 2). There will be worlds where Rosa goes to London today, worlds where she goes to Australia instead, worlds where she murders her neighbours, worlds where she wins the Wimbledon tennis tournament, worlds where there is no Rosa at all, and so on. We can impose some restrictions, however. No two worlds are exactly alike in every respect, since the second such world would be redundant in the theory; though Lewis (1986a: 224) himself was pretty agnostic on this issue. And if the identity of each world is fixed by its contents—that is, by what is the case there—two worlds with all the same contents would really be one world.

Some see Lewis's realism about other worlds as ontologically extravagant, not so much in suggesting new types of existent, since all the other worlds are just as real and concrete as ours and differ only in what particular things happen there, but it is a quantitatively extravagant theory. The usual intuition is that there is only one world—a single spatiotemporally connected universe—and there is no reason to believe in any other. There is no empirical evidence for such worlds—there could not be, since each world is spatiotemporally disconnected from every other, so there cannot be any kind of observation of any other world. We usually favour parsimonious theories over prodigious ones, but still, this consideration might not be decisive. Nor does an incredulous stare, as Lewis's view initially provoked (Lewis 1986a: 133), constitute an argument. The theory might sound incredible, but new theories often do. Nevertheless, if we could explain mere possibilities in a more down-to-earth way, such as with a theory that posits only a single world—our own—then we might have reason to prefer the latter theory.

5.2.2 Powers realism

The powers metaphysics claims to offer such an alternative. As in the case of modal realism, a powers theory is about more than just modality and certainly about more than just mere possibilities. Nevertheless, as in the previous case, I will concentrate on what the theory has to say about the subject at hand. I will also offer a version of the powers theory that reifies possibilities since the present concern is with how credible such a view could be. Accordingly, the powers view would be that the real possibilities are those that are nested within powers (Mumford 2004: ch. 10, Borghini and Williams 2008, Jacobs 2010) or, as some say, potentialities (Vetter 2015). A simple example illustrates the view. The dissolving of a sugar cube might be a mere possibility in the case where the sugar cube is undissolved. On this version of the powers view, however, the dissolving is real in the sense that it exists within the sugar cube in virtue of it having a power to dissolve. Powers are real, in this ontology, and the mark of a power is that it is *for* some further outcome. Those possible outcomes then exist within the powers. Something is soluble, for instance, only if the possibility of dissolving is contained within it.

This view could reasonably be extended to the other cases of real possibility. Rosa could have gone to London today, for example, because she has the power to do so. She has free will, for instance, and a body that is capable of making the journey whenever she really wants to. Going to London would have been the manifestation of the power, where this exists unmanifested within the power until it is called upon or stimulated to reveal itself. The term ‘manifestation’ has been adopted within the literature for precisely this connotation. Something is manifest when it has become apparent, appears into view, which implies that it could have been there previously but without being so apparent. The mere possibility of being in London is real, for Rosa, in the sense that it resides unmanifested in her powers, but it is there nonetheless, hidden and waiting to be revealed. It is perfectly consistent with powers theory, however, that such a real possibility is never revealed. Rosa might never visit London before she dies. It is true, nevertheless, that she could have done because during her life she had the power, containing that real possibility.

We will see that a powers view has a lot going for it. Among its theoretical appeal is the fact that it invokes only one world in order to explain possibility. Unlike modal realism, the mere possibilities have being in our own world, and we will later see that this is the basis of an objection to Lewis’s view. The possibility of dissolving, and the possibility of visiting London, are real within our world because they are parts of the powers that are held or instantiated within our world. By comparison, in modal realism what is a mere possibility for us has no reality in our world. Its reality is to be found only in some other world. In order to

say that something is possible here, we need to invoke the theoretical device of inter-world similarity and a counterpart relation (Lewis 1973: 91–5 and 1986a: ch. 4). Using this notion of similarity, it is possible that Rosa visited London today, even though she didn't, because in another world that is similar to ours, someone who is very similar to our Rosa (her counterpart) decided, on a day very similar to our today (the day's counterpart), to visit a city very similar to our London (its counterpart), and she arrived there successfully. Both theories offer reification of possibilities, therefore, but they differ over the question of at which worlds those possibilities are real. Quantitative extravagance is again an issue for modal realism here, however, since there will be many other worlds in which it is actual that a Rosa visits a London. There is a world in which Rosa does so where Hillary Clinton is president and another in which another Rosa does so where Sarah Palin is president, for example. Indeed, there are as many worlds in which a Rosa actually visits a London as there are other possibilities that are consistent with this being a fact, including whether this Rosa has a hundred thousand hairs on her head or a hundred thousand and one. Modal realism does not just reify the possibility in question, therefore. It reifies countlessly many versions of that possibility. The one possibility escalates into multiple tokens, each equally real.

It might seem that in pointing to these escalated possibilities, I am advocating the more parsimonious realism about mere possibilities that the powers theory offers. However, the matter is not quite as simple as that, since the powers view has a problem of its own. The sensitive question for the account is whether a power or potentiality actually contains its possible manifestation or does so only possibly. If the latter, then no progress has been made by invoking the powers theory. The issue is that the manifestation remains only a mere possibility until the time that it manifests. It would be a mistake to think of the manifestation as actual prior to its manifestation. The power could be real and actual, and the power could be for, or directed towards, a certain manifestation but still it seems most natural to understand the situation as one in which the manifestation has no being until it actually appears. This is apparent in the causal theory that comes from a powers metaphysics (Mumford and Anjum 2011) in which the power, when suitably situated, produces or causes its manifestation: a manifestation which appears only at that point. Given that the manifestation is produced by the power, then, it is clear that it did not exist prior to that occurrence. This is not to criticize the powers view per se, but it does undermine a simple claim that the powers view offers a reification of mere possibilities. The question of how a powers view can distinguish genuine from non-genuine possibilities remains open at this point, then. But if we are simply looking for reified accounts of mere possibilities, then this version of the powers view does not seem to provide any such thing.

Perhaps there is a different version of the powers theory that would succeed in reifying mere possibilities. Indeed, there is one powers view that might do so.

Marmodoro (2017) disagrees with the idea that a power causally produces and brings into being its manifestation, which would mean, as we have seen, that the possibility of the manifestation had no existence prior to that point. She thinks instead of the manifestation as 'the activated state of the very same power, and not the occurrence of a new power' (Marmodoro 2017: 57). The manifestation is in the power, therefore, but as an earlier, 'unactivated' state of it. The mere possibility of the manifestation at least has reality in this respect.

However, even this account turns out to be unsatisfactory as a reification of mere possibilities. For a start, as a theory of powers, Marmodoro's view is controversial. Powers are usually considered property-like, and it is typical of properties that they are identical in their instances: a One running through a Many. Her view offers an unorthodox model in which we can have the same power in two different states. This effectively allows change within a property, the difficulty of which is seen when we have the power in its unmanifested state and the power in its manifested state and have to accept them as the same power, even though the two states might look very different and might involve different modal features. For example: solubility in its unmanifested state contains solidity and the possibility of dissolving. But in its dissolved state, assuming it is entirely dissolved, there is no such possibility at all and no solidity, and yet these are supposed to be the same thing. It is hard to see that trope theory or even nominalist theories of properties could make much more sense of this view. Whether or not it is a coherent view, however, the difficulty of whether a possibility is real before it is manifested would resurface in any case. For, one would have thought, it is not until the power makes the transition from unmanifested state to manifested or activated state that the possibility in question attains actuality. It might only look as if the possibility resides in the power because of the theory's pretence of sameness but difference within the power.

5.3 Fictionalism

Perhaps, then, it is best not to reify mere possibilities and to stick with the initial intuition that they are nothing at all. It is sensible to say, as we have seen, that these mere possibilities become real only at the point when they are actualized. It was stipulated at the beginning that as soon as a possibility is actual then it ceases to be a mere possibility. There is thus no situation in which something is both a mere possibility and an actuality.

How, though, could we defend such a view? It seems that we are talking about something when we talk of mere possibilities, so what are we talking about? One such view (for example, Armstrong 1989: 49) is to say that the unactualized possibilities are mere fictions, which have no real being at all. The attraction of such a view is that it allows us all the usual possibility-talk but without any of the serious

ontological commitment that we have just explored. The non-actual traffic queue is no more than a fictional queue and Rosa's trip to London was a purely fictitious one since nothing in the world—the non-fictional world—constitutes any being of her London trip. On this view, we do not need to be fictionalists about possible worlds specifically. There is a view called modal fictionalism (Rosen 1990) that takes possible worlds to be fictional entities. Armstrong's fictionalism is not based on that but, rather, fictional recombinations of the existing elements of our world.

A mere possibility is thus like a story or part of a story. That the fictitious Oliver Twist ate porridge is granted, since the novel of that name states it, but this in no way requires any real being called Oliver Twist to have eaten any real porridge. Similarly, if I say that Rosa could have had croissants for breakfast, instead of porridge, it does not require that any person named Rosa ate any croissants, as modal realism would require. According to your theory of fiction, we are engaging in a pretence or make-believe (Hicks 2015); we are supposing that something were the case, as if Rosa ate croissants. We are imagining a case—and when one imagines something, one typically does not mistake it for reality.

What might also lend some credibility to fictionalism about mere possibilities is that it would explain how we are able to make all sorts of fanciful suppositions, such as counter-nomic possibilities, advanced modalized statements, and even counter-logical possibilities. If fiction is mere pretence, then it can explain how such language games occur without any ontological seriousness. For example, the earlier claim that Rosa cannot jump to the moon under her own power seems true only if we assume certain background conditions, for example, that the gravitational force exerted by Earth on her body is of a certain extent. But there is a sense in which it is still possible that she jumps to the moon if the laws of nature were different, such as if the gravitational force were much weaker than it is, if human beings could breathe in outer space, and so on. Her jumping to the moon in these circumstances is thus a counter-nomic possibility. Now it is hard to take seriously the reality of any counter-nomic situation, since the laws of nature apply at all times and places in our universe. But if counter-nomic possibilities are merely fictional contrivances, then they require no such being. This is not to suggest that the possibilities, so conceived, are of no use or interest, since fictions are often both. Second, advanced modal claims concern statements that cannot be made within the framework because they are about that framework (Divers 1999, 2002: 47–50, and Lewis 1986a: 6). For example, one might consider whether it's possible that there be an empty world. One might say either that it is, or that it's not possible for there to be two identical worlds. Or one might assert that there could have been a non-Hume world. If we are only entertaining fictions, then it needn't really matter if what we say isn't sayable in terms of only the worlds. We are merely considering a fictional case. In the third sort of example, we can even contemplate fictions that violate the laws of logic. Jenkins and Nolan (2012) suggest that we would be surprised if we encountered a round square. We can accept

that this is true even though a round square is geometrically impossible: indeed, it would surprise us precisely because it is impossible for there to be one. Similarly, one would be shocked to find a genuine exception to the law of non-contradiction. These cases involve ‘possibilities’ of some kind. If they are counter-logical possibilities, there would be serious concerns if that were to require any kind of being for them. If they are mere fictions, however, then we can consider and make claims about such possibilities with relatively little cost.

However, this case for fictionalism does not get a free pass. There are still some questions that demand adequate answers. As stated thus far, this form of fictionalism does not explain why some claims of mere possibility are true and some are false. If our theory were to be simply that mere possibilities were nothing but fictions, then there would be no constraints upon them: no grounding that tethers the possibilities to reality. Seemingly, one can go around proposing any fiction irresponsibly, but there seem to be at least some constraints on what we are prepared to say is possible and what we will say is not.

This might seem difficult to square with the aforementioned point that there can be counter-nomic and even counter-logical possibilities. But it can be squared. Claims of possibility are made within a framework of shared tacit assumptions dictated largely by context. When it is stated that Rosa cannot jump to the moon unaided, there is seldom confusion over whether this assumes the actual laws of nature or not. It can be taken as a default assumption that we are thinking within that framework. But the assumption is cancellable. There could be contexts in which it is clear and does not need stating that we are assuming a different set of laws of nature, or even different laws of logic. There will be various ways of making such contextual assumptions clear or leaving them implied.

5.4 Grounding possibility in what there is

We can return now to our concern that some claims of possibility seem true and others false. We might mark the difference by saying that some are genuine or real possibilities and some are not, where ‘real’ possibility need not be taken to mean real in the ontological sense of having being or existence. We have already seen that it is not easy to reify mere possibilities in that sense. What we need, to supplement our fictionalism, is an account of how some fictions are acceptable, and others are unacceptable. We believe that it is really possible that Rosa travels to London because that fiction is in some sense acceptable. The fiction that she jumps to the moon is in some sense unacceptable. How can we explain this difference?

What I will now propose is an account that satisfies all the requirements. It will satisfy the requirement of Parmenideanism because it will not reify any non-existence. It will satisfy the initial intuition that mere possibilities are nothing at all. Importantly, however, it will satisfy the requirement that, while mere possibilities

are nothing at all, since they exist only as fictions, nevertheless there is a distinction between genuine and spurious possibilities. This distinction, we will see, is grounded in what there is.

I will consider three proposals for how possibility can be grounded in reality. The first two proposals will be familiar: concrete other worlds and powers. A third—combinatorialism—will be introduced. Combinatorialism is combined with fictionalism in Armstrong's work. We will find that some mix of combinatorialism and the powers view is best for satisfying our requirements. This need not mean that it is the best theory of possibility absolutely. Those who are not Parmenideans might feel that they have no need to impose some of our constraints on their theory. If one is looking for a theory of possibility that seeks to avoid reification of non-being, however, then I am suggesting this option.

5.4.1 Modal realism again

First, let us look again at concrete other worlds, since we have not yet finished with them. Previously, they were considered insofar as they reified mere possibilities, though at an alleged cost of multiplying them unreasonably. We then considered the case for deflating such possibilities, treating them as mere fictions. The remaining problem, however, was how to explain why some possibilities are real (or true) and some are not. To solve this problem, the suggestion is that some possibilities are grounded in reality and others are not. What we need to do is explain what that grounding is. Could it be that the real possibilities are those grounded in the facts at all the concrete other worlds? In exploring this proposal, I am happy to set aside the lack of quantitative parsimony and consider whether the plurality of worlds would constitute an otherwise adequate grounding of possibility in reality. My conclusion, however, is that it does not. There are reasons to be sceptical about other worlds providing groundings for the possibilities.

We should note, in the first place, that there is no evidence for, or reason to believe in, the plurality of worlds other than their explanatory power (Lewis 1986a: 3). If we take them as a metaphysical postulate, then we can justify the postulation if it is explanatorily useful. There can be no empirical evidence for these world's existence, since we are spatiotemporally connected only to our own world. But the worry, then, is that possible worlds get the order of explanation the wrong way round, when it comes to what is possible.

Plato (*Euthyphro* 10a) wondered whether something was good because the Gods loved it or the Gods loved it because it was good. We can ask a similar Euthyphro-type question for possibility. Suppose that we could put every possibility into a one-to-one correspondence with truth in another world (though we will see shortly that this might not be the case). We can still wonder whether (a) *p* is possible because it is true in at least one world or, instead, that (b) it is true

that *p* in at least one other world because *p* is possible. Lewisians urge that (a) but there are some reasons to think that (b). Here are four reasons why (b) is to be preferred.

The first is that our judgements about what would hold at other worlds is dictated by our knowledge of what is the case in the present world (Richards 1975: 109–10). We cannot inspect any other world. We can ‘consider’ another ‘world’ but that is an act in our minds only and those considerations are bound to be informed by what we know already. We believe that it is not possible for Rosa to jump to the moon under her own power, for instance, and nor can any other human being. Within the possible worlds framework, this means that there exists no world (with the same law of nature as ours) in which someone jumps from the Earth to the moon. But how do we know that there is no such world? Isn’t that judgement formed on the basis of what we already know, from this world, of gravity, human anatomy, the size of the planet, the hostility of outer space, and so on? In contrast, when we judge that there are possible worlds in which Rosa visits London, our judgement is informed by what we know of our Rosa and London, human society, and travel connections: and it is then an easy judgement to make. How else could it work? How could I know, prior to any modal knowledge, whether there are other worlds, isolated from my own, in which someone who looks like Rosa visits somewhere that looks like London? I simply cannot. Lewis’s alleged explanatory power of the possible worlds framework then looks chimerical. It is far more likely that our modal beliefs explain—and perhaps even ground—our beliefs about possible worlds, rather than vice versa. Rather than being concrete existents, this would then make possible worlds look more like an expressive device: a way of articulating possibilities rather than either constituting or grounding them.

Second, even if there is a real, existent, concrete world corresponding to every possibility, of what relevance would it be to our world? In other words, it is not Rosa that travels to London. It is, at best, a counterpart of Rosa, who looks like her, who visits a counterpart city that looks like London. We wonder whether Rosa could have visited London. If we think she could, we are judging about Rosa herself, and London itself. We are saying that *she* could have gone to London. What relevance is it, then, that someone else, elsewhere, visits a city that is similar to London? This is known as the Humphrey objection, after the case of Hubert Humphrey wondering whether he could have been president (Kripke 1980: 45). He wanted that *he* be president, not someone else, so what matter is it to him that in another world, with a different set of circumstances, there is someone like him who is president? Similarly, suppose Rosa has an identical twin in the actual world who does visit London. This cannot constitute the fact that it is possible for Rosa to visit London even if it might inform our judgement of its possibility. For the same reason, it matters little to me that I have a counterpart in another world who murders his neighbours. According to modal realism, this is the case. Should

I have any care about this other Stephen Mumford? Should I be concerned about his pain when he is executed in the electric chair? I might have a worry about how close to our world is such a world where Stephen Mumford is a murderer. As stated above, however, any judgement I make as to how close that world is to ours is almost certainly formed on the basis of my knowledge of this world and myself, such as how often I have murderous thoughts (for clarification, never). This second problem for possible worlds suggests that even if they are real concrete existents, as accepted by a metaphysical postulate, then they still would not be able to do the explanatory work demanded of them. It simply does not seem that truths at other worlds connect, influence, dictate, decide, or constitute the real possibilities at our world. This, I suggest, gives us a reason to believe that our beliefs about such worlds are determined by our prior beliefs about possibility.

There is a third argument for this, too, which is that the truths of possibility are empirically discoverable and this cannot be explained by the existence of other worlds, since other possible worlds are not empirically accessible. There are good grounds for naturalism about the mere possibilities. For example, we can discover causal structures empirically, as well as causal laws, and these go a good way to telling us what is really possible. On the basis of empirical data, for example, we could formulate a mathematical function that related two variables, x and y . The formula might satisfy the known values of x and y . But suppose we consider some value for x that has not yet been measured. Using the formula, which has otherwise been well confirmed, we could be justified in assuming what the merely possible value would be for y , given x . Similarly, I could inspect some new mechanism with no experience of it in action but still be in a position to conclude what it is capable of doing: the, as yet, merely possible behaviour that the mechanism could produce. Perhaps some modal knowledge is intuitive, but these cases show that much of our modal knowledge is probably learnt. Modal realism provides no account of how our modal knowledge can be learnt, since it is impossible to investigate empirically, or perform any experimentation, in other possible worlds. Shalkowski and Bueno (2000) once issued a plea for a modal realist epistemology. The problem is that there can never really be one. This consideration certainly seems to apply to all natural or nomic possibility, even if other things would have to be said about logical possibility or legal possibility. Legal possibility is to be known through consulting the relevant laws and statutes, rather than other worlds. Logical possibility is not an empirical matter but nor, in judging whether something is contradictory or not, can I settle it by knowing whether or not it is true at some other world.

A fourth and final consideration for favouring interpretation (b) of the modal Euthyphro question is that, as we have already suggested, it seems that there are some possibilities for which there cannot be a possible world, or which we cannot articulate within the possible worlds framework. This consideration denies the coextension that leads to a Euthyphro question in the first place, but it is nevertheless worthwhile reviewing it here. We have encountered one of the instances

already, in chapter 3. It is possible that there was nothing. As we saw earlier, however (§3.8.4, above), there is no empty world, so we have a modal belief that we cannot faithfully express in terms of possible worlds. This is to accept the intuition that there could have been nothing but also that the possibility could not be expressed since, on the available conceptions, a ‘world’ containing nothing would not meet the basic requirements for being a world. A committed modal realist would see this as a basis for judging it to be necessarily false that there could have been nothing, or that it is necessary that there is something. We have already seen enough evidence to reject this conception of possibility, however. Instead, we can judge—perhaps using subtraction considerations, as in Baldwin (1996) and Rodriguez-Pereyra (1997)—that there could have been nothing, but this does not commit us to saying that there could have been an empty world or that the possibility of there being nothing resides in, or is equivalent to, the possibility of there being an empty world. Assuming, fairly, that there is no necessary existent, so each existent is a contingent being, then each might not have existed. Recalling an earlier analogy, each pebble can be lost from the collection, but it was said that not every pebble can leave the collection and it still be a collection. That is correct. Freed of commitment to possible worlds, however, nothing stops us allowing the possibility of all the pebbles being lost and there being no collection. It is possible that there be nothing even if there is no empty world. We have encountered other modal claims that cannot be accommodated within a possible worlds framework, for instance, the possibility of there being more than there actually is. The modal realist automatically rules against the possibility of there being more worlds than there are.

It would need a book-length study, at the very least, to resolve this issue to any specialist’s satisfaction. Our primary focus here is on non-being and the danger of its reification within various subject areas. We do what we can within the finite limits we have, and all our conclusions are provisional and defeasible. With the weight of argument we have been able to amass within those constraints, however, we will be right to judge that possible worlds provide a poor grounding for the mere possibilities. Instead, it looks as if our beliefs about possible worlds are grounded, rather, in prior modal knowledge. This still leaves the question of where that modal knowledge comes from. Speaking more ontologically, we are still right to wonder what constitutes the reality of some of the mere possibilities, if not possible worlds. It is here that a powers view becomes helpful once more, though we need not restrict ourselves to this metaphysical doctrine exclusively.

5.4.2 Natural groundings

The main point is that the truths of possibility may be grounded in this one reality that we know: the natural world of which we are a part. The reason powers gain a mention, in this respect, is that there are some conceptions of the natural world in

which useful modal knowledge would be restricted. We have seen Lewis's idea that the world is just a vast mosaic of particular matters of fact. From what we know of what there is, there would be no way of knowing what there could and couldn't be. There is just what is. As Hume argued, there are no rational grounds, within this framework, even on the basis of past experience, for saying that one thing is possible and another is not. There is no reason to think that the future will be like the past. Instead, Hume (1739: I, iii, 6) says that anything could follow anything else, so the possibility of Rosa jumping to the moon is no less real or unlikely than her visiting London. The powers view, though not the powers view alone, tells us instead that there are natural potentialities, capacities, and capabilities within things in the world, which make only some things possible. Rosa is capable of travelling to London, but she is not capable of jumping to the moon. Because powers are instantiated in particular things, they have a concrete existence. This existence is all this-worldly, and thus the objections held against possible worlds do not apply to powers. They are empirically accessible, for instance. Knowledge of them can be a basis for our judgements as to what is possible. I accept a metaphysics of powers, some of which can be primitive and irreducible, but all that matters for our account of mere possibilities is that they have some real-world grounding, for which other conceptions of powers could serve perfectly well. If we take Armstrong's (1983) view, for example, we have a world in which there are categorical, non-dispositional properties that can be related by a law of nature to other categorical properties. Hence, the property of being sugar can be related by a law to the property of being dissolved when placed in liquid. This would analyse the putative power of solubility in nomic terms but, if it worked as an account, it would be enough to ground the possibility of a particular sugar cube dissolving. In using the term power, I will remain neutral as to whether it is a primitive or analysable in other such terms. But it allows us to say that there is some such power behind everything that is possible, while there is no possibility of something if there exists no power for it (for a detailed demonstration, in terms of the similar notion of potentiality, see Vetter 2015).

It might be noted, however, that this account, while having the advantage of being naturalistic, suffers from a disadvantage because of that. It suffices for accounts of so-called natural possibility; that is, what is possible given the actual laws of nature, but it cannot apply to the unnatural cases of mere possibility, where we are assuming that different laws of nature hold or even that the laws of logic are otherwise. I could weigh a lot less than I do, for instance, if gravity were a weaker force than it is. I call these extended modal claims and, as noted earlier, there are some contexts in which possibility claims will be made against such a background. One might even learn something about the actual world if one considers how things would change if the laws of nature were different. For example, one might understand the effect of the actual gravitational force better if one considered the possibility in which it was weaker or stronger than it is.

There are various points that we can add, however, in order to provide an account of extended modal claims. For one thing, we can employ Armstrong's combinatorial principles. Armstrong generates possibilities by recombining the existing elements; not in actuality, of course, but fictionally. Hence, horses exist and the property of having a horn exists. These can be combined into the fiction of a horse with a horn, giving us the idea of a unicorn. Unicorns do not exist naturally, but under this fictional recombination, they are possible. The conversational context should signal that we are engaged in a combinatorial practice rather than saying that unicorns might have naturally evolved, though this latter claim could also be decided on combinatorial principles.

We can also make a fictional assumption about the laws of nature being different. For such a possibility to be coherently held, however, we still bring to bear much of our natural knowledge of this world. If absolutely anything goes, then we lose a grip over what possibilities we are stating. In assuming the law of gravitational attraction to be different, however, we will typically hold everything else the same. This corresponds to Lewis's (1986a: 20–7), move of considering the closest possible world in which something is true that is false in our actual world: don't make gratuitous assumptions about differences other than the one that is entertained and differences that have to follow from it. Hence, although I consider the fiction of things being radically different, as they would be if gravity were different, I still assume that I exist, the world exists, there are weighing machines, and so on. Furthermore, our knowledge of the effect of gravity is very well empirically grounded and so too, by extrapolation, would be our knowledge of how things would be if gravity were different. We can do the calculations, for example, and know what the attractive force would be between two masses that were a certain distance apart. Hence, even counter-nomic possibilities can remain tethered to the real world. There is no doubt that counter-logical possibilities take us even further from reality, since it is very fanciful to imagine a fiction in which logically or mathematically impossible things happened. But even here, where one might have to adopt a paraconsistent logic for purposes of entertaining the fiction, there is still some grounding in reality. To assume that I would be surprised if I saw a round square is to assume that a lot stays the same, such as my psychology: otherwise, why would I be surprised? Even in a logically impossible world, we still shouldn't presume that anything goes. Everything follows from a contradiction in classical logic, but when it comes to assuming a fiction, there are still constraints on what makes a coherent story.

5.5 Parmenidean possibility

We have covered much ground on the nature of possibility, so we should not forget what was at stake for our Parmenidean project. Our concern was with mere

possibilities, which were seemingly nothing at all, but nevertheless had some being. If so, this would be a threat to Parmenideanism. What we have been able to do, however, is maintain the view that a mere possibility has no actuality while also explaining how some possibilities can be real and some not. They can be real or true in the sense of being grounded in what there is without themselves having any being because they are fictions. Soft ontological Parmenideanism thus remains intact.

Nowhere is it more important, though, to be reminded that this is an account of possibility within a soft Parmenidean framework only. We saw in chapter 1 that Parmenides himself denied the possibility of change, and with that we must assume that he denied that there were any mere possibilities at all. A hard ontological Parmenidean would then have no conception of mere possibility. All possibility would be collapsed into actuality. We could add this to the reasons for preferring a soft Parmenideanism over a hard one.

It should also be made clear that this chapter is not offered as a complete theory of modality. Nothing has been said about necessities, for instance, and a host of related issues. The concern has only been to show what a soft Parmenidean can say about mere possibilities, allowing them no being while acknowledging that they can be grounded in what there is. Other issues relating to absences and nothingness have also been invoked. What is a fiction, and how a fiction can be true, are matters that will be revisited in chapter 8. An account of what to say about claims that are not grounded in what is, will emerge from chapter 9.

Perception of Absence

7.1 Seeing what is not there

To illustrate the idea of perceiving an absence, Sartre used the example of seeing Pierre is not in the café (Sartre 1943: 9f). I will offer my own example because it is likely to be of an even more striking kind. Jane parks her car on the outskirts of London. She catches the underground train and spends a day enjoying the sights. Later, she returns, knowing exactly where she had parked. She rounds the corner, her keys already in her hand, but the car is not there. She stops in her tracks, wide-eyed. It is a shock to her, but the car has gone, as she can plainly see. The realization seems to hit her right between the eyes.

Cases such as Jane's stolen car, or seeing Pierre is not in the café, are what we can call perceptions of absence. It might be better, here, to call them alleged perceptions of absence since whether they are genuine cases of perception, or of something else, is one of the questions we will have to resolve. Sartre's example was a bit less dramatic but involves the same philosophical problems. How can one see someone or something that is not there? Don't we only see things that are there? Perhaps we never do more than infer the absences of things or people from what we do see; so we only see *that* Pierre is absent, rather than seeing an absent Pierre. It is arguable, however, as I will show, that there is a distinctive phenomenology that accompanies the perception of some absences, which makes such cases hard to simply dismiss as inferential. The example of Jane's car draws our attention to such a phenomenology, which might be more pressing since it involves a valuable loss. Farennikova (2013: 430) uses the example of a missing laptop, which is similarly compelling. You are working in a café, you step outside for a break, and when you return you see that your laptop has gone from your table.

How one describes such an experience is less than obvious, however. It might be relatively simple to detail the experience of a presence, in contrast. For example, a red ball enters into your visual field from the left and rolls down the road. Did you see it? Certainly, you answer. What did it look like? 'Well, it was the size of a football and it was red. It was moving in a straight line down the slope of the hill. As it was a medium distance from me, it was pretty clear and distinct and I looked directly at it, my eyes followed its movement. It was the same colour as the jacket on the book I had been reading: my favourite colour,' and so on.

Now consider what an absence looks like. In some sense, the table on which there is a missing laptop looks just like a table upon which there is no missing

laptop. Indeed, another patron could enter the café at the same time as you and look at the same table. You know that there used to be a laptop there while the other customer does not. Do the two of you see the same—a table with a dirty coffee cup and salt and pepper pots on top—or do you see different things? What could it be that you are seeing and the other person is not? You see the absent laptop, you might insist, but what does that look like? Can you describe it as well as you could describe the presence of a red rolling ball? Isn't there something ineffable, to say the least, about perceptions of absence? Seeing an armchair, for example, doesn't 'feel' the same as seeing the armchair in which your recently deceased spouse used to sit. But it is by no means clear how such a difference should be described in purely phenomenological terms, even if we grant that the difference between these two experiences is real.

Apart from this uncertainty about the supposed phenomenological difference between seeing an absence and seeing the same scene without an absence, perception of absence raises other problems that we will need to consider. If we can understand perception of absence, however, there is every prospect of a significant gain. The gain will be twofold. In the first place, it will assist our soft Parmenidean project if the perception of absence can be explained in a way that contains no commitment to negative entities. Perception of absence presents a *prima facie* challenge to our project as 'perceive' is considered a success term: if you perceive a table then there must be a table that is perceived. If there is no such table, then you can only misperceive one. It seems to follow, then, that if you can perceive an absence, there must be an absence that is perceived. Does this require ontological commitment to a real absence? The second potential gain from an account of perception of absence is that it could explain how we have a concept of nothing in the first place, thus giving us an answer to Parmenides' (Piv) view that non-being is unknowable. The challenge of the classic empiricism of Locke and Hume was to show the origins in experience of the contents of the mind. There is a legitimate idea of a thing only if that idea can be traced back to some original sense impression. If we can explain how absences can be perceived, then potentially we might meet this challenge and be able to show how we can have an idea of nothingness. This second potential gain would be significant, since how else can we have a concept of nothing given that everything we otherwise know is a something?

Explaining perception of absence is not a simple matter, however, as I have already indicated. I propose, therefore, to spend a little more time outlining the target of our account. I will then look at two theories of how perception of absence occurs: the perceptual account and the cognitive account. I argue that neither of these is satisfactory since they show too great an adherence to a rational-intellectualist reconstruction of the phenomenon. Instead, I argue for absence perception as involving a non-rational mechanism. I invoke Daniel Dennett's multiple drafts model of consciousness as a useful heuristic in explaining some of

the mystery of perception of absence. Distinctive of such an account will be scepticism regarding the inner theatre of consciousness and also of a sharp division between sensation and propositional attitudes among mental contents. This will blur the traditional distinction often invoked between seeing and seeing *that*.

7.2 Our subject matter

There are a number of phenomena that are either related to perception of absence, close to it, or possible to confuse with it, so I will try to make clear exactly the subject matter under investigation.

First, and most obviously, we must distinguish perception of absence from absence of perception. Sorenson (2008: 17), for instance, is clear that we can see darkness, and that seeing an absence of light is not merely an absence of seeing. A failure to see something is different from our target case. Sartre could look around the café and not see Pierre even though Pierre is there. Sartre might even be told later that Pierre was in the café at the time he was looking and, hence, he could believe that Pierre was in the café even though he did not see him. Such a case is distinct from ours though still worth mentioning since it can have a feature that will be of interest to us later. Suppose you are asked to get the ketchup from the fridge. You open the door, have a good look inside, but then report that you can't see any ketchup. You might ask where it is. In response, the other person walks up and instantly reaches for the bottle of ketchup that was in plain sight. You might feel stupid in such a case, and be admonished, since a competent looker should have seen the object in question. Nevertheless, these cases of so-called inattentional blindness do occur. What is noteworthy about them is that it is perfectly possible for an object to be within one's visual field, presumably having an effect on one's retina, and yet you don't see it. There are other reasons why you might not see the ketchup bottle or Pierre in the café. The bottle was hidden. Pierre had visited the restroom at the moment you entered; or he hid from you behind a sofa since he wanted to avoid you. Those are relatively uninteresting examples of absence of perception, philosophically speaking. There can also be absences of perception in cases where you close your eyes or avoid the scene, perhaps if there has been a nasty accident. Inattentional blindness is where nothing is hidden, there is nothing to prevent you from seeing something, and you are looking towards the object, but nevertheless you do not see it.

It might also be, though not necessarily, that in cases where there is an absence of perception, there is an accompanying perception of absence, even though the two cases are distinct. You can be struck that something is not where you expected it to be, and then be surprised when somebody shows you that it was there all along, so it looks like both cases have occurred. It is not clear whether we should describe this as a misperception of absence, however, given that there was no




Figure 7.1 Kanizsa triangle

absence to perceive. But we will leave that matter as spoils for the victor, to be tackled (by someone else) only after we have given an account of perception of absence.

A second class of case we can distinguish are those of seeing what is not there: cases often dismissed as mere illusions. These could include, for instance, so-called ‘filling in’, where the mind completes some missing piece of information. This type of case shows us that seeing what is not there might at times be useful and perhaps provide even an adaptive advantage, a matter to which we will return later. There could be good, historically evolved reasons why the mind can work with an incomplete picture and still see things. It means that occasionally the mind can be tricked, as with the Kanizsa triangle in Figure 7.1, but at other times it is useful that we can see things with just a glimpse or partial view.

This gives us a distinction between seeing what is not there and seeing an absence. We have to be careful how we state such a distinction, since the two cases sound similar, but they are distinct. The distinction, we want to say, is something like this. In seeing what is not there, there is something that you see: the presence of a triangle, for instance. And you can still see this even when you know it is not really there. In seeing an absence, you don’t see the presence of anything. It is the absence of something—a person, a car, a ketchup bottle, a laptop—that is putatively seen. The absence of a person does not look like the person, for instance, hence it is extremely unlikely that the absence of something could be mistaken for that same thing’s presence. Nevertheless, it does seem possible to see both a thing and its absence, as cases such as Figure 7.1 illustrate. You can ‘see a triangle’ in the figure and see also that the triangle is absent. Whether seeing the absence of the triangle requires a closer inspection or some other such mental operation is not clear at this stage. When we have an account of perception of absence, we ought also to be able to explain cases such as this. What does one do in order to see the absent triangle that you are not doing when you see the triangle? Can one see both the triangle and its absence at the same time or must one oscillate between two modes of perception (and what would such ‘modes’ be)?

A third issue, in identifying our subject matter more precisely, is to make it clear that our topic is not merely limited to the perceptual faculty of sight. *Prima facie*, absences can be felt with a number of the sense faculties and these cases have just as good a claim to being perceptions of absence. Can't you hear silence (Sorensen 2008: ch. 14), for example? Suppose there is a burglar alarm going off in the night or some neighbours are having a party. You hear that the noise has stopped or you might just suddenly become aware that it has. In a piece of music there might be a pause, of a specific duration, when there is no sound. The length of the silence is dictated by its 'boundaries': the sounds on either side of it. It is perfectly plausible that this silence is a part of the music, since it would be a different piece without the pause. And it is also plausible that you can hear it, even though it is an absence of sound (at a public lecture at the BFI in London, I asked the Oscar-winning Hollywood sound engineer Randy Thom whether he considered silence a sound. He saw the philosophical problem that lurked within the question but accepted that silence was a part of his sound palate). In respect of touch, it seems that you could feel your way in the dark, stretching out with your hands, but feel that nothing is there. Sometimes there is surprise when this happens; for example, if you were walking upstairs in the dark and reached with your foot for a top step that wasn't there, probably leading to a stumble. Cavedon-Taylor (2017) gives the example of feeling with your tongue a missing tooth that has been extracted. There can also be olfactory absences (Roberts 2016); for example, if you smell that some perfume has worn off. Such a case need not mean that you smell nothing at all: only that a specific smell that you were checking for is no longer there. In respect of taste, you could notice that salt has not been added to a sauce, and in the case of proprioception, it could be noticed that a weight is absent. You go to lift a heavy case from a bed but are surprised to find that it is empty and that the lift is much easier than expected. I do not want to assume in advance that a single account of perception of absence can be given that covers cases from all the different sense faculties. That is a thesis that would have to be determined following consideration. It would be unfortunate if separate accounts were required for sight, hearing, tasting, and so on. Nevertheless, a unified theory of perception of absence is something that has to be won rather than assumed.

Fourth, I wish to set aside cases of simple logical consequence and concentrate solely on the alleged cases where one perceives an absence directly. The sorts of cases I am looking to exclude are such as, for instance, you see a square and, given that nothing can be both a square and a circle, in seeing a square you do not see a circle. Similarly, if you see something that is red all over, you are thereby seeing something that is not yellow all over. One might think that these are cases of seeing the absence of a circle and seeing the absence of yellow; but that is not a compelling interpretation. These examples are trivial, and I want to ignore them because, presumably, in seeing something that is all-red, you are seeing

something that is not yellow, whether or not it even crosses your mind that the object is not yellow. Something all-red is not yellow whether you perceive it or not, and whether you draw the inference from what you see or not. Our concern is with perception of absence where the absence is known experientially, not deductively. There might, of course, be genuine cases where absence of yellow is indeed experientially perceived, if perception of absence is eventually validated. For example, a snooker referee calls for a yellow ball but is handed one that is red. In such a case, it might strike the referee perceptually that the ball is not yellow. At least, the plausibility of such a case is still up for grabs.

This example raises a fifth issue, which is that the absences perceived can presumably be as varied as absences are generally, ranging over the sorts of diverse negative phenomena we have been considering so far in earlier chapters. It is one thing to perceive an absent object or person, but, as the snooker ball example shows, similar issues could arise with the absence of a certain property. The perception of such a thing might vindicate the notion of a negative property, which we considered in chapter 2. Regular properties are in principle perceivable, one might say, whereas negative properties, on account of their lack of unity, are not. This argument is challenged if the absence of a property can be perceived directly. Likewise, it is worth considering whether negative states of affairs could be perceived, and so on.

Sixth, we should also distinguish the perception of absence from the alleged case of seeing things that cannot exist. The claim has been made by Priest (1999) that we can see impossible objects, ones that depict contradictions, such as an object that has two prongs at one end but three prongs at the other, despite no apparent discontinuity in the middle, or an Escher-type staircase that is always going upwards but loops round to meet itself. These things cannot exist in reality and yet we are able to see them nevertheless, Priest argues. I separate these as cases where we see something—there is something there to be seen—but we interpret it as looking like something that could not exist. This is not quite the same as seeing an absence, although in the same paper Priest does allege that he ‘can see directly that something is not green’ (Priest 1999: 444). He does not explain how he does this, however.

With all these qualifications and distinctions made, we can understand our target phenomenon more precisely. We are considering, with all the above caveats, whether absences can be genuinely perceived in some similar kind of way to how presences of things are perceived. Alternately, is there some deflationary account available that explains that we never really perceive an absence? An account of perception of absence will also, presumably, explain what the difference is between perceiving an absence and perceiving a presence.

Before defending such an account, it is worth examining the two main theories of perception of absence. When we see the shortcomings of these theories, we will have a better idea of what is needed in a satisfactory alternative.

7.3 Perceptual theories

What I call a perceptual theory is one that accepts that absences can be perceived—known directly through experience—just as much as presences can be, as in, for example, Sorensen (2008). We can see, feel, smell, taste and hear absences of things rather than merely infer *that* there are absences. There is a phenomenological feel, for instance, which shows what it is to perceive an absence, as in the case of the stolen car or laptop. A perceptual theory has been offered recently by Farennikova (2013, 2015) and we will consider some of the detail of her theory shortly.

There are, however, challenges to a perceptual theory. I will detail them here so that we can then go on to see whether there are possible answers within some variant of the theory.

According to a perceptual theory, it is possible to see the absent dot in box B of Figure 7.2, as opposed to merely inferring the absence of a dot (Taylor 1952: 444–5). This is possible in the same way that you can see a missing car or an absent Pierre in the café. Seeing an absent dot in box B feels different from just seeing that box B is empty. Seeing that something is empty, unoccupied, or vacant might well, of course, itself count as a variety of perception of absence. But even if it is, seeing an absent dot in box B is, presumably, different from merely seeing that B is empty. The absence of a dot is a specific absence, and it seems to follow that seeing the absence of a dot would look different from seeing the absence of a cross, triangle, cow, and so on.

7.3.1 Phenomenal collapse

It is possible to use this line of thought, however, just as much to undermine perceptual theories of seeing absence as to support them. A sceptical response has bite. Is there really an absent dot in box B that you can see? Surely not. What you




Figure 7.2 Absent dot in box B?

can see is a box, or a square. At most you can see an empty square. Isn't the absence of a dot merely inferred from what is seen? An affirmative answer to this last question suggests a cognitive theory of absence perception, to which we turn in the next section. And doesn't the thought of a dot being absent occur to us only because of the suggestion made by the presence of a dot in box A? Suppose there had been no box A. Would you still then have seen that there was an absent dot in B? Or suppose that box A had contained a cross. Would you then say that you could see an absent cross in B instead of an absent dot? But box B would look the same whether A contained a dot or a cross, or anything else, or if there was no box A at all. It appears, then, that although there is some *prima facie* appeal to there being a notable phenomenology of perceiving absences, any such commitment we have to the perception of absence could melt away once we subject that phenomenology to sceptical scrutiny. Farennikova (2013: 432) calls this the problem of phenomenal collapse. We encountered it earlier with the idea that a tabletop and a tabletop where there used to be a laptop seem to look the same. They could be indistinguishable. And yet Farennikova still defends the view that the table is phenomenally different for someone who is looking for their missing laptop. Similarly, the chair and the chair where your deceased spouse used to sit seem different. There is an additional feel to the latter.

We should look in a little more detail at how Farennikova defends this view and offers a perceptual theory of absence perception. Her view is one with which I have some sympathy, but, I will argue, it suffers from the imposition of overly rational constraints. A sensible starting point is that we often have expectations when we make a perception (Farennikova 2013: 440). This would offer a clear way of understanding what is occurring in Figure 7.2. Seeing a dot in box A can give the perceiver an expectation in relation to box B. Hence, you can approach your perception of box B already with an idea of a dot in your mind, and the anticipation or expectation of a dot is then disappointed. One can see how this basic idea can also make sense of some of the other cases. Jane expects to find her car when she rounds the corner, and it is then a strange feeling, with its distinctive phenomenology, when that expectation is disappointed. Similarly, the regular place in which your deceased spouse used to sit will have many associations attached to it, and could produce the expectation of seeing someone, which makes their absence even more imposing on the senses.

7.3.2 The Mismatch Model

But how does such an expectation operate on perception and, more to the point, is it definitely a perception (of absence) that is produced rather than, for instance, a belief that accompanies the perception of the presences? After all, an expectation does not itself seem to be perceptual. It is a cognitive state that might

accompany a perception, but it is dubious that it is a perception itself. Farennikova (2015: 623) concedes this point but argues for the view that cognitive states such as expectations can penetrate perceptions. The expectation can thus change what is perceived. The idea that perceptions can be cognitively penetrated seems a sensible one so we can accept Farennikova's view thus far, but how does the expectation change the perception?

The mechanism Farennikova offers is what she calls a *Mismatch Model (MM)*:

MM: Visual experience of O's absence consists in an object-level mismatch between O's template generated by visual working memory and a percept of the observed stimulus. (2013: 444)

Hence, the expectation generates what Farennikova calls a template that is then matched against 'a percept of the observed stimulus'. The recent widower will have a template in which his deceased spouse is sitting in the chair, the laptop owner will have in their working memory a template in which their laptop rests atop the table or, to take a non-visual example, the man having a tooth extraction will already have a template of the feel of the inside of his mouth and all its current teeth. The strange feeling that we get when we perceive an absence is, on Farennikova's account, when this template mismatches the reality, which would be the current stimulus that is actually received.

How plausible is such a model and can it really provide us with a credible account in which absences are perceived? As I have already indicated, the phenomenology of these cases is not decisive, and I will explain why I think this is so in due course. There seems to be some difference between seeing a table and seeing a table with an empty laptop but it is not clear what exactly this difference is phenomenally, and the problem of phenomenal collapse exploits this uncertainty. However, when we examine the Mismatch Model, we can find a good reason to reject it, and this will be a verdict that could generalize to any other theory of its ilk.

Consider the following example from Farennikova, which she takes as an illustration of her theory:

For instance, if you are looking for a yellow cup in the kitchen, its template will activate the visual feature that will enable the most efficient search (e.g., color). The template of the object will then be projected and matched against expected locations of that object (e.g., the kitchen counter). If the object is not there, the visual system will establish a mismatch: a dissimilarity between the image of the sought object and what is observed. (Farennikova 2015: 628)

In particular, we should look at Farennikova's claim that the visual system can 'establish a mismatch'. The question to ask is how it does this. Here is the most

obvious explanation. One compares the template ('the image of the sought object') with 'what is observed'. If there is something in the template that is not in what is observed, then we have a mismatch and therefore, according to Farennikova's theory, a perception of absence.

There is a clear problem here, however. What the above account seems to rely on is that there is something in the template that is absent in the observation. You are supposed to compare two things and see that some element is missing from one of them, specifically from the observation rather than from the template. But this is meant to be an account of how perception of absence occurs. Supposedly you can see that something is absent from the observation (but which was in the template). But how does one do that? The very thing that is to be explained—how one can see an absence—is an essential part of the explanation, since the mismatch requires seeing something that is absent from the observation. Any theory that relies on us being able to see an absence in its putative explanation of what it is to see an absence clearly fails on grounds of circularity.

Could we salvage Farennikova's theory on the basis that it need not avail itself of a prior seeing of an absence? Her theory is in terms of a mismatch, which she describes above as a *dissimilarity* between the template and observation. The suspicion, however, is that the terms 'mismatch' and 'dissimilarity' are being used to cover up the fact that what the theory actually requires is that one sees an absence in the observation of something present in the template. To grasp this, we can note that it would not count as a perception of absence unless the mismatch went in a specific direction. It would be a mismatch, after all, if there was something in the observation that was not in the template. You walk back into the café, equipped with a template for the table on which your laptop sits, but then see that a cushion has been placed on the table as well. Here, there is something in the observation that was not in the template; or, to put it another way, there is something absent in the template that was in the observation. But this is not a perception of absence. For that to occur, it must be the other way around. The mismatch must have the specific direction where there is something in the template that is absent from the observation. Use of words such as mismatch and dissimilarity cannot disguise the fact, therefore, that what the theory needs is that an absence be detected perceptually. The requisite mismatch and dissimilarity must consist precisely in this.

I cannot see how the theory can be saved from this objection. It might seem that I am being unduly harsh with Farennikova and it will seem even more so once I have presented my own explanation of perception of absence since it might be thought to bear similarities with hers. Our views will vary in some specifics, however, which is why I have had to consider the detail of her mismatch theory. My own account will differ from hers to the extent that it does not attempt to impose rational and normative constraints on perception of absence. The preferred view is one in which perception by absence is the result of a non-rational, evolved mechanism.

7.3.3 A causal theory of perception

There is a further objection that applies to any theory that affirms the possibility of absences being perceived. The objection arises from a causal theory of perception. Such a theory, a version of which I hold, maintains that to perceive something is to enter into a causal relation with it (for example, Grice 1961, Strawson 1974, Pears 1976). To qualify as a genuine perception, the object perceived must be the cause of the perception. The most plausible such theory would hold only that it is a necessary condition that the object has caused the perception, not a sufficient condition. A dose of LSD might cause a hallucination of a banana, but this will not qualify as a perception of a banana, nor a perception of the LSD. Hence, other constraints might need to be added, such as that the perception faithfully resembles its cause. But I need not consider any such additional conditions here since it is the causal condition on perception that perception of absence would already struggle to meet. Why is it even a necessary condition on perception? Suppose someone dreamed, imagined, or hallucinated a hand in front of their face. None of these would count as the perception of a hand, since there was no hand in front of the person's face that caused the dream, image, or hallucination.

Causal theories of perception are not problem-free (for instance, deviant causal chains: what if the cause of you dreaming about a hand in front of your face is the fact that you are lying in a strange position with your hand in front of your face?), but a causal condition is a powerful and intuitive constraint on what can rightly count as a perception. You succeed in perceiving a table partly because light bounces off the surface of the table and hits your retina. Interaction with the object has thereby caused a change in your body. Similarly, if you feel for a table in the dark, you perceive it when it touches your body and impresses itself upon you. Similar accounts can be given for the other senses, where perception always seems to reside in the object perceived causing a change in your body. More complicated stories will have to be told for perceiving things through microscopes, oscilloscopes, fMRI imaging, feeling the resistance of a notepad with the end of a pen, and the need for the perceived object to contrast with its background (Sorensen 2008: 54, 55), but these should nevertheless come out as causal explanations even if the causation involves intermediate steps and accompaniments.

Given the above account, it should be clear why absences are incapable of satisfying this condition upon perception. Light cannot bounce off an absence and strike your retina, an absence cannot press itself into your body, absent salt cannot affect a change on your tongue, nor can an absent sound beat upon your ear drum. The objects supposedly perceived cannot cause changes in the perceiver.

Now a believer in the possibility of perception by absence might, of course, also accept the possibility of causation by absence. If absences can be causes, then perhaps an absent laptop, which used to be on a table, can, after all, have an effect

on a viewer's retina. Apart from the concerns we raised in chapter 4, above, however, perceptual cases bring home the sheer empirical implausibility of causation by absence. For how could an absent object cause a change to one's retina? The absent object cannot reflect any light, for instance, and reflected light is the usual way in which perceived objects affect the retina. Similarly, a present object can impress itself on my body, giving me a way of perceiving it, but clearly an absent object cannot press against my body. And flavours in the sauce can affect my tongue but surely the absent salt cannot. It does not seem, therefore, that there is an easy way of redeeming perception by absence simply by upholding causation by absence, since this latter notion is at least as problematic as that which it is meant to redeem. This should really come as no surprise, since it might strike one immediately that perception of absence is a variety of causation by absence.

The causal constraint on perception affects other accounts of perception of absence too. Martin and Dokic (2013: 121), for instance, oppose Farennikova's view that absences can be perceived, but in the metacognitive account that they offer as an alternative, in apparent cases of perception of absence there is a metacognitive state, labelled a feeling of surprise (FoS), that accompanies the things that are perceived (and for which a standard account of perception suffices). Hence, you round the corner and see bare asphalt where your car used to be. What you perceive is the asphalt, but the perception is accompanied by the metacognitive feeling of surprise, since you were expecting to see your car there instead. Again, this explanation has its attractions, but there remains a problem of causation by absence even with this. If it is problematic to account for how an absence can cause a perception, is it not equally problematic to understand how an absence can cause a metacognitive feeling of surprise? In offering any alternative account of perception of absence, therefore, we had better ensure that it does not at any stage require acceptance of causation by absence.

There is one last attempt that might be made to salvage a perceptual theory of perception by absence, which is that there are some cases where absences can be seen not because of any change that they cause but because of the holes they leave behind and which can affect the senses. Let us start with the perception of an actual hole and see if such an account can be generalized to other cases. One can see a hole punched in a sheet of paper, for instance, where the hole could be understood as the place where the paper is absent. Here, one sees the hole, supposedly, not by the hole itself leaving an impression on the senses but by the effect on the senses of the paper that surrounds the hole. If this is a valid case of perceiving an absence then perhaps it can be generalized to other cases of perception of absence.

There are doubts that this is possible, however. As we saw in chapter 3, there are some reasons for accepting holes as real things and not merely absences. In particular, their existence was parasitic upon that of their hosts (Casati and Varzi 1994), and their hosts create a place at which there is an enclosed boundary.

This coincides with the boundary of the hole, and it is essential for a hole that there is such a boundary. An empty area of space is not sufficient for a hole unless there is something that bounds that area. This constraint leads us to say two things that cast doubt on this as a possible explanation of perception of absence. First, there are grounds on which to maintain that, in seeing a hole, what you are actually seeing is its boundaries, upon which it depends, and which are made out of real stuff, perceptible in the normal way. Hence, no distinct perception of an absence is required in order to see a hole. Second, there are no good reasons to assume that we could generalize from the case of holes to all cases of perception of absence. The missing laptop on the table has no perceptible boundary, like a hole does. The laptop's being taken away does not create a laptop-shaped gap in reality whose boundaries you can see. Absences would be much easier to see if that were the case. But if the theft of my car left a car-shaped boundary on the road then presumably anyone else would be able to see it and not just me. They can't. Similarly, the missing dot from box B, above, has no boundary. There is nothing at all in box B. And the extra eleventh planet in our solar system, which does not exist, is not seen by seeing anything around it. Its absence leaves no trace. So there are reasons to believe that this sort of model offers no additional hope to a perceptual theory of perception of absence.

7.4 Cognitive theories

Given the difficulties detailed above, in relation to perceptual theories, a tempting thought is to say that perception of absence does not involve absences being straightforwardly seen. Rather, according to the cognitive view, absences are effectively inferred or deduced from what is seen. You cannot perceive something that is not there, but you can certainly have negative beliefs, and these can be derived from what is seen. In other words, you cannot see the absent laptop, but perhaps you can see *that* a laptop is absent; that is, you can hold a belief that the laptop has gone based on what you have seen. 'The Cognitive Account' is Farennikova's (2013: 434) term and, although she does not attribute it to anyone directly, she notes it as implied by the perceptual theories of Gibson (1966) and Dretske (1969).

What is meant by a negative belief is the topic of chapter 11, below. Although, as we shall see, this is no simple matter, what we mean approximately by a negative belief is a belief *not*: that such-and-such is not the case, for example, that it is not that case that *a* is F or that *b* is not-G. We cannot, however, adopt a solely logico-linguistic criterion of what is a negative belief since negative beliefs can be expressed without using the explicit word 'not', such as when I believe that my drink is sugar-free.

A possible theory of what is really going on with putative cases of perception by absence, then, is that various things are perceived—a table top, a coffee cup, a

floor, a chair, a room, and so on—and from them, you conclude that your laptop is absent. For such a conclusion to be drawn, there might have to be a background belief that there was a laptop on the table when you left the room, but this will be acceptable on a cognitive account since other beliefs in addition to what is perceived could be responsible for a further belief. In the cases in question, the initial belief seems required, since another person looking at the same scene might not ‘see’ that the laptop is missing, by which we mean, on a cognitive theory, they might not infer the absence of the laptop from what is seen. Clearly a contextual cue is needed, in addition to what is seen, in order to infer an absence. The contextual cue could be a belief or a desire, such as if someone wants a particular book that they believe to be on their bookshelf but then fails to find it. After looking at every shelf, they might conclude that the book is not there. Have they then seen the absence of the book? Not according to a cognitive theory: they have merely judged from the books that they have seen that the one they wanted is not among them. Similarly, when you expect to meet Pierre, or anyone else, in the café, you do not really see his absence. You see various other people and furniture in the café. You have as good a look around as you feel you need in order to establish whether he is there. And if you don’t see him, you draw the conclusion that he is not. An absence of perception is not a perception of absence, but you are certainly at liberty to draw conclusions from what you don’t see, which is the thought behind a cognitive account. This is clearly not the same as seeing an absent thing. When someone says, ‘the silence was deafening’, no one really thinks this means there was a loud, deafening sound. It means that something could be read into that silence: it was awkward; someone should have said something but didn’t.

The cognitive account has an advantage over a perceptual theory in that it does not require a problematic theory of absent things being literally seen. It does not require the perception of absences in the usual sense of perceive that we apply when we see presences. But there are two major problems with a purely cognitive account, nevertheless.

7.4.1 The phenomenology of deduction

The first major problem is that a pure cognitive account does not respect the phenomenology. Perception of absence feels like something whereas holding a belief does not. Some time was spent at the start of this chapter explaining what that feeling might be. The feeling of perceiving an absence might be ineffable; nevertheless it feels like something when you see that your car or laptop is missing, or that the chair of your late spouse is unoccupied. Although we are yet to describe that feeling adequately—and we might be best reserving judgement on the matter, in any case—we can nevertheless offer some explanation of the force of this

argument. Mental contents are often divided into two kinds: sensations and propositional attitudes. We can see that the perceptual theory explains our target phenomenon in terms of sensations while the cognitivist theories explain it in terms of propositional attitudes. One important difference between these two types of mental contents is that the former are occurrent while the latter are dispositional. As traditionally conceived, a sensation has a dated beginning, a duration, and an end. You might have to explain when your toothache began and how long it lasted. And once you no longer feel it, then it is deemed to have gone. Indeed, it is only there, and only real, as long as it can be felt. Beliefs are not like that. They do not have to be felt in order to be held; indeed, in standard cases, they feel like nothing at all. Because of this, they are not considered to be occurrences. Of course, there might be a time at which a belief is acquired, and some beliefs are given up, but a belief can be held over a relatively long time regardless of whether it is present before the mind. Consider your belief that the Earth is spherical, which has no doubt been held continuously from an early age and enters only rarely into your consciousness. If we understand the cognitive account as one in which the perception of absence is reduced to the holding of a belief, then it is not in a strong position to account for the alleged phenomenology, simply because beliefs have no feel.

It might be said, in objection to this claim, that some beliefs are indeed capable of feeling like something. Perhaps a strongly held religious belief, for example, is accompanied by a feeling. And a supporter of a cognitive theory of perception of absence might argue that beliefs that something is absent could be among the counterexamples to a general claim about beliefs having no feel. But there are a few things we could say in reply.

First, simply asserting that these negative beliefs have a feeling is not adequate. What we need is a plausible account of how a belief could feel like something and, in particular, how this kind of belief—that something is not the case—can feel like something. How is this possible at all?

Second, we can point out that there is a prior distinction that is usually made between believing something and knowing it experientially and the supporter of the cognitivist view ought to explain why this standard distinction does not straightforwardly apply to the case of perception by absence. Cavedon-Taylor (2017) gives the example of when a dentist removes a tooth under local anaesthetic. You are conscious throughout the procedure and immediately acquire the belief that the tooth is no longer there. Perhaps you see it being taken out. But it is only once the anaesthetic wears off can you actually feel, with your tongue, that it is absent. There seems no special reason, then, why this distinction, between believing something and experiencing it, shouldn't apply to perception by absence. I can believe that Pierre is not in the room because I have been told he is not, or I can visit the room and look for myself. The conclusion to be drawn from this is that cognitive accounts are not really explaining perception of absence

since they simply ignore what is important about it and describe a related but different doxastic state instead.

Third, in response to any cognitivist's claim that beliefs about absence can feel like something, we can say that the claim fails to establish what it needs: that beliefs have a feeling. At most, we could conclude that some beliefs are accompanied by a feeling, but this is consistent with Martin and Dokic's metacognitive view. The feeling of perceiving an absence is a metacognitive state, they say—not merely the holding of a belief but a feeling of accompanying surprise. This is consistent with the view that beliefs feel like nothing. Anyone who wanted to insist that it is certainly the belief that we feel, instead of an accompanying metacognitive state, faces a burden of proof. (For the record, Martin and Dokic present their metacognitivist view as opposed to both the perceptual and cognitivist theories.)

There is one further aspect to the objection from phenomenology. The criticism to the cognitivist account so far has been based upon the idea that there is a phenomenology to the perception of absence that a cognitivist account misses. But another concern is that typically there is no apparent consciousness of drawing any kind of inference, and this will relate to the next topic considered. The basic idea of the cognitivist view is that absences are not immediately perceived but only inferred from what is perceived. One worry, then, is the possible objection that we have no awareness of any such process of inference occurring. It is not an inference that we consciously make, in many such instances. A belief might not feel like anything, but an inference is something that we can experience, for example, when performing a mathematical calculation, such as completing a Sudoku puzzle. We can be perfectly aware that we are inferring and deducing. Yet in some of the most striking examples of perception of absence, there is no such awareness. Among other things, there does not even seem to be time for such an inference. You can see straight away that your car has gone, feel immediately that a tooth is missing, or feel instantly that there is nothing in front of you in the dark. Now it might be objected that perhaps there is some very fast subconscious inference that is being drawn and is responsible for the feeling of a perception of absence. On such a view, I will not make a judgement here, since it bears some similarity to the account ultimately I shall give. Before that, however, there is a further consideration that is relevant.

7.4.2 Deductive validity?

I said that the cognitivist account faced two major problems. The first was that it did not respect the phenomenology, and that case has just been made. The second problem concerns the nature of any such inference. Specifically, there is no deductively valid inference to be drawn from what is perceived to what is not there, unless some further negative is invoked.

This point can be understood purely logically or intuitively. I begin with the intuitive. You look around the Parisian café. You see chairs, Paul, white tablecloths, a waiter, windows and drapes, chandeliers, and so on. The perception of these things cannot conclusively determine that Pierre is not there. He might be hiding behind a sofa, seeking to avoid you, or perhaps he is wearing a new hat and you don't recognize him, or he was stood behind you as you looked. The inference to his absence is not a secure one, therefore. Seeing all these things does not mean for certain that Pierre is not there. This seems to undermine the cognitivist view that seeing he is not there consists in such an inference since no such inference is strictly warranted. I have a theory, which I can in no way prove, that a famous dispute between Russell and Wittgenstein was over this very point (see Russell's letter of 1st November 1911, in McGuinness 1988: 89, and also Russell's obituary for Wittgenstein, in Russell 1951: 297). Wittgenstein stubbornly refused to accept that there was not a hippopotamus in the room, no matter how much of the room Russell showed him. While it is easy to dismiss this as characteristic stubbornness on Wittgenstein's part, we should also note that strictly he was right to decline any such inference.

It might then be no coincidence that Russell (1918: 207) subsequently made play of that very same point when he formulated the logical version of the argument: one that we already encountered in §3.2, above. The logical argument says that what there is cannot entail what there is not, except for some special cases that we will consider shortly. Russell realized that in addition to atomic facts, higher-order general facts would be needed to account for general truths. We can now present the Russellian view in a way that is adapted for the case of perception of absence. The presence in a room of a table, a chair, a man, a coffee cup, and so on, cannot entail that another thing, such as a laptop, is absent, since the presence of those things is consistent with the laptop also being there. The existence of those things does not exclude the existence of another thing, therefore there is no logical warrant for inferring the absence of that thing merely from the presence of those other things perceived.

Now it might be objected that there are sometimes good grounds for saying that the table, chair, man, cup, and so on, are all the things that are in the room. And if these are all the things, and the laptop is not among them, then we can infer from the presence of them to the absence of the laptop. However, this is the issue that Russell addressed directly. His point, as we saw in §3.2, is that what is in the room cannot entail that it is everything in the room. For it to do so, it would have to exclude the existence of any extra thing, which it cannot do. Indeed, there may have been times when those things were in the room plus something else, so the current things are consistent with the presence of some further thing. This argument applies even to the whole universe: what there is cannot entail that it is all that there is, since it is possible for there to be further things.

There is no deductive path from what is seen to what is absent, unless a general or totality fact is added. This is problematic for a cognitivist account. The additional required fact is negative: that there are *no more facts*. Even if this does succeed in securing the requisite deduction, the price is too high. A general or totality fact is another negative and how is this general absence known? It seems even harder to account for this than the original absence it is meant to explain. Instead of having to see an absence of one thing, we now have a task of seeing an absence of many things; indeed an absence of any thing. To know such a fact experientially would then amount to seeing the absence of an infinite number of things. As this would include the absence of the first thing in question, then this task is at least as difficult as the one with which we started.

No doubt there are several other arguments that we could consider, if we wanted to cover every issue related to our investigation into perception of absence. However, the negative part of the study can stop here, not because it is complete but because enough has been said to show the difficulties attached to the subject and enough points have been made to motivate the need for a better view. We can now attempt an account that avoids some of the claims that have generated the difficulties. Foremost of all of these, afflicting both the perceptual and cognitivist accounts discussed, is an over-reliance on a rational explanation. Perception of absence occurs, I claim, not because of any rational process, following from expectation or beliefs. Rather, it occurs due to an evolved, non-deductive mechanism and an illusion, albeit a useful one, of the mind.

7.5 A Dennettian account

I call my account Dennettian not because Dennett has offered a theory like this. Indeed, he does not discuss explicitly the problem that concerns us but only a few related cases in passing. My account is Dennettian in the sense that it rejects use of the Cartesian inner theatre of consciousness and it accepts certain elements of Dennett's multiple drafts model to explain the apparent phenomenology of absences.

In order to understand the positive account, however, let us just pause and consider some of the constraints on an acceptable explanation of perception of absence that come out of the above discussions. In sum, we saw that:

- a. There is a causal constraint on a theory of perception: a perception has to be caused by the object perceived. But there is no causation by absence.
- b. No informative account can rely on a prior perception of absence, such as if some expected element is missing (i.e. absent) from an observation.
- c. Similarly, we should avoid explaining perception of absence in terms of any item, such as a feeling of surprise, that would only exist once an absence had been perceived.

- d. There is no valid deduction of what is absent from what is present, even when what is present is seen.
- e. There is a phenomenology to the perception of absence that must be respected, even though the nature of this phenomenology is uncertain and elusive. There is a difference between knowing something to be true and knowing it through experience. The cases we seek to explain are when something is known experientially.

In order to give an account that satisfies all these constraints, I will explain perception of absence as a variety of user illusion that arises from an evolved mechanism of non-deductive but overall reliable inference that is experienced as having phenomenal character. Deploying Dennett's multiple drafts model, I will also mention a distinctive temporal feature of this experience, which justifies why I classify it as an illusion, though it is an adaptively advantageous one. I see no reason why such an account could not be extended to human infants and animals, who seem to have awareness of absences even though they are thought to have relatively limited intellectual capabilities.

First, let us consider the causal constraint again. A perception is a result of a causal interaction; but we must also consider what is caused and how. A sense data theory tells us that the impression of external objects on our senses causes some inner representation or sense datum and a succession of such sense data is what we experience. Sense data theories seemed dead for some time (Austin 1962), but they have experienced a surprising revival in recent decades (for example, Chalmers 1996). A notorious problem with them is that they do not really explain how perception occurs because the question of how external objects could be perceived directly has only been swapped for a question of how internal objects can be perceived directly. A line of objection more pertinent to our case, however, is that sense data theories do not account for the way in which experiences are cognitively laden or mediated. What we believe can affect what we experience. Dennett gave a number of examples of which sense data theory struggles to make sense. For example, you probably hated the taste of your first sip of beer (Dennett 1991: 395f). Now you like the taste. What has happened? Does beer taste the same to you as that first sip but you now like that taste? Or should we say that the first taste really was horrible and that it now tastes better? If sense data were real, then there ought to be a definite answer to that question. But there isn't. Similarly, I have elsewhere (Mumford 2011: 57) given the example of two people looking at the same object where one believes it is a work of art and she has an aesthetic experience while the other thinks that it's a regular fire extinguisher and duly has no aesthetic experience. These two people apparently receive the same stimuli and, let us suppose, have identical retinal images. But is there a very real sense in which they see different things? Is this because they experience different sense data? That question seems unanswerable. What they see is, sort of, the

same; and yet they see it in different ways and have different experiences. It also seems implausible to say that they experience the same but merely have different beliefs about what they have experienced. There is no way we can know that. I give the above arguments as a way of motivating the search for a better alternative; I am not claiming that they definitively put paid to sense data theory.

Something is caused in perception: an experience. But we should not just think of causation in a simple, linear stimulus and response way. Admittedly, many theories of causation are like that: successions of causes and distinct effects, often represented with neuron diagrams (against which, see Hitchcock 2006). C. B. Martin (2008: ch. 5) offered an alternative, though: the mutual manifestation model. Here, an effect is produced that is the mutual manifestation of various empowered disposition partners. In the case of perception, those partners will be some object of perception and the perceiver, complete with a psychology. A perception is an action rather than a passion; that is, we are active with respect to it rather than passive recipients of sense data. We can determine what we see—what we attend to in a visual field—whether that be through conscious choice or from the evolved non-intellectual mechanisms we bring to the perceptual process. We act upon the object perceived and the experience comes only at the end of this interaction between object and perceiver. This is how there is a possibility of cognitive mediation upon our experiences. A mutual manifestation account does not take us to subjectivism, however. Worldly objects of perception have their own contribution to make: they will have their own dispositions or tendencies to be seen a certain way (Anjum and Mumford 2018: ch. 7). But nor does this account support a naïve realism since experience can to an extent be perceiver- and theory-dependent too (Hanson 1958).

This understanding of causation is important for what is to come. As Dennett shows, the connection between sensory inputs or stimuli and experiences is not a simple one. In particular, an experienced order of events need not match the order in which the stimuli were received (Dennett 1991: 115–26). The experience is an outcome of a sometimes messy causal process in which the perceiver's own faculties play a role. In this case, I claim, the relevant faculty is an evolved, unconscious, non-deductive mechanism—a causal rather than normative one—that creates an ineffable experience of absence. This also allows us to accept a feature of Farennikova's (2015: 623) account, that in some instances of perception of absence, cognition penetrates experience. Either expectation, or unconsciously held templates, can influence the feel of what is perceived and become a part of the experience.

Let us consider the point, then, that the relevant cognitive mechanism is causal rather than normative. By this, I mean that it has more in common with a reflex response than completing a Sudoku puzzle. There is no deliberate, conscious, and rational inference drawn from what is seen to what is absent. We have already detailed how there is no support for any such deductively valid inference and nor

is there always an experience of reasoning occurring. However, what this does not exclude is an automatic and direct non-deductive inference drawn from what is seen to what is absent. Furthermore, it is easy to see how such a direct inference would be selected in evolution if it is overall reliable, by which I mean that it tends to get things right. Both Farennikova and Dennett (2017) use evolutionary arguments to explain the mechanisms of the mind, and it is easy to see how they apply to the case of perception of absence, there being an adaptive advantage in knowing what is not there as much as in knowing what is there (Farennikova 2013: 425). For example, seeing that there are no predators or dangers present allows you to feed, while putting it even more bluntly, seeing that there is no solid ground ahead of you prevents you from stepping off a cliff. The latter example is, of course, one that can be felt tactually as well as seen, hence there is no reason why this non-deductive inferential mechanism should not work on all the sense modalities and not just vision.

I classified this non-deductive mechanism as causal rather than normative. What does this mean and why is the inference not normative if it is reliable? The answer to this is that the inference is one over which we have no deliberate control; one that we only know about through its effect on our experience and of which we have no consciousness until it is complete. It is true that we can sometimes deliberately and consciously draw non-deductive inferences that we know to be reliable to a degree and that rational norms can guide us in this. For example, a philosopher might draw an inductive inference from some limited cases that have been experienced to the nature of the next case to be encountered. The philosopher knows that the inference is not deductively valid, but can choose to draw it nevertheless, guided by some non-deductive set of norms concerning when it is useful to draw such an inference. These norms could be along the lines Bacon (1620) described: look for a wide variety of cases, at different times and places, with no counter-instances; and then an inductive inference is more reliable than one, for example, drawn from a solitary instance. But in perception of absence, there is no such normatively guided process that we go through. Perceptions of absence are immediate. One could choose to look for Pierre, certainly, but the experience of his absence is not something one chooses but is dependent upon what you see and don't see once you have chosen to look. It is possible, of course, to be very deliberate in looking for Pierre. One could search all around the café, looking at every single thing, and then decide to employ inductive norms before concluding that Pierre is not there. But this is not the perception of absence that is our subject matter: the one that has a pressing phenomenology. As we said, there is a difference between believing something and experiencing it. The former might be a normative matter in the case of absence, but perception of absence is a case of the latter and this is where the mechanism seems purely causal and non-normative. Think of the immediacy, again, of rounding a corner to find that your car or laptop is gone.

However, it might be objected that even if the inference is immediate and non-deductive, that this still sounds like a cognitive account. What about the phenomenological aspect of perception of absence, which we said we should respect? How is that explained if we speak in terms of a non-deductive inference?

It is here that we may invoke Dennett's account of consciousness and the notion of the user illusion. The sort of illusion that is in mind is like the user illusion we have when we operate a computer. The computer's workings are understood by programmers in terms of machine code and physically are a system of on-off switches. They also have a user-friendly interface for those of us who understand none of what lies beneath. Dennett's suggestion is that there is something similar occurring with the minds that we use in order to navigate our way around the world. In his words:

The brain's task is to guide the body it controls through a world of shifting conditions and sudden surprises, so it must gather information about the world and use it swiftly to 'produce future'—to extract anticipations in order to stay one step ahead of disaster. (Dennett 2017: 144)

This quotation emphasizes an important truth about perception. It is not there just for us to enjoy the sensations: to bask in the phenomenology. It is there for a reason: to guide us through a potentially hazardous world. In order to do so, it must guide us quickly. A mechanism that get things right 100 per cent of the time, but taking an hour to do so, is far less useful than a mechanism that can very quickly get things right almost all of the time.

With that thought in mind, we can return to perception of absence. How will such a mechanism, if it is to provide an adaptive advantage, do so quickly enough? Here is where we can argue that presenting the information experientially, as a perception of absence, wins out over other alternatives. And yet this must be an illusion. What is not there cannot be seen. Only presences—what is there—can be seen. Yet absences seem experientially accessible. They have a phenomenology, even if an elusive one. And how could it be otherwise? As urged above, it is not as if there are laptop- or car-shaped holes in reality that *could* be seen. In the instances that concern us, there is nothing there. The phenomenology of these cases could not be as clear and definitive as seeing a deep-blue sky, tasting cinnamon, or smelling a rose. If this account is right, the experience of an absence is just how a hastily grasped belief presents itself to us. It is advantageous for such beliefs to be experienced as such, if it can stop us from walking off a cliff, for example. The haste can produce mistakes, though, of course. We can recall the case of looking in the fridge for ketchup and not seeing it even though it is right there in front of us.

Were there not, however, some serious concerns afflicting this sort of account when others presented it? In discussing Farennikova's Mismatch Model, for example,

we noted that an absence from the observation already had to be experienced in order to know that it didn't match the relevant template. Similarly with Martin and Dokic's FoS account: it seems to make more sense to say that the feeling of surprise arises because an absence has already been perceived. The feeling of surprise cannot then be responsible for the perception of absence. Both accounts seem to get matters the wrong way around, deploying a mechanism that requires the absence to have been already perceived. Why does the same objection not apply to the account I am presenting now?

My account has a noticeable difference from the above. It is not trying to intellectualize perception of absence. It is not trying to provide it with a coherent, rational structure, whereas both the perceptual and cognitive accounts were in that business. Yet some explanation will still need to be given of how perception of absence occurs, even if it is an explanation of a possible causal mechanism rather than a rational structure.

A further example from Dennett will help in describing such a causal structure, while also helping to illustrate the idea of the user illusion of consciousness. In one experiment that Dennett (1991: 114) describes, which was first documented by Wertheimer (1912), two discs separated by up to 4 degrees of visual angle flash on and off alternately. When the discs are on, they are visible but invisible when not. The experiment revealed that if the discs appear alternately within a range of intervals (that is, with a certain speed), it did not look to a test participant as if there were two discs flashing on and off alternately but that there was a single disc moving or 'swinging' from side to side.

The effect of the alternating discs on an observing subject raises a puzzle for the understanding of consciousness, which threatens the idea of the Cartesian inner theatre and what philosophers of mind call transparency. We tend to think that the contents of our own minds are that with which we are most familiar, and can know with certainty, incorrigibly. To the observer, it looks like a single disc moving from side to side. Conspicuously, it looks as if the disc moves through the intermediate places between the disc's left and right extremes. We know from the experimental set-up, however, that the disc is never in these places, despite the appearances. We can conclude right away that some kind of illusion is occurring, therefore. Taken alone, this is perhaps not so dramatic a finding, since we are subject to many optical and other illusions on our senses. But there is something distinctive about this one since it seems to mean that our brains have reordered the sequence of events that have impressed themselves upon us.

The events that occur, and impress themselves upon the observer's retina, can be described as follows: 1. Left disc on; right disc off, 2. Right disc on; left disc off, 3. Left disc on; right disc off, 4. Right disc on; left disc off...and so on. The experimenter, in setting up the experiment and controlling the conditions, knows that this is the real order of events. In experience, however, there are extra stages inserted. The experienced sequence could be described as: 1. Disc at the left, 1*.

Disc moves from left to right, crossing the centre, 2. Disc at right, 2*. Disc moves from right to left, crossing the centre, 3. Disc at the left...and so on.

The issue here is that the disc appearing in the centre, during 1*, must be caused by its appearing on the right, at stage 2. We have a very good reason to believe this since when we switch the experiment off, the disc goes to one side and then just vanishes. It does not start out across the middle and fade away somewhere around there. This heavily suggests that it appears to be in the middle, moving right (stage 1*), only because it has already appeared at the right (stage 2) and the eye has received this information. It follows, then, that the brain must have created the appearance of stage 1* only after 2 has occurred (similarly, 2* will have been created after 3, 3* will have been created after 4, and so on). The experienced sequence of events has been reordered from the factual sequence, therefore. The eye must have 'seen' the disc on the right, at 2, but then 'told' itself that 1* occurred before 2, even though the experience of 2 must have been the cause of 1*. Like many other illusions, this one is involuntary from the point of view of the subject, since we still experience it this way even when we know that the experience differs from the reality.

In the light of this interpretation of Wertheimer's experiment, we are in a position to return to the case of perception of absence and dispel some of its mystery if we accept that there is also a mechanism involved that is capable of reordering a sequence of events in experience. Hence, a feeling of surprise when an expectation mismatches an experience could accompany the drawing of a non-deductive experience from what is seen but, nevertheless, be experienced as simultaneous with the seeing, and even as that which marks it out as a perception of absence. Hence, we could have a rapid sequence in which the scene in front of a subject makes an impression on their senses, a non-deductive inference is drawn from what is present to what is absent, where this mismatches some expectation or template, but then the feeling that follows from that is experienced as if it is a part of the original seeing. We do not have a puzzle, then, of what is experienced first (an ineffable phenomenological feel) being dependent on what must already have happened (information gathered via the senses), since we allow that it is part of the causal mechanism that it reorders the sequence of events in the user's experience. And it is worth noting that perception of absences and Wertheimer's discs would not be the only such examples of this reordering. Dennett's details many other such examples (see his 1991, chs 5 and 6, for example). Note also that this account is not relying on a prior perception of absence in order to explain the current one. A perception (of absence) does not occur at the moment light hits the retina or when a signal reaches the brain. On the mutual manifestation model, a perception does not occur until the causal process of perception, including cognitive and other mechanisms, has occurred. Information is available to the subject, from which the non-deductive, non-conscious inference is drawn, but it is

only once that inference is drawn and presented to the user in experience that we can say that a perception of absence has been made.

I said that such an account could be extended to non-human animals and human infants. Does the above account assume too much of them? The hardest cases of perception of absence, for the Parmenidean to accommodate, are those where there is an apparent phenomenal experience of absence and I have tried to explain these. I have done so by invoking an evolved causal mechanism. I cannot say for sure whether animals have a phenomenal experience at all nor, if so, whether it is like ours. It is undeniable that animals perceive their environment, however, and empirically plausible that they perceive absences. Whether they do so through a phenomenal experience or not, an evolved, non-rational mechanism is still likely to play a role in this ability. We cannot rule out the possibility that different species developed different mechanisms for perceiving absences. In our own case, the empirical evidence that even very young infants seem aware of absences, such as in visual cliff experiments, strengthens the case for a non-rational, causal mechanism being responsible rather than there being any norm-governed rational inference involved.

7.6 Where this leaves us

We are now in a position to draw some conclusions and present a verdict on the perception of absence.

The first thing to note, given the account that has been developed, is that the reality of perception of absence, as understood here, does not present a threat to our soft Parmenidean programme. A number of our other topics would present such a threat. If we allow negative entities, negative properties, facts, instantiations, truths, and so on, it seems that we are committing to the reality of negatives: allowing them some being. But not so here. There might be some theories of the perception of absence that hold this danger: for instance, if one allowed that it involved causation by absence. But our account doesn't do that. The negatives involved in the present account are entirely epistemic, since perception of absence does not require the reality of any absence perceived. Instead, as we have seen, perception of absence is merely a useful non-deductive inference drawn from what is there to what is not there, cognitively mediated and presented in a certain way in human experience.

It might yet be wondered whether this account sides with a perceptual or a cognitive theory, as those two options were outlined above. My task was mainly to explain what perception of absence is, and I am satisfied with the account presented. In examining those other theories, my aim was to show what was attractive and unattractive about each. The theory offered is one that retains the

attractions by keeping elements of both the perceptual and cognitive theories. Yes, absences are experienced perceptually but only because of a mechanism where an inference is presented phenomenally. By taking something from each of these theories, the shortcomings of both are also avoided. The new account thus blurs the distinction between a perceptual and cognitive theory.

What recommends such an account is that it reconciles the constraints, a–e (§7.5, above), that came out of explicating the problem through previous theories. The causal constraint (a) is satisfied as the perception is caused by what is seen. Importantly, though, there is no requirement of a causal interaction with the apparent absence that is perceived. Nevertheless, the experience of absence is causally related with what is seen, so it is not as if the causal interaction is of a non-perceptual kind, such as when a drug causes a hallucination. The experience of an absence is a way in which perceptual information is presented to the person concerned.

Second, this account does not rely on a prior perception having occurred, in order to explain our first perception of absence (b), nor a prior feeling of surprise whose best explanation is that an absence has already been perceived (c). Instead, the account says that information can be gained through the senses, which is then mediated cognitively, but results in an experience of absence with which any sense of mismatch or surprise is simultaneous phenomenally. This is perhaps the most difficult part of the account, so I used some of Dennett's cases to give us a grip on how such a model of the mind could work.

Third, we require no rationally valid deduction of absence from what there is, since none is available (constraint d). In its place, we have a causal mechanism that takes us from visual information to an experience of absence. Although this type of experience is not one that is strictly supported by the evidence, we can understand why it would have evolved since the ability to experience absences in this way will have conferred an adaptive advantage.

Finally, the account accepts the reported phenomenology of perception of absence (constraint e) and the distinction between knowing something (propositional knowledge) and knowing it through experience (experiential knowledge). Absences can be known experientially. However, the account also explains why there can be doubts about that phenomenology: why it is rightly considered to be uncertain and elusive, and why it might even be nothing more than a feeling of surprise or disappointed expectation in what is seen. Perception of absence is not the same as perception of presence where, in large part, the phenomenology is more clear-cut. On the account advocated here, the cases involve two different mechanisms where perception of absence is a useful user illusion. Indeed, this could explain what we noted in respect of the Kanizsa triangle, in Figure 7.1: that it was possible both to see the triangle and to see that the triangle was not there. There are multiple cognitive mechanisms at work in perception, and it is possible to switch between them, seeing the same object in different ways. Of course, if

Dennett's account of consciousness is right, even regular perception—even subjectivity is general—is a sort of user illusion. But it is a different illusion.

The theory satisfies constraints a–e, therefore, and we can note that the mystery of perception of absence is to a large extent resolved. What supports this view, for our purposes, is that there is at least one viable theory of how perception of absence occurs without commitment to the reality of absences. Our soft Parmenidean principles are not threatened here, then. Furthermore, this is a theory that seems to have at least some empirical content, despite being largely philosophical. It might, then, be possible to test it. Even if we await that, the security of knowing there is at least one viable theory might make us ready to avail ourselves of some of the advantages of permitting perception of absence. As noted above, the reality of perception of absence could pay epistemic dividends; for example, in explaining how we ever had a concept of absence or nothingness in the first place and, later, how we are able to hold beliefs about what is not.

8

Empty Reference

8.1 Can we talk about nothing?

Parmenides (*Fragments* 6.1, see ch. 1, above) thought that non-being was unthinkable and unnameable. It is tempting to disregard this as a typical statement of the unfathomable Presocratics, but there is, in this case, a serious challenge behind it. How can anything we say or think refer to what does not exist? It is hard to deny that:

- (1) Centaurs do not exist.

But what is this statement about? Its truth means that there are no centaurs; so to what does the word ‘centaurs’ refer here? If there is no referent, how is (1) meaningful? Sentences such as ‘Beyoncé sings,’ ‘a rectangle has four sides,’ and ‘kindness is a virtue’ are all about things. ‘Beyoncé sings’ is about Beyoncé, the person. ‘Kindness’ refers to a character trait or action. Not all words have a referent—‘is’ and ‘a’ don’t—but many of them do, and the reference determines what those sentences are about.

The problem addressed in this chapter is how we talk about non-existents, as it seems we do in ordinary statements such as:

- (2) Father Christmas gives gifts
(3) Oliver Twist ate porridge
(4) Mermaids are half fish and half human
(5) There are no square circles.

(1) and (4) concern mythical creatures. What is said of them seems to be true. But can there be truth if the referent that is the subject of the sentence is non-existent? How, if at all, are these about anything? Similarly, (3) concerns a fictional subject. What is said of the subject seems to be true in a way that:

- (6) Oliver Twist was King of Rome

is not. It is possible, of course, in the cases of (2), (3), and (4) that someone could utter those statements without knowing that they have no subject of reference, such as when a child believes (2). Someone could have heard the story of Oliver

Twist and assumed him to be a real, historical figure. (2), (3), and (4) have the same appearances and forms as other statements that are about real things. But that similarity doesn't really tell us how we can understand statements that lack a referent since there are other cases, like (1) and (5), with a different form, which express negative existentials. They specifically say that the subject of the statement does not exist: they wear their lack of reference on their sleeves. Crane (2013: 25, 27) accepts a generalized negative existential that is perhaps the most obvious illustration of our quandary:

(7) There are things that do not exist.

This uses the quantifier expression 'there are,' which is usually understood as existence entailing, but then says of those 'things' that they are not.

My focus here is with the problem of empty reference: where a term is used although there seems to be no corresponding referent. There can be different aspects of this problem found in different kinds of referential statement. There is singular and general reference with their corresponding empty cases. A singular empty reference is found in a proper name or definite description, such as 'Oliver Twist,' where the concern is with the semantic value of a singular term. A general empty reference is found in a common noun, such as 'centaurs,' where the concern is with the semantic value of a general term. The details of how we treat singular and general empty reference might vary. One might say that the reference of the general claim in (1) is the empty set, for instance, or that the reference of (4) is everything, since it has to be understood as a generalized conditional, true of everything, that if it is a mermaid, then it is half fish and half human. I do not see how such treatments of the generalized cases avoid the problem of empty reference, since that issue still has to be answered eventually. We will still need to know to what the term 'mermaid' refers in the conditional 'if anything is a mermaid,...' Perhaps it is the empty set, again. But then we must also understand to what the empty set refers, which we saw (§3.8, above) also has questionable existential credentials.

The problem goes deeper than making sense of individual statements. If statements containing so-called empty names are really about nothing, then, as Braun (1993) argues, one empty name ought to be substitutable for any other *salva veritate*; that is, preserving any truth of the statement. But this does not seem to be the case. Let us assume that:

(8) Father Christmas has a beard

is true. Its truth is not unproblematic, of course, but many would accept it as true. It contains the empty name, 'Father Christmas,' which seemingly has no referent, since there is no Father Christmas. If (8) is about nothing, Braun argues that we should be able to substitute in place of that empty name any other empty name we

choose. The assumption is that there is only one nothing, just as there is only one empty set (Braun's argument could be run in terms of the empty set, *mutatis mutandis*). If all empty names have no referent, or refer to the empty set, substitution with another empty name should not affect the truth value of the statement since they would be about the same thing: nothing. Nevertheless, we would not consent to:

(9) Oliver Twist has a beard.

Typically, (9) would be rejected as false. Something has gone wrong, then. Maybe empty terms do have some kind of referent, after all; otherwise how could (9) be false? Or perhaps sentences containing empty terms are never really true (but then what about (8)?); or the substitution principle does not hold.

If we cannot give an account of empty reference, then Parmenides' prohibition of thinking or speaking about nothing is largely justified. Yet it seems that we can say meaningful and truthful things about what does not exist, including that it does not exist. There ought to be, then, a tenable account of empty reference. We should be able to explain what is happening in these cases where we appear to be referring to something that does not exist and saying true things about it. Perhaps the appearance is false, but, even if so, an explanation ought to be forthcoming of why it seems true.

My approach is as follows. I will uphold what Searle (1969: 77) called the axiom of existence, namely that whatever is referred to must exist. As Crane (2013: 9) says, reference involves a *real relation* where it is possible to refer to *x*, for any *x*, only if *x* exists. Various strategies for claiming that the referents in empty reference cases do exist, contrary to superficial appearances, will be rejected. Crane's distinction between aboutness and reference will then be accepted, however, allowing us to say that while we cannot refer to something that does not exist, there can be statements about non-existent things. I will then explain how, adopting a broadly fictionalist account in which apparent reference to the non-existent involves a pretence. However, I do not accept that we have to endorse the reality of fictional objects. The pretence is located not in the object of reference but in the referring itself: in other words, we pretend to refer when there is no object of reference. We must accept something like Donellan's (1974) distinction between fiction and error, however, since a knowing and deliberate empty reference is a different matter from a failed reference. In the latter cases, one is not knowingly pretending to refer, so we must say something else about those instances. We can indeed say something since there are explanations of aboutness other than pretence. We need a plausible account of aboutness, however, that is not merely in terms of the largely mental phenomenon of intentionality, which is Crane's focus. Some of the problems associated with empty reference would re-emerge if statements are supposed to be about mental representations, thus we need an account of aboutness in terms of a public and shared meaning.

8.2 The axiom of existence

The axiom of existence tells us that anything to which we refer exists. This is because referring is, as Crane puts it, a real relation. A real relation is one that holds only if all of its relata exist and reference seems to be such a case. For a word or words to refer to something, both the words and the referent must be (for an early expression of this view, see Frege 1892: 58). This holds whether the words form a proper name or a description, hence 'the tallest woman in Britain' refers only if a tallest woman in Britain exists. It could fail to refer if there is no tallest woman in Britain; for example, if there were several equally tall women and none taller.

The reason reference involves a real relation is that it is a success term; that is, to refer implies doing so successfully. One can attempt to refer but fail, in which case one has not referred at all. Reference is not alone in being a success term. For instance, it is true that *a* caused *b* only if *a* and *b* exist or existed. Thus, an entirely conclusive ground for denying that *a* caused *b* is that there is no *b*: for example, Bill did not cause Jane to be happy if Jane was not happy. Another way of expressing this feature is to say that the relation is factive, implying that Jane's happiness is a fact, for example.

There is no reason to belabour this point since it appears to be a no more than an analytic matter, following from the concept of referring. Thus, if someone said that they referred to *a* but there is no *a*, we see it as problematic. The temptation is to say that they did not refer since there is no such thing. We can say that they failed to refer, that they merely attempted to refer, or that the speaker misunderstands or is misusing the term 'refer'.

However, despite this, the axiom of existence has indeed been disputed. Martinich and Stroll (2007: 101f) acknowledge it as a long-standing orthodoxy, finding it expressed in many places, but still reject it. They think it right to call the view an axiom since it is assumed without proof, but they reject it on the grounds that it is contrary to ordinary opinion. They assert that 'It is a fact that in everyday discourse we do use language to refer to non-existent (including fictive) objects by name and to make true (or sometimes false) statements about such objects' (Martinich and Stroll 2007: 102).

This finding jumps the gun, however, casting aside the axiom of existence prematurely. In particular, it assumes that it is specifically reference that is involved in talk of non-existence rather than anything else. Shortly, we will look at a distinction between aboutness and reference that allows us to draw a different conclusion. Their verdict also assumes that we can only account for the truth of statements such as:

(10) Odysseus was married to Penelope,

and the falsehood of:

- (11) Holmes was married to Watson

if Odysseus, Penelope, Holmes, and Watson, despite their non-existence, are objects of reference rather than anything else. This ought to be a matter of suspicion since someone who assents to (11) is also likely to agree that there wasn't really a Holmes or a Watson. There is at least some problem to be solved, therefore, of to what those names refer. This difficulty seems acknowledged in the fact that while we have an initial inclination to accept:

- (12) Sherlock Holmes lived at 221b Baker Street

as true, upon further reflection we want to say that it is not *literally* true or that it is true only *in the stories*. If we are looking at ordinary usage, then, which Martinich and Stroll invoke, we ought to consider the full picture. There is enough evidence to suggest a difference between (12) and a similar statement that we accept as unequivocally true, for example, that Margaret Thatcher lived at 10 Downing Street.

For completeness, however, it should be registered that there is no universal agreement that statements such as these (10 and 12) are true in any sense. The Fregean tradition rules that there is a truth value gap created by an empty term (see Crane 2013: 57), so it is neither true nor false that Sherlock Holmes lived at 221b Baker Street (see also Strawson 1950, Kaplan 1989). This verdict seems implausible, however, in cases like (2) and (12), and outright wrong when it comes to negative existentials like (1) and (5). Even if 'square circle' fails to refer, it is hard to maintain that 'square circles do not exist' is indifferent between truth and falsehood. We will all be inclined to the former truth (even if some explanation would then be needed for negative truths, see chapter 9). There are also so-called extra-fictive uses of fictional names that are uncontestedly true, for example:

- (13) Oliver Twist was a fictional character created by Dickens.

And there are cases analogous to (11) that concern neither fictional characters nor negative truths, such as:

- (14) Mermaids are mythical

and

- (15) Square circles are commonly used as philosophical examples.

One reason that it is hard to dismiss statements such as (13), (14), and (15) as true is that they can be used as part of the explanation why reference to non-existents is problematic; that is, if it were not for the truth of statements like (13), (14), and

(15) it is not obvious why other statements containing empty reference are a concern. *Oliver Twist* does not exist because he is only fictional, mermaids are not real because they are only a myth, and square circles are discussed at all by philosophers only because they cannot exist. Hence, to say that statements like these have no truth value undermines the very idea of there being a problem of empty reference.

Let us accept, then, that there are statements containing empty terms that are *prima facie* true. Of course, if we cannot find an explanation for the apparent truth of statements like (10), (11), and (12) without allowing reference to non-existents, then we might have to revisit the axiom of existence and side with Martinich and Stroll in rejecting it. But we are not there yet and, as we will see, we need never be.

8.3 Proxy referents

There is nothing to which certain names or words refer yet those names or words do seem to be about something. Perhaps, therefore, there is something, after all, to which they refer. *Oliver Twist* was a character in a book, so maybe the referent of the name is not a real boy who existed but something like an idea of a boy that we have taken from the book and subsequent films. Perhaps there is some such suitable proxy for a real *Oliver Twist*, and it is this proxy to which we refer. Various proxy referents for non-existences have been suggested. However, none of them really provide us with a credible account of the reference that supposedly involves them.

A view that Crane (2013: 10, similarly Kroon 2005: 180f.) attributes to no one in particular is that when we refer to non-existent things what we are really thinking about are sets or properties of descriptions. This would have the advantage, if it worked, that it wouldn't matter if nothing actually instantiated that set of properties and there was nothing of which that description were true. There could still be a coherent thought, not about a particular but about, for instance, the combination of being a half-human half-fish woman, or the set of properties of being brought up in an orphanage, eating porridge, and being schooled by pickpockets. And it doesn't matter that nothing has these properties or matches the description since the properties are real whether or not anything instantiates them. Similarly, a description is still a description whether or not it is true of anything.

It is dubious that descriptions or sets of properties really would be satisfactory proxy referents, though. For one thing, it is dubious that reference occurs via a description, as direct reference theories emphasize (for example, Donnellan 1966, Kripke 1980, Kaplan 1989). When we refer to Ed Sheeran, we refer to him—the thing itself—rather than via a description. After all, different people might know a host of different things about him and understand different descriptions of him. How can different people refer to the same person or object if they hold different

understandings and descriptions of him or it? Some of the information in a particular description might even be mistaken, and yet still we succeed in referring. For example, I might ask a colleague who the woman is that is talking to John across the room. I still refer to her even if it's not John she is talking to but Jill. Furthermore, it is psychologically implausible that when we refer, we do so via a list of properties. I might have to think quite hard in order to come up with a description of Ed Sheeran, but I can surely refer to him in an instant without much thought at all.

Now one possible reply to this (what Crane 2013: 11 calls the moderate view) would be that while we refer directly to things when a real referent exists, given that there is no such referent in the cases of non-existences, then there we refer only to the proxy: a set of properties or description. But this defence is unlikely to be persuasive. Again, it is psychologically implausible that this is what we do. The same problems hold in that two people could succeed in referring to *Oliver Twist* despite having different descriptions of him in mind. But we have an additional problem, now, that people sometimes talk about non-existents unknowingly. Someone might mistakenly think that *Oliver Twist* was a real historical figure and then, according to this reply, refer to him directly rather than via a description. But how could this happen, on such a view, if there is no real *Oliver Twist*? Or someone could make a different mistake, talking about the highest prime number, for example, in the belief that there is one.

This problem shows us something more general about the problem of empty reference. We can unknowingly talk about things that do not exist. Such cases suggest that there is no detectable psychological difference between reference and failed reference. Any theory which says that we are doing psychologically different things in these two cases looks implausible.

Let us consider another proxy referent for non-existences. The early Russell drew a distinction between being and existence in an attempt to solve the problem that concerns us. He made the distinction as follows:

Being is that which belongs to every conceivable term, to every possible object of thought...Numbers, the Homeric gods, relations, chimeras, and four-dimensional spaces all have being, for if they were not entities of a kind, we could make no propositions about them...For what does not exist must be something, or it would be meaningless to deny its existence; and hence we need the concept of being, as that which belongs even to the non-existent.

(Russell 1903: §427)

Everything that we can conceive has being, according to this view, but having being does not entail existing. Russell cites the possibility of reference as a reason for a division between being and existence. Non-existences can nevertheless be objects of thought. They have to be 'entities' of some kind in order for there to be

propositions about them. The referent of a term for a non-existent, then, is an entity that has being but not existence.

One question is to ask whether this remains a plausible theory specifically of reference. Do we know whether we are referring to an existent or only a being? And does that matter? The most likely concern with such a view, however, is its metaphysical implausibility. How could there be a division between being and existence? What would it amount to? As Routley (1980: 42) said, 'there is only one way of being, namely existence'; and it seems that we should say too that the only way to exist is to be. Admittedly, the problem that led Russell to make the distinction is a real one, but his solution works only if a division between being and existence can plausibly be maintained as more than a merely verbal difference, which Quine (1948) insisted it was.

Furthermore, Russell's early view, that every conceivable object of thought has being, means that the golden mountain and even the square circle would have being. The only way to avoid such a conclusion is to claim that such things are inconceivable, which is not convincing. As Crane (2013: 24f) argues, Russell (1959: 64) tried to foist upon Meinong the view that the round square had being. But Meinong's (1904) view was more refined than that. While Meinong draws a distinction between being and existence, it was because he thought that only spatiotemporal things exist. Non-spatiotemporal things have another mode of being, namely subsistence. Round squares have neither existence nor being. We can still say and think things about round squares, but there is no need for them to have being, or existence, in order for us to do so. Meinong's view is thus more credible and less metaphysically revisionary than Russell's early view, particularly because Meinong gives a clear statement of what he takes being to consist in, rather than it being what any term must have in order to be conceivable.

Let us then consider this conceivability some more since it leads to what is a natural proxy referent for a non-existent. This would count as another purely relational conception, to use Crane's (2013: 10) term. The view is that an empty term (a term for which there is no existing referent) refers purely to an object of thought rather than a mind-independent existent. In defending such a view, one might say that while it is obvious that Father Christmas, Oliver Twist, and the square circle do not exist outside of our minds, what we must be talking about when we refer to them are the ideas we have of such things. Thus, Father Christmas does not exist as a concrete individual, spatiotemporally located and causally powerful, but he does exist in my mind and the minds of many other people. It is this mental representation that we are talking about when we talk about Father Christmas. In further defence of the view, we can invoke another term from Meinong (1904), which is his notion of being-so (*sosein*). This allows that the properties a thing has are independent of its being. Hence, Father Christmas can *be-so* that he is jolly and red-faced even if he does not exist or has no being. The representationalist can explain this since the idea of Father Christmas is of a man

with a jolly and red face even if that man is not real. Similarly, the idea of Oliver Twist comes with an associated *sosein* since he is described in some detail in the novel. All his properties belong to the representation alone, however, since there was no Oliver Twist in reality. Even if there was an actual orphan boy called Oliver Twist who lived in London in 1837, it is clear nevertheless that most of our thoughts are not about him, since any such resemblance is coincidental. We are thinking specifically of the fictional Oliver Twist, who exists only in our thoughts.

Some reflection reveals that this representationalist view will not stand, however. It is true that there are objects of thought, but this does not mean that the reference of such thoughts can be to the thoughts themselves. The representationalist view fails specifically as a theory of reference. In the first place, if 'Father Christmas' refers to the idea of Father Christmas, then the negative existential 'Father Christmas does not exist' comes out as false. The *idea* of Father Christmas exists, according to this theory—and does so in order to avoid the problem of empty reference—but then that leaves us unable to assert a true negative existential. If this account is generalized, then any question of whether *x* exists would have to be answered trivially in the affirmative, since its account of reference is that we refer to the idea of *x* whenever we refer to a non-existent. Indeed, any statement that *x* does not exist would be trivially false. It is clearly true, however, that some things do not exist, even if we are struggling to explain how best to say this in a non-contradictory way.

This is not the only problem with the representationalist move as a theory of reference, and the second problem is one that will be important later, in what is to come. No theory of reference is plausible that locates the referent in the head of the referrer since it makes communication impossible. On this theory, when John refers to Father Christmas, he refers to the idea of Father Christmas in his head and, when Jane refers to Father Christmas, she is referring to an idea in her head. This means that John and Jane are talking about different things when they use the name Father Christmas. Thoughts are private while meaning and language are public. For the latter to function, we must all be able to talk about the same thing. And given that each person's mental life is private, we can only know our own. Hence, it is not even possible to nominate one authoritative thinker and say that the referent is her idea, since only she has that idea.

This was the point of Wittgenstein's (1953: §1.293) beetle in the box example, which can be understood as directed against a Lockean theory of language. Locke (1690: III, ii, 2) had said that the meaning of a word was an idea in the head of a speaker. Wittgenstein likened this to us all having a box with the word 'beetle' written on it, but where we are allowed only to look in our own boxes and no one can look in each other's. Somehow, we are expected to accept that we are all talking about the same thing when we use the word 'beetle', but since we cannot check this by comparing what is in someone else's box with what is in our own, then we

have no way of knowing that we are talking about the same thing. Without talking about the same things when we use words, it seems that we cannot communicate since we cannot know what anything means. This problem places a constraint on any viable theory of reference. It must be possible for reference to be public in the sense that we can refer to the same things, whether it be particular objects, such as people, places, and things, or properties and attributes such as redness and courage. Meaning has to be shared. Ideas, located in private spaces in our heads, just do not seem capable of delivering that.

A further proxy object of reference for empty cases would be fictional objects, but as I am going to consider fictionalism in more detail shortly, I will leave fictional objects aside momentarily.

8.4 Aboutness

A distinction between being and existence was rejected where it has no plausible basis. Another distinction can be defended, however, which I take to be useful for the problem of empty reference. This is the distinction between reference and aboutness.

Thus far, I have spoken as if the only way in which a statement can be about something is by referring to it. Now we can see that this is too loose a way of speaking, however, and that there is a distinction to be made. If we maintain the view that reference is a real relation, which I think we should, then we see that we can have thoughts and statements about things without referring to them. This is because we can have thoughts and statements that are about things which do not exist.

Some see this aboutness, under the name intentionality, as the defining characteristic of the mental, which is the Brentano (1874) thesis. A part of his notion of intentionality is that of intentional inexistence, and Crane (2013) defends the distinction between aboutness and reference on the basis of intentionality. One can fear a burglar downstairs when none exists: it was only the cat knocking over a lamp. The fear is real even if the burglar is not, and the fear is about or directed towards the burglar in one's house even if there is none. Similarly, one can hope for world peace, even if it will never come. One cannot fear the burglar if one knows that he is not real, hence when he is not real there must be at least the tacit belief that he is, if one is to be afraid of him.

The notion of *aboutness* is a semi-technical notion. Of course, being *about* something is not technical: that is a very commonplace notion, involving directedness towards an intentional object, which might or might not exist in reality. However, the usefulness of the notion is justified by it being able to make sense of statements such as (1) to (5) but also, importantly, those like (7). We want to be able to state negative existentials, including the generalized one in (7). While such

statements are common and routine, there is admittedly a tension within them. The tension resides in the apparent contradiction between the two parts:

(7a) There are things ...,

and

(7b) ...that do not exist.

(7a) seems to affirm existence of something, while (7b) denies existence of the very same thing. Indeed, the tension appears whenever we say of a thing that it does not exist. If it does not exist, how can it be a thing? At best, isn't it a 'thing', where scare quotes suggest that it is not quite that?

This apparent tension is resolved if we accept a distinction between aboutness and reference since aboutness is not existence entailing, whereas reference, being a real relation, is. Hence (7a) can be taken to mean 'There are' in the aboutness sense, such as when we can have thoughts about square circles, burglars downstairs, *Oliver Twist*, and mythical creatures without it thereby requiring that those things exist, which is what (7b) says.

The acceptance of aboutness allows that we can use quantification in a non-existential way (Crane 2013: ch. 2). The Quinean orthodoxy was that we are ontologically committed to that which we quantify over (Quine 1948, 1969: 91–113), with quantifier phrases such as 'there are...', 'there is at least one...', 'most...' and so on. 'There are...' is typically taken to be an assertion of existence. Perhaps this is so only in philosophical circles, under Quine's influence. Statements such as (7) clearly cannot operate that way since it would render them self-contradictory. In a noneist view (Routley 1980, Priest 2005), there are non-committing quantifiers. But a variant on this is to say, as Crane allows, that we can quantify over non-existent objects, which we do with (7).

What exactly aboutness consists in, and how it is possible, is a matter we can put aside now, but we return to it in 8.8, below. First we should consider what we are able to do with such a distinction. One thing we can do is explain what it is for a thought or statement to be about something that does not exist. This is to say what is happening when we apparently refer to a non-existent.

8.5 Fictionalism

One way in which we can speak and think about things without referring to them, I maintain, is in a fictionalist way. Fictionalism has found a wide variety of applications in philosophy (see Kalderon ed., 2005: 1–10), but my concern here is purely with fictionalism as an explanation of empty reference (for which see also

Kroon 2005). This does not mean, for instance, that I am advocating fictionalism over mathematical entities, though I am not denying it either. I explained in chapter 5 how fictionalism had an appeal in the area of modality, but my approach in this regard is a piecemeal one: to adopt a fictionalist stance where it is explanatorily useful to do so, and only then.

What should fictionalism look like for empty reference? The idea behind fictionalism is that an area of discourse is not literally true but is to be taken only as we would a fiction. Kalderon characterizes it thus: 'The distinctive commitment of fictionalism [is] that the acceptance in a given domain of inquiry is not truth-normed and does not involve belief in the content of the accepted content' (Kalderon 2005:3). Thus, it is not literally true that the Artful Dodger picked pockets since no such incident ever occurred. It is entirely sufficient to establish this that the Artful Dodger did not exist since non-existents cannot pick anyone's pockets. But we can take it as fictionally true. This could mean that it is true in the fiction that Artful Dodger picked pockets, but it is better still, as I will go on to explain, to say that it is a fiction that it is true that Artful Dodger picked pockets. The discourse that interests us is not just literary fiction, however. If we apply a fictionalist stance to empty reference generally, then we can say the same of mythical beings. That Medusa had snakes for hair is not literally true, as almost everyone knows, but we speak, within the discourse, 'as if' it were true, in 'story-telling mode', to use Azzouni's (2010: 113) phrases.

What is useful about the parallel with literary fiction is that there can be fictions concerning all sorts of incredible and unrealistic scenarios but which we easily grasp and accommodate once we realize that it is not reality but a form of make-believe (Walton 1990). Harry Potter performs magic, for instance, yet even those of us who seriously doubt there is any such thing as magic are capable of adopting a fictionalist stance towards it. Doing so allows us to understand what is said and then enjoy the story, part of which comes from engaging in the fantasy, with an 'engaged perspective', says Kroon (2005: 185), suspending our disbelief. Our ability to comprehend unrealistic scenarios thus allows us to consider counter-nomic fictions as we might, for example, when we consider what the world would be like if the gravitation law were an inverse-cubed law instead of inverse-squared. The inverse-cubed gravitation law has no reality whatsoever—it does not exist—and yet we are able to talk about it and put it into seemingly truth-apt statements.

An extreme example of the flexibility of fictions is counter-logical discourse. We saw in 5.3, above, that Jenkins and Nolan (2012) think, reasonably, that we would be very surprised if we encountered a square circle. This seems right even though we know that there cannot be a square circle. Indeed, the fact that there cannot be one is the very reason why it would be a huge surprise to encounter it. If there cannot be a square circle, though, how can anything we say about it be truth apt? Fictionalism answers this question. We are able to engage even with fictions about the impossible such as this one in which contrary properties are

instantiated in an object. We can think about what we would do if we had a magic wand—a physical impossibility—and evidently we can engage in stronger impossibilities too. Still, it might be wondered what it is to have any kind of impossible object as a referent, since such a thing having any being at all seems problematic, for which a solution will be offered shortly.

While it seems an advantage of fictionalism that it can accommodate outlandish cases, we should not lose sight of its application to the more standard cases of empty reference. The Artful Dodger did not exist in reality, for example, so we cannot refer to him, given that reference is a real relation. But we know that the fictional Artful Dodger picked pockets, so this is a fictional truth rather than a literal truth. Within the fiction, (12) is true but (11) is false, and for them to be so depends on there being a fictional reference to Sherlock Holmes. This allows us to say that there are books and statements that are about Sherlock Holmes even if they don't refer to him, because he does not exist. The suggestion is that we try to extend this, as far as we can, to all empty terms; that is, all cases where an apparent referent is non-existent.

Thus far, this is offered only as the basic insight of fictionalism. We quickly see, however, that the account cannot be extended to every case. There are some examples of empty reference where there is no obvious intention to treat the referent as a fiction, for instance, where someone talks about something without knowing that it is non-existent. There are also cases where one talks about fictions but not within a fictional discourse. To understand such cases, it is necessary that we probe deeper into what a fictional reference is, and we must also utilize a further distinction.

8.6 Pretence

A Platonist might think that fictional objects are real since centaurs and fictional characters could exist as abstract entities on such a view. While such objects are abstract rather than concrete, a Platonist sees that as no reason to deny their reality. Indeed, the spirit of Platonism would treat abstracta as even more real than concreta. Likewise, a modal realist could allow that the Artful Dodger is real in the sense that there are some possible worlds in which he exists, or counterparts of him, and it just so happens that he is not real in our world. But I am deliberately setting aside these options as I am neither a Platonist nor a modal realist. What do I say about fictional objects, then?

Following Walton (1978, 1990), I want to say that engaging with a fiction is taking part in a form of pretence. The fiction only pretends that an Artful Dodger picked pockets. Likewise, it is only a pretence that the gravitation law is inverse-cubed, only a pretence that Medusa had snakes for hair and only a pretence that there are square circles.

But what is the pretence? Where does it reside? How would such pretence work specifically in relation to reference? The best answer is that the reference is pretended. The other option is to pretend that the objects of reference are real. This is not necessarily to say that fictional objects really exist, as the Platonist thinks. It is only to pretend that they do. The option is not entirely persuasive, however. In the case of counter-logicals, it is not clear what it is to pretend that an impossible thing is real. One can understand what it would have to be to be a square circle. But to understand that is know that there could be no such thing, as a matter of absolute necessity. What would then even be the point in pretending that there is an impossible object? Is that something that can be pretended? Doing so seems to be acceptance of a burden that is impossible to fulfil but also an unnecessary one. In particular, there should be no need for us to be forced into accepting any sort of world of fictional objects.

Perhaps such difficulties can be managed but there is an attraction to the first answer we gave, above. It is not that we are genuinely referring to something and that the something is a fictional object. That would count as another proxy theory of reference, where the proxy is the fictional object rather than the real deal. Rather, the pretence resides in the referring rather than the referent. Given that there is no referent, in these cases, then we cannot refer to them, as we have seen, since reference is a real relation. But we can pretend to refer to such things, and this is what I am suggesting that we do in fictions. We are not referring but pretend referring. This is where we get aboutness rather than reference. Aboutness is pretend reference.

What, then, is it to which we refer in fictional discourse? Nothing, is the correct answer. We do not refer to the Artful Dodger nor any proxy for the Artful Dodger. The term has no referent. There should be no temptation to seek one if we know that he does not exist. Nevertheless, thoughts, fictions and statements can be about the Artful Dodger—about something that does not exist—since aboutness does not require the existence of its subject. Statements about square circles do not require their existence or being in any respect in order to be about them, which is just as well since this allows us to say, without contradiction, that they do not exist, even that necessarily they do not exist. To deliberately talk about a non-existent is to pretend to refer to it.

Negative existentials are thus again a good test case of the theory we are seeing emerge. If one offers a version of fictionalism in which the reference is genuine, then it requires a referent, even if that referent is a proxy, such as a fictional object. Saying that square circles do not exist should not require us to say that something exists of which we then deny the existence. There is no necessity to do this if we see such cases as involving pretended reference only. Since we think we are talking about something that does not exist, we should accept categorically that it does not. Our something is a ‘thing’, rather than a thing, in that it is something to which we can only pretend to refer.

The same account also suffices for the case of extra-fictive empty reference, which Azzouni (2010: 114) calls fiction-external statements. Negative existentials are just a special case of this. Fictive uses of empty terms (fiction internal) are those which our core theory has addressed. When we say that the Artful Dodger picked pockets then we are engaging in the pretence of fiction and clearly with a fictional discourse. I agree with Azzouni that fiction-internal statements are only pretended true or pretended false, since they are true only in the fiction. It is also clear that there will be many fiction-internal statements that are neither pretended true nor pretended false. For example, Dickens does not say in the novel what colour underwear Oliver Twist was wearing on the day he first met Fagin, so it is neither pretended true nor pretended false that his underwear what white that day. Extra-fictive statements can be about the same things as fictive statements but in statements that are not internal to the fiction, for example, our (13):

(13) Oliver Twist was a fictional character created by Dickens.

There is nothing in the story itself that says whether the character was fictional or by whom it was created. It is possible to write a novel in which a fictional writer creates a character: the character David Copperfield becomes a novelist, for instance. It is rarer, though still possible in a surrealist novel, for a writer to create a character and acknowledge that she, the author, is its creator. Setting aside such special cases, however, it is clear that nothing in the novel *Oliver Twist* says that (13) is true.

Extra-fictive uses of empty reference have to be distinguished, however, because it is clear that (13) is not merely pretended true but really is true, literally. Now one might wonder how (13) can be literally true if it contains an empty term, which refers to nothing. But we see that we have the resources to answer this question, now, without revising our theory. 'Dickens' refers to Dickens, but 'Oliver Twist' refers to nothing. In using that latter term, we only pretend to refer to something, since we acknowledge that there is no real Oliver Twist. But pretended reference within a statement does not automatically mean that the truth of that statement is also only pretended. It is literally true that Dickens created a popular character to which we pretend to refer with the term 'Oliver Twist'. We see, then, that it is not sufficient for a truth to be pretence that a reference within it is pretended. We cannot refer to Oliver Twist, since he does not exist. We can make statements *about* Oliver Twist, though, when we pretend to refer to him, and the statement being about Oliver Twist is enough for the possibility of (13) being literally true, given that it is extra-fictive (some other conditions have to be met for it to be true, of course).

There are some other types of extra-fictive empty reference, though they need not detain us long since the same account is applicable. For example:

(16) Oliver Twist is more famous than Uriah Heep

(17) Le Verrier thought about Vulcan,

as well as our earlier (14) and (15). The truth of (16) is not a fictional truth since it is not stated in any fiction. It compares the relative fame of characters from two different novels. It involves two pretended references, but it is a literal, empirical truth that one fictional character is better known than the other. (17) is also literally true since Le Verrier was a real person and he certainly had thoughts about Vulcan, even though Vulcan didn't exist. Our (14) involves a concession that what it is about is not real, and it is literally true that mermaids are mythical. Given that there are no mermaids, then, we can no more than pretend to refer to them. Finally, (15) is true because philosophers have frequently used square circles as examples. They are used so often precisely because they are things that cannot exist so, again, there can be no reference to them.

8.7 Reference failure

There remains one more potential hazard for our pretence theory of empty reference. We saw that (17) is true. Vulcan is the 'planet' that Le Verrier thought existed, somewhere between Mercury and the sun. Its presence there would explain the perturbations in Mercury's orbit in much the same way that Neptune was posited for a similar reason before it was observed. From 1859 to 1915, others were still searching for Vulcan and occasionally reported its observation. However, there was no such planet, and Einstein's theory allowed the motions of Mercury to be explained without the need for an additional planet.

This presents a potential problem to us. When we consent to (17) we do so knowing that there was no Vulcan, so, on the present theory, take it to involve an empty reference where we only pretend to refer to something with the term 'Vulcan'. Le Verrier himself, however, believed that he really was referring to Vulcan rather than only pretending to do so. This raises some questions. Was Le Verrier doing something different from us? Does it matter whether or not one knows that a potential referent is non-existent? Can one engage in pretended reference unknowingly? The latter problem occurs where someone believes that an intended referent exists when it does not. Can we tell them that they are engaging in the pretence of reference even in these cases? Martinich and Stroll (2007: 61) use this sort of consideration against the plausibility of a pretence theory of reference generally since it is said to give one account for genuine reference and another account for empty reference even when many people are either unsure or mistaken about whether a potential referent is real or not.

This is where another distinction is required, as it will help us disentangle some different cases. The distinction is not ad hoc since it is one that we can grasp independently of the current questions to which it is applied. The distinction is like the one in Donnellan (1974: 5, see also Crane 2013: 15) between fiction and error.

This corresponds to the distinction I now make between empty reference, which includes more than just fictions narrowly defined, and failed reference.

Thus far, I have been discussing primarily empty reference, which is distinguished by the speaker talking about something that they believe does not exist, or where a statement acknowledges the non-existence of something, as with negative existentials, or a discourse is understood to be fictional in the broad sense, including myths, counter-nomic, and counter-logical discourses. Such discourses are understood through participation in the pretence. Kroon (2005) calls these cases 'deliberate use' of empty reference but says relatively little about their accidental use (see Kroon 2005: 189).

Reference failure differs from deliberate use in that a speaker does not know of the non-existence in question and is not engaged in a pretence. There are a variety of reasons for this. As with *Le Verrier*, one might believe an intended referent to be real when it is not. Nevertheless, *Le Verrier* fails to refer if there is no referent. In other cases, someone might speak of the highest prime or the tallest woman in the room without realizing that there is none such. Reference fails because of the truths of mathematics, in the first case, and because of the contingent fact that there are two equally tallest women in the second case. Nevertheless, neither our speakers nor *Le Verrier* know they are failing to refer. What makes it the case that they fail to refer, therefore, is not some psychological fact about a specific speaker or believer. It is that 'refer' is a success term that, in these cases, is not satisfied.

To say this, we need not deny that there is any psychological difference between someone knowingly pretending to refer and someone unknowingly failing to refer. But what is not a psychological matter is whether they refer since the non-existence of a referent is alone sufficient for a failure of reference (which is not to rule out the possibility of a referent existing and someone still failing to refer to it: non-existence is sufficient for non-reference, but existence is not sufficient for reference). There is also a difference between thinking about something and not thinking about that thing. We thus have adequate conceptual resources to distinguish general empty reference from reference failure, but where does the pretence come into it?

A key difference between empty reference and reference failure is that the latter is a mistake while the former is not. The mistake of the latter is to attempt to refer to something that does not exist instead of pretending to do so. When we now read (17), most of us can see it for what it is, where there is no *Vulcan* so we can only pretend to refer to it. *Le Verrier* would not have seen (17) this way. However, his statements concerning *Vulcan* were not about nothing. However, their being about something did not consist in anything he pretended. This, however, is as it should be. Were he only pretending to refer, all would be well since pretended reference does not commit to the existence of the referent. Because he intended literally to refer, he was left with reference failure. We still need an account of how his statements were about something, and that must be an account that does not

involve pretence, since there was none in his case. An account of aboutness is still required in any case since the theory of pretence was never advertised as supplying the theory of aboutness.

8.8 Aboutness again

Accordingly, then, we must finish this chapter by revisiting the subject of aboutness. Crane (2013) gives an account of aboutness in terms of intentionality, the traditional mark of the mental. I have some reluctance to defend a purely psychological account of how thoughts, words, and statements can be about something. This is not because I doubt that thoughts can be about certain contents but because of reservations about treating intentionality as a special, irreducible power of minds, as Searle does in places—comparing ‘mentation’ to lactation as a special biological power of brains (Searle 1981: 372). Naturalistic accounts of this mental capacity should be preferable and possible. The value of such an account is that it would allow us to overcome Wittgenstein’s beetle in the box problem that was discussed above. For us to have thoughts about the same thing and, indeed, for words or thoughts to be about something at all, we need an account that grounds them in a public discourse of shared meaning rather than a private realm or Cartesian inner theatre. We cannot have any shared meaning, and meaning *simpliciter*, based upon the private ideas that Locke invoked in his philosophy of language.

Armstrong’s (1968) theory of intentionality, for example, explains it in terms of belief, action, will and perception. These form an interconnected cluster of psychological abilities that must come together. Crucially, they are tethered to the world around us, which we encounter in experience. Now it is relatively easy, in such terms, to explain how a thought can be about some object that one actually encounters: an apple that one perceives, for example, and has a desire to eat. Thoughts are about objects such as apples because our psychology is action-oriented. Perception and belief can play a role in action only if they lead to thoughts about things. This is well and good but what about intentional inexistence? How do we have thoughts about things that do not exist, since we cannot claim that such things have been experienced in the way that an apple is experienced?

The complexity of thought plays a role here, which is another, more acceptable Lockean thesis. Complex ideas can be built from relatively simple ones. This includes the constructions of things that don’t exist: unicorns, centaurs, fictional characters, and so on. Even square circles can be constructed from experiences, conjoining that of a square thing with that of a circular thing (see Crane 2013: 72 for a more detailed account of constructing the idea of a square circle). Our thoughts can then be about those constructs and we can pretend to refer to them. The sense of aboutness that is required is a public one, though, and the sense of construction is social construction. The concept of a circle which partakes in this

construction is a shared one: shared by those who are language and geometry users. We said that we could not accept proxy referents such as mental representations in individuals' heads to explain empty terms, since we would be unable to refer to the same things as each other. We see, now, that this means an account of aboutness is required that is explained in terms of language and meaning since these are public matters. I have not given myself the task here to offer a full theory of meaning nor theory of language, since that would be too ambitious. I am content to use a relatively uncontroversial claim that will be common to many theories: that meaning is public. The key is to accept that while some terms are based on experience, others are constructed as fictions, allowing us to partake in a shared language game of pretend reference. We can, thus, think and speak about the same things as each other.

We can still use the metaphor of directedness, which is commonly used to explicate intentionality (see Molnar 2003: ch. 3). Our thoughts and words are directed towards that which they are about. The target phenomena is necessarily a public one, I have argued, instead of an idea in the head. Thoughts can be about the things to which they refer or about the things to which they pretend to refer, and public, shared norms of proper usage govern both the referential and the pretence cases. We can deliberately direct our thoughts towards what we know does not exist, speaking about those things but not referring to them. Failed reference, however, is not simply a mental error: such as failing to execute the correct mental process of referring. Consequently, we cannot tell the difference between successful reference and failed reference by first-person introspection alone, since they can seem indistinguishable to the thinker. The difference lies in an external and objective component, whether or not that which the thought or statement is about is also an existent. In reference failure we are trying to refer rather than pretending to. The thoughts and statements are still about something, but the mistake is thinking they are about a referent rather than just a fictional or mythical social construction.

In the light of our discussion, we may revisit the original Parmenidean claim with which we began, that we cannot name or think about what-is-not (Pvi, 1.3, above). We noted there that Parmenides' position could be accepted or rejected according to how it was interpreted, and we now have a more detailed account of how that is so. In summary, we cannot refer to what-is-not in either thought or word and, under that interpretation, Parmenides was right. But he was wrong if Pvi is taken to involve aboutness since, as argued here, we can have thoughts about, and thus terms for, non-existents, hence statements can be about them too. Our soft Parmenidean should then accept the Pvi claim under one reading but reject it under another.

Adding a capacity to think and speak about what-is-not to our ability to perceive absences (ch. 7, above), we are well on the way to understanding how it is possible to have negative beliefs, concerning what is not the case. But how do we say truly what-is-not? This will be our next topic.

Negative Truth

9.1 The hippo not in the room

It is true that there is a table in the room in which I now sit. There is something in the world in virtue of which this truth is true. In front of me stands a table: a physical object that I can bang with my knuckles and rest my coffee cup upon. The temptation is to say that the table makes true the truth; that is, an object in the world relates to a statement about the world in such a way as to make it true. The matter is slightly more complicated than that, however. The statement in question is:

(1) There is a table in the room in which I now sit.

(1) is not made true simply by the existence of the table but by the rather more complex fact of the table being located in the room in which I now sit (Mulligan, Simons, and Smith 1984, Armstrong 2004). What makes (1) true, I am assuming, is not only the existence of an object but the existence of a fact or state of affairs in the world of which the table is a part or constituent. Other constituents of the fact would be the room, me, and, seemingly, the relation of being inside. Furthermore, these constituents must be ordered in a certain way to constitute the fact in question. The cat being on the mat is a different state of affairs from the mat being on the cat even though they have the same constituents: in the two cases, the constituents are organized in different ways.

We can understand truth to be a feature of some but not all statements, sentences, utterances, or propositions. Following Dyke (2008), I have some reluctance to commit to propositions as things in the world. There are undoubtedly utterances and sentences, and I want to allow that different people can mean the same thing using slightly different words or sometimes very different words from different languages. To avoid the metaphysics of propositions, however, it might be best to use the term *truth-value bearer* (or sometimes just *truth bearer*) for the sort of thing that is sometimes true and, presumably, true in virtue of how the world is. The truth-value bearer can also be false, when the world is not how the truth-value bearer says it is. I will follow the convention of putting angled brackets around this truth bearer, hence (1) becomes <There is a table in the room in which I now sit>. I will assume that there are only two truth values—true and false—but I also accept that some strings of words have no truth value.

Primarily, this truth value-bearer has to be understood as some kind of logico-linguistic entity whereas the things that determine its truth value frequently are not. The table is not a purely logico-linguistic entity, for example, nor is the room in which it stands. The table is a physical object that you can kick or drill into. The room is an extended space with physical boundaries. Even the relation of *being within* refers to something worldly and not merely linguistic. There is a difference between physical containment and logical containment, for instance. Truthmaking is described by Armstrong (2004: 6) as a cross-categorical relation precisely because, when it holds, it holds between something worldly and something logico-linguistic. Perhaps relations that span these two very different categories are mysterious in some ways but anyone who has a realist theory of truth seems committed to them in some form. A minimal requirement on such a theory is the idea that what is true is dependent in some way on what there is in the world or how the world is. There has to be some connection between the truth of a statement and the world, otherwise truth would float free of reality, which is not what we realists want the concept of truth to be.

There are exceptions to this, of course. Some truths are true purely in virtue of the language used, as with self-referential statements such as:

- (2) <This statement contains five words>

and analytic statements such as:

- (3) <Bachelors are unmarried men>.

I will set these aside as they do not help us to identify the problem at hand, and there is, in any case, plenty of attention in philosophy given to a priori and analytic statements and logical truths. To get to the problem of negative truth, I must instead focus on contingent truths, of which there are many, such as:

- (4) <Buckingham Palace is in London>,
 (5) <There are six coins currently in my pocket>,
 (6) <The Battle of Hastings occurred in 1066>,
 (7) <Dogs exist>.

In these cases, as truthmaker theorists say, the true statements are made true by something in the world or how some local matter of fact is in the world. When I say that there is a table in the room (in which I now am), it is clear that my statement is of this kind: true (when it is) in virtue of something in or about the world. Presumably, if it were false that there is a table in the room, this would also be false in virtue of the way the world is (or, at least, the local portion of the world where the room is located). The broadly realist suggestion behind this is that

truth supervenes or depends on being, hence there is a relation that holds between the world and certain statements when they are true that does not hold when they are false. Statements such as (6) concern the past, so we will need to say more about them. What we will say depends on our other commitments concerning the metaphysics of time. My own preference is for a growing-block theory in which the past is treated as part of reality, just as much as the present, hence there is indeed a fact within this growing four-dimensional reality that makes (6) true.

Perhaps this account of truth can be challenged, but let us put any such challenge on hold and consider one specific issue that would arise from it. What of a truth such as the following: there is not a hippopotamus in the room (in which I now sit—a qualification I shall henceforth assume without stating)? Hence:

(8) <There is not a hippopotamus in the room>

(8) looks true, I would have thought: as true *prima facie* as anything could be. But what is there in the world that makes it true? The more I think about it, the less confident I am that I can find any existent thing in virtue of which such a truth is true, as I will detail below.

The problem with a statement such as (8) is that it seems to be a truth not about how the world is but about how the world isn't. It states that something is not the case. So how can a truth about what there isn't be made true by what there is? Can truths ever be made true by what there isn't? And then it seems we really have a problem: how could what there isn't—something non-existent, presumably—bear a relation to something else: a statement? The requisite cross-categorical relation looks even more problematic in the negative case. We usually think that for a relation to hold, its *relata* must exist. If we include what there isn't in a relation, therefore, we would then have reified a nothingness: a *what-is-not*.

I have posed this problem, which Horn (2001: 49) calls the paradox of negative judgement, in the terms of truthmaker theory: of identifying the truthmakers of negative truths. But the problem is more general than that since it applies to any broadly realist theory of truth, such as correspondence theories. Take this example, for instance, which was posed prior to the development of contemporary truthmaker theory:

The correspondence theory of truth seems to require the postulation of something in the world to serve as the correspondent of a true proposition about the world. Since there are true negative propositions about absences, lacks, and privations in the world, there must be such negative events or states to serve as their correspondents. If it is true that a forest fire does not occur now in Cook Forest, there must be a negative event consisting in the non-occurrence of a forest fire or a forest fire not occurring now in this Forest. (Gale 1976: 1)

This is the problem of negative truth. It seems that there are many things that are both negative and true, for instance: there is no native life on the moon; Che Guevara was not English; Angela Merkel is not identical with Joe Biden; there is no hippopotamus in this room, nor a rhinoceros, nor a painting of the Mona Lisa; electrons are not positively charged; base metals cannot be turned into gold; pigs cannot fly under their own power; there are no square circles; Leibniz did not write the *Meditations on First Philosophy*; the internet did not exist in 1965; there are no unicorns, nor griffins, nor centaurs; the Cuban missile crisis did not lead to nuclear war; and so on. How can any of these truths be made true by what there is in the world? And, if they cannot, where does that leave our understanding of truth, especially the *prima facie* plausible idea that truth depends on being?

9.2 Molnar's problematic

The revival of interest in the problem of negative truth is in large part down to George Molnar's (2000) lucid articulation of the issue. Molnar set out the problem in the form of four claims, each of which appears to be true, and which taken together entail that there must be 'positive' truthmakers for negative truth. The four claims are:

- (Mi) The world is everything that exists
- (Mii) Everything that exists is positive
- (Miii) Some negative claims about the world are true
- (Miv) Every true claim about the world is made true by something that exists.

For purposes of argument, I will accept Molnar's problematic as a way of articulating the issue. Parsons (2006) is right that there is an equivocation between (Mii) and (Miii) over the meanings of positive and negative, but I do not think it seriously affects anything that I say. I will also accept that (Mi)–(Miv) jointly entail that there are positive truthmakers for negative truths. The difficulty, then, is that Molnar's problematic remains unresolved. It compels us to pursue one of the following options. We can either (a) take up the challenge and succeed in finding the positive truthmakers of negative truths, or (b) reject one or more of the claims (Mi)–(Miv) that jointly entailed that there must be positive truthmakers for negative truths.

Molnar himself does an excellent job of exploring these options before ending with the negative conclusion that there is, as yet, no successful solution to the problem of negative truth. For the purposes of completeness, I must nevertheless summarize these difficulties while also updating the discussion with some of the subsequent proposed solutions, which I find just as unsatisfactory now as Molnar did in his day. I will look first at what I call the proposed *solutions proper*, which

are those that purport to uphold all of (Mi)–(Miv) and find positive truthmakers of negative truths. These do not succeed. Then I will look at the proposed solutions to the problem that would be based on a rejection of at least one of (Mi)–(Miv). It is this latter kind of approach which holds the most promise of dissolving Molnar's problematic and, indeed, my own preferred solution falls into this category. Of course, it matters greatly which of (Mi)–(Miv) is rejected. We will be left with very different accounts of the world depending on which is. A Parmenidean would not welcome any theory that rejected (Mii), for instance, since (Mii) states a version of the Parmenidean (Pi) principle. The power of Molnar's articulation of the problem, however, is that rejection of any of (Mi)–(Miv) seems to come at a cost. If so, then you have to choose which cost you are willing to pay, unless you can after all succeed in finding positive truthmakers for negative truths.

Before that, however, it will help to explain briefly the motivation behind each of (Mi) to (Miv), as this shows us what would be at stake in their rejection.

(Mi), 'The world is everything that exists,' evokes the opening to Wittgenstein's (1921: proposition 1) *Tractatus*. Its place in Molnar's problematic serves to prevent any tricks or dodges. We are looking for worldly truthmakers and that must mean truthmakers among everything. If it exists, it is part of the world, for these purposes. What you cannot do, then, is admit that all existents are positive (under Mii) but then suggest there is something extra in the world that serves as the truthmaker for a negative truth: something that wasn't included in our initial inventory. If it exists, it is part of the world, and there's no avoiding this. Molnar cites Meinong and the Russell of *Principles of Mathematics* as those willing to abandon (Mi). As we have seen, Meinong's thinking is more subtle than commonly supposed, but what Molnar has in mind is looking for something that doesn't quite exist—perhaps it quasi-exists—that can play a role in truthmaking. Molnar correctly rules out any such equivocation on the meaning of 'exists' as an attempt to solve the problem.

Claim (Mii), that everything that exists is positive, is the Parmenidean principle. As I argued in chapter 1, it might be possible to be less strict concerning this principle. Perhaps there can be good reasons why we have to grant the existence of a negative. A resolution to the problem of negative truths, we shall see, is one reason why some otherwise naturalistically inclined philosophers think that (Mii) has to be relinquished. But Parmenides and those who have followed have made a good case for retaining (Mii). Why should reality be split into positive and negative components? What would it mean for something to be a negative existent? How is negativity a feature of what there is rather than just a feature of the way we talk about what there is, and what there is not? We have a much neater metaphysics if we say that there are no degrees of being—that everything simply exists or it doesn't—and there is no need to posit negativity as part of the world. Of course, if it is puzzling how negativity can be a feature of what exists, perhaps it is equally puzzling how positivity can also be such a feature, as Parsons (2006: 591) says.

Once one rejects the cleavage in nature, then a positive–negative division seems to have no application. That point can be accepted without threatening the Parmenidean principle, however. Parmenides is able to state his case without applying the qualifier ‘positive’ to what exists. For stricter accuracy, Molnar could instead have stated that nothing can both exist and have a negative nature. It would then be up to anyone who rejects this claim to explain in what sense something can be both existent and negative. Instead, as stated, it would be creating an unnecessary burden on someone who holds (Mii) to justify calling an existent positive. Most of us will just want to call it existent.

(Miii), that there are some true negative claims about the world, seems undeniable, though this is not to say that it is a clear or simple matter what count as the negative claims. Often such claims will contain the words ‘no’ or ‘not’, but not necessarily. To say that ‘my pockets are empty’ or that ‘Jessica will be absent from school today’ are arguably negative claims of some variety, even though they do not use the word ‘not’. Not every case will be clear-cut. But then there are also some cases that do seem clear-cut, such as ‘there is no golden mountain’. This is a negative existential claim, and claims like this are especially problematic in respect of their truthmakers. As well as negative existentials, we should also give an account of predicate denials, such as that $\langle \text{Theaetetus is not flying} \rangle$, simple negatives such as $\langle \text{There is not a hippopotamus in the room} \rangle$, and universal quantifications too, as we have already seen, since statements like $\langle \text{All men are mortal} \rangle$ and $\langle \text{Everyone in the room is a philosopher} \rangle$ have an irreducibly negative aspect. We also need to give an account that respects the difference between external and internal negation since there is a distinction between:

(9) $\langle \text{It is not the case that } x \text{ is soluble} \rangle$,

and

(10) $\langle x \text{ is insoluble} \rangle$.

As I will argue below, there is not a single account that explains both of these. For completeness, it should also be observed that there are negative analytic ($\langle \text{Bachelors are not women} \rangle$), a priori ($\langle 2 + 2 \neq 5 \rangle$), and necessary ($\langle \text{water is not identical with } \text{H}_2\text{SO}_4 \rangle$) truths. As stated above, however, these latter kinds of statements will not be my main concern, and the reasons should now be clearer since they avoid much of Molnar’s problematic. Analytic statements, for example, if they are about the meanings of words, do not have the difficulty of trying to explain a truth in terms of what there is and what there is not. Negative analytic, a priori, and necessary truths are not entirely unproblematic, but their problems are not quite our metaphysical one.

Claim (Miv) is the key truthmaker claim, and it is a truthmaker maximalist one: all truths are made true by something that exists. This articulates what is attractive about a correspondence theory of truth but also any view along the lines that truth supervenes or depends upon or is grounded in the world. Some speak of a principle that truth supervenes on being, but this technical notion of supervenience is not really up to the job. If A supervenes on B, then there cannot be a change in A without a change in B, and if any two particulars agree in all their B-type respects then they must agree in all their A-type respects too. But when we consider truths and what in the world makes them true then we find that these must co-vary and any supervenience would therefore go both ways. The truth <dogs exist> supervenes on the existence of dogs in that were dogs to cease to exist then it would no longer be true that <dogs exist>. The opposite holds too, however, since if it ceased to be true that <dogs exist> then it would have to be the case that dogs no longer existed. Supervenience fails to capture the sort of asymmetric dependence that truthmaker theory is evoking. It is the world that makes truths true rather than the truths making what there is in the world. I will stick to the notion of dependence to name this asymmetric relation between the world and truth. Now some, as we shall see, think that there are ways of articulating the idea, that truth depends on being, in a way other than truthmaker theory, and there are some who accept the idea of truthmaking but not truthmaker maximalism: perhaps only some truths are made true by the world and some others are true by default. But as these are potential ways of resolving Molnar's problematic, we will address them in the following sections.

9.3 Solutions proper?

Here I will examine two attempted solutions to the problem that aim to resolve the matter by identifying the worldly truthmakers for negative truths. We will see, however, that they really fail to do so, and both effectively collapse into the other kind of solution, namely rejecting one of (Mi)–(Miv).

9.3.1 The exclusion solution

The first of these putative solutions seeks to identify the 'positive' real existent that acts as truthmaker for the negative truth, where it does so by being inconsistent with the negation of the (negative) truth in question. The abstract structure of such a solution is that < $\neg p$ > is true, where < $\neg p$ > is some negative statement, because there is a positive fact in the world, call it q, where q could not be the case if <p> were true. For example, suppose that we consider it true that <This pen is not blue>.

What in the world makes it true? On inspection, we see that actually the pen is (all) red. It is thus true that <This pen is red> but what really matters here is the worldly fact that the pen is red rather than the truth made true by it (Demos's original 1917 version of this solution placed the emphasis more on the truth than the truthmaker, which is why this kind of clarification is now apt).

That the pen is red certainly seems to be a positive fact or state of affairs about the world. It is something that can exist, assuming that facts are existents. And if it is indeed a fact, then it would seem to necessitate the truth of <This pen is not blue>, assuming that we successfully refer to the same pen. Here, then, looks to be our positive truthmaker for a negative truth.

Although such a solution dates back at least as far as Demos (1917), a modern version of such a view has been presented by Cheyne and Pigden (2006). They seek to broaden such a solution to cover all cases since they think that with every negative truth, its falsehood would be inconsistent with the world as it is. For example, consider the situation where it is true that:

(11) <Theaetetus is not flying>

The world as it actually is, they maintain, is such that it could not be this way were Theaetetus to be flying. Theaetetus is actually sitting, let us suppose, and to be sitting is incompatible with flying (under one's own power). The fact, *q*, of him sitting necessitates the truth of (11) because *q* cannot be true at the same time as the opposite of (11). If we negate (11), and apply double negation elimination, we get:

(12) <Theaetetus is flying>.

(12) will be false when Theaetetus is sitting (and there is indeed a truth that <Theaetetus is sitting>).

Cheyne and Pigden even think that this argument can be extended to negative existentials, such as:

(13) <There are no unicorns>.

Were it to be true that there were unicorns, then the world would have to be different from how it actually is. In other words, the existence of unicorns would be inconsistent with the way the world is, therefore, the way that the world is necessitates that there are no unicorns.

Despite an initial appearance of credibility, however, this solution fails as an attempt to provide positive truthmakers for negative truths. It might resolve into a different kind of solution, such as one that jettisons Molnar's (Mii), and we can come to that option in due course. But considered as an attempt to retain all of (Mi)–(Miv) and provide the positive truthmakers of negative truths, it cannot succeed.

The weakness of the putative solution is revealed by the name we use to label it. It is an incompatibility solution (Russell 1918: 186–9), which we might also consider a case of exclusion (Molnar 2000: 73). Such names are appropriate. In the first case, the solution alleges that there is some positive fact q that is incompatible with the negation of the negative truth: hence $\langle \neg p \rangle$ is true because q is a fact and q is incompatible with $\langle p \rangle$. Presumably this is meant to be some kind of modally serious, *de re* incompatibility given the cross-categorical nature of the truthmaking relation. If q is, $\langle p \rangle$ cannot be true. It is this modally serious relation that Molnar describes as exclusion: there is a fact that excludes a particular truth, and it is because of that exclusion that the negation must be true. If q excludes $\langle p \rangle$ then it must be that $\langle \neg p \rangle$. (Of course, this exclusion most likely would go through an intermediate stage where the fact, q , excludes another fact, p , such that when q , then it cannot be that p . If it cannot be that p , then $\langle p \rangle$, the correspondent truth value-bearer to p , cannot be true, hence the negative $\langle \neg p \rangle$ must be true.)

The first problem with this is that it is not clear that every putative negative truth is one whose negation is excluded by something that exists, as Parsons (2006) pointed out in response to Cheyne and Pigden. A key difficulty concerns the negative existentials. Certainly, a world in which there are unicorns would be different from a world in which there are no unicorns, which is where the appeal of their solution lies. But the contentious point is whether the world as it actually is *excludes* the existence of unicorns. It looks like it doesn't. There seems to be nothing in this world that is actually preventing the existence of unicorns. It is just that they don't exist. As best as we know it is merely a contingent truth that, given how the world is, there are no unicorns. Parsons makes the point in terms of possible worlds. Our actual world is a possible world in which there are no unicorns. But, as far as we can tell, there could be another world that is exactly like ours, with all the same existents or their counterparts, plus unicorns as well. This argument invokes Russell's (1918: 206–8) old case for general facts: all that there is does not entail that it is all that there is. There could be all that there is plus one or more things besides. We will resume this thread of argumentation shortly, as it is at the heart of the other proposed account of truthmakers for negative truths.

There is a second and likely more damaging objection to an incompatibility solution. This objection concentrates on the nature of incompatibility or exclusion. When one thing excludes another, or is incompatible with it, this can reasonably be taken to mean that those two things cannot both be at the same time. If A is incompatible with B , then it cannot be the case that both A and B . There are two ways that this can be interpreted. One could take such exclusion to be purely logico-linguistic and the other takes exclusion to be metaphysical. The logico-linguistic interpretation can be taken as meaning that A and B cannot both be true or: $\langle \neg(A \ \& \ B) \rangle$. If incompatibility means this, then we can see that it fails as an attempt to provide positive truthmakers for negative truths as it evidently

involves nothing more than swapping one negative truth for another. Explicitly, we wanted to know what made it true that $\langle \neg p \rangle$ and were told that it was made true by q . But then we saw that q made $\langle \neg p \rangle$ true only because it was also true that $\langle \neg(p \ \& \ q) \rangle$. So then we are left with the question of what makes $\langle \neg(p \ \& \ q) \rangle$ true. If the same kind of answer is given—that there is some other fact r , and $\langle \neg(r \ \& \ (p \ \& \ q)) \rangle$, then it is clear that a regress is underway. If there is a regress, then we have not really provided the positive truthmaker for the (original) negative truth. This problem was identified by Price (1929: 106), Taylor (1952: 449), Buchdahl (1961: 164), and Prior (1967: 459). However, Mabbott (1929: 79), Raju (1941: 595), Heinemann (1944: 151–3), and Russell (1948: 122) all argue that incompatibility is not a negative. As Gale (1976: 19ff) says, one cannot resolve this debate until one has established what it would be for something to count as positive or negative. I have argued (2007a) that this is ultimately a metaphysical judgement that cannot be discerned from logic or language alone.

The metaphysical interpretation of exclusion takes it to be a real feature of the world. There are some credible cases in which one state of affairs excludes another, the case of determinates falling under the same determinable being the best example. If an object instantiates the shape of being spherical, then it cannot at the same time be a cube. This is, no doubt, because being spherical and being cubed are two determinate ways that something can have the determinable property of having shape, where shape is a property at a higher level of abstraction. How and why determinates exclude each other when they fall under the same determinable is a matter of which I need not give a complete account (even if I could), but it seems that it relates in some way to the notions of being and becoming. According to one, not implausible account of how determinables and determinates relate (Armstrong 1978b: 111–13), something can have a determinable property only by having one of its determinates. The way for a particular to have shape is for it to be cubed, or spherical, or pyramidoid, and so on. Being any one of these entails having shape. And while having shape entails having one of those determinate shapes, there is no entailment which one it is. In becoming a determinate shape, an object gains specificity. It becomes one thing and cannot then be a different shape (at the same time) than it is. There is, therefore, at least some coherent account of exclusion in the case of determinates under the same determinable. But does the exclusion of one truth by a fact have this same kind of structure? Possibly it does if we are willing to say that the world could be several specific ways something is, and if it determinately becomes one of those ways, then it cannot be another specific way. Certainly, such an interpretation is consistent with one example I gave for this, where the pen being entirely red excludes the pen being blue, and hence is a truthmaker of $\langle \text{this pen is not blue} \rangle$. But there were other cases of negative truth where an account in terms of determinable and determinates has a less plausible application, for example with $\langle \text{there is not a hippopotamus in the room} \rangle$ and $\langle \text{there are no unicorns} \rangle$. If there is something

about the room that excludes the presence of a hippopotamus, or something about the state of the world that excludes unicorns, then we had better know about it. Whatever it is, it seems as if it will have to introduce a worldly exclusion: a kind of negative existent. I will say no more about this, other than that it seems to be effectively a rejection of Molnar's point M(ii). It is up to those who wish to take that line to explain what a worldly exclusion is and why dropping (Mii) is then the best solution.

A general lesson to be taken from the above is that it seems hard to account for the negativity of a truth without introducing a negative element elsewhere. As one more example, consider Barker and Jago's (2012) account of anti-instantiation. Barker and Jago permit the reality of negative facts explained in terms of anti-instantiation. For example,

(14) <The lake is not frozen>,

is made true by the fact of the lake anti-instantiating frozenness. They tell us very little about what this anti-instantiation is or how it can be in the world. They justify use of the term rather by what work it can do, including the work of providing truthmakers for negative truths. Whatever it is, though, it certainly looks like a negative component in reality, violating Molnar's (Mii). In addition, it seems that this is weak as an account of negative facts because, without saying more of what anti-instantiation is, it looks as if it is defined as whatever it is that accounts for negative facts: the name for a solution to a problem rather than a solution.

9.3.2 The totality solution

There is a second kind of account for the truthmakers of negative truth that suffers from the same problem: that it effectively works only if it relinquishes Molnar's (Mii). This is the totality fact solution of Armstrong (2004: 54–60). Totality facts have been mentioned earlier, but the point here concerns their use in a solution to the problem of negative truth. The proposal begins by acknowledging a matter that we encountered above; namely that no total ever entails that it is a total.

Armstrong proposes to provide the truthmaker for a negative truth in the following way. Let us take (8) again as our example negative truth:

(8) <There is not a hippopotamus in the room>.

There will be lots of facts about the room in question, for instance that it contains a table and chairs, that it is of certain dimension, that there are two philosophers in it, and so on. I will assume it established that the Cheyne and Pigden account is

wrong in that there is nothing about this collection of facts that excludes the presence of a hippopotamus. Certainly, Armstrong, and Russell before him, accept this. There could be all the actual facts about the room but also an additional fact that there is a hippopotamus inside it. That being the case, the actual facts do not necessitate that there is no hippopotamus in the room so cannot make (8) true. (Armstrong, as several other truthmaker theorists do, take it that the truthmaking relation is one of necessitation: the worldly fact must necessitate the truth of the truth-bearer.)

The best way to get the requisite necessity, according to Armstrong, is to accept that as well as all the first-order facts about the room, there is also a higher-order fact of totality about it. The totality fact totals the first-order facts, and it is higher order in that it is a fact about the facts: namely, the fact that they form a total. Hence, if we have a set, Σ , of all the first-order facts about the room, where $\Sigma = \{F_1, F_2, F_3, \dots, F_n\}$, plus Σ^T , the totality fact that Σ is all the first-order facts, then we get the necessitation from Σ plus Σ^T that \langle There is not a hippopotamus in the room \rangle .

There are two difficulties with this as a solution to Molnar's problematic. The first, Armstrong acknowledges. The totality fact Σ^T might be part of reality, even though it is a higher-order part of reality, but it is inherently negative. It says that there are no more first-order facts than those gathered in Σ . Armstrong concedes, then, that the totality fact is a negative fact. The totality solution thus relinquishes (Mii). But Armstrong claims at least a gain in economy over Russell's version of the solution. Russell also reluctantly allowed negative facts to account for negative truths but in a way that seemed to require one negative fact for every truth. Hence, there would be separate facts of there being no hippopotamus in the room, no rhinoceros in the room, no Abraham Lincoln, no field of tulips, no ghosts, no woolly mammoth, and so on. Given the enormous, if not infinite, number of negative truths about the room, then the world would be heavily populated with negative facts making them true. Armstrong's quantitative gain in economy is that with the addition of a single negative fact, Σ^T , plus all the first-order positive facts that exist unproblematically, we can get all the negative truths without any further addition. Given Σ , plus the fact that there are no more facts than Σ , then every negative truth follows. A single negative fact does all the required work, and we need not have a world containing more negative facts than positive ones.

If Armstrong's solution succeeds, then, there is at least some mitigation in its favour. It might be that every solution has its price and our task would then be to identify the least costly solution. Armstrong can claim a significant advantage over a solution that permits more negative facts than necessary, although he acknowledges that such a solution was easily within Russell's grasp, since he had already allowed general facts, which could do much the same work as totality facts. But there is a second concern about Armstrong's solution that would

undermine its claim to be a viable solution in the first place. The worry is that we get the requisite necessitation of the negative truth in question only if Σ does not already contain the corresponding positive fact that, were it to exist, would render the negative truth in question false. In other words, saying that Σ , plus Σ^T , necessitates that \langle There is no hippopotamus in the room \rangle is true, works only if we specify in advance that Σ does not include the fact that there is a hippopotamus in the room. Σ being all the facts about the room, on its own, does not guarantee that Σ excludes this fact. We must allow in advance of adding the totality fact Σ^T , that the facts included in Σ do not contain a specific one: the one that would make (8) false, hence the corresponding positive claim of (8) true. And it is clear that we would need to make this sort of qualification in respect of every negative truth that is supposed to be made true by the addition of Σ^T . What this then suggests, is that it is not Σ^T that is really doing the work of making all these negative truths true. It is the negative fact smuggled in to Σ , that means that it does not contain what would make (8), or whatever negative claim, false, that is doing the work. Given this problem, it then seems that even the claim that the totality solution is more economical than just permitting as many negative facts as needed, fails to hold. To get the requisite truths, we need a negative fact that Σ does not contain a fact of there being a hippopotamus in the room, and it does not contain the fact of there being a woolly mammoth in the room, and it does not contain the fact of there being Abraham Lincoln in the room, and so on. The economic advantage of a totality solution is, then, suspect.

9.4 Attempted solutions that reject at least one of (Mi)–(Miv)

Given the preceding discussion, there are grounds to doubt that we can find positive truthmakers for negative truths. Perhaps someone will succeed eventually, but there might be a good reason for scepticism about this. As Martin (1996: 61) said, 'In general, the attempt to account for the absence of things in terms of the presences of other things is hard to accomplish.' What, then, of rejecting one of Molnar's (Mi)–(Miv)? After all, if we have a seemingly irresolvable problem, perhaps it has been created by an erroneous or unhelpful initial set-up. As Wittgenstein (1953: §309) suggested, our task should then be showing the fly the way out of the bottle: identifying which commitment generated the problem in the first place.

Recall, in that case, Molnar's four propositions:

- (Mi) The world is everything that exists
- (Mii) Everything that exists is positive
- (Miii) Some negative claims about the world are true
- (Miv) Every true claim about the world is made true by something that exists.

Can we reject any of them?

Again, it must be admitted that there is a cost in doing so. To reject (Mi) would be to take some kind of quasi-Meinongian route, allowing that there was something that didn't quite exist but was nevertheless able to make a negative truth true. This approach faces the huge problems of, first, explaining how there is something that is part of reality and yet nevertheless not fully existent (I refer readers to the discussion of Meinong's view in chapter 8). Such a move would come up against the major presupposition of naturalistic metaphysics: that existence is univocal or, in other words, that there are no degrees of being (Parmenides' Pii). Second, there would also be the not inconsiderable task of explaining how this not-quite-existent thing was somehow able to enter into a relation with something that existed in order to make it true. So one way that such a 'solution' might go would be to say that the non-existent hippo was nevertheless 'some kind' of thing and that it was in the room and this is what makes (8) true. I will leave it to supporters of such a view, if there are any, to provide the details. Suffice it to say that I do not take this to be promising as a solution and Molnar's (Mi) should be retained.

As we have seen, one common way out, albeit reluctantly, is to drop Molnar's (Mii). In chapter 1, I explained that the Soft Parmenidean was someone who resisted allowing negative entities or components into their ontology but was nevertheless willing to countenance the possibility of doing so if certain conditions were met. It will be recalled (from chapter 1) that those conditions were that the entity was irreducibly negative and indispensable to our best theory. Accounting for negative truth could, then, be exactly such a case. Our best theory of truth, it might be said, is one in which truths are grounded or dependent upon the way the world is and, therefore, if some negative claims are true about the world, then they must also be grounded in this kind of way. And that about the world in which they must be grounded seems to be irreducibly negative. The attempts to ground them in something positive appear to have failed, the most serious proposal being in terms of incompatibility or exclusion. But we saw that facts of incompatibility or exclusion are themselves irreducibly negative. There is a negative component to them (that it cannot be both A and B) that cannot be eliminated or treated in entirely positive terms. It looks, therefore, like the soft Parmenidean's two conditions for granting the existence of a negative entity are met. There are reasons, then, to drop (Mii), notwithstanding the intuitive resistance to doing so.

The soft Parmenidean is well aware of how problematic it is to jettison (Mii) but will countenance it, especially if there seems no more viable solution to the problem of negative truth. The price is a schism in reality: a cleavage between positive and negative existents. The proponent of such a solution will have to tell us what a negative existent is, for instance. It is not inconceivable that someone will attempt this, especially as some see an independent motivation for believing

in the reality of absences and other negative entities. Martin (1996) takes absences and lacks of things as real, for example, while not treating them as things. They are ‘localized states of the world or universe and, therefore, though not things or natural properties of relations of things, they can serve as truth-makers for negative existentials or false-makers for positive existentials’ (1996: 58). For example, ‘A void is not a *thing*, but it may be *how a space-time region* is’ (1996: 62). Hence, it could be a localized state of the region occupied by the room that it lacks a hippopotamus, or a hippopotamus is absent from it, and this is a real state of affairs even though it is inherently negative. Absences also figure in Lewis’s (1992) account of negative truths since he says that they are true because of an absence of falsemakers. Hence, (8) is true since there is nothing that makes it false: specifically, no hippopotamus in the room. Like Martin’s, Lewis’s account requires taking absences seriously. It is also clear, however, that there are many remaining battles. There are still compelling reasons to think that absences cannot be real. They are causally inert and imperceptible, to name just two. But we have also seen that these claims are themselves controversial, since some think that absences can after all be causes and can be seen. We can certainly say that the rejection of (Mii) is arguable, even if there would be work to be done in its defence.

Are there possible grounds for rejecting (Miii), that some negative claims about the world are true? As stated before, it is *prima facie* self-evident that some negative claims are true, for example (8), and there is surely an endless supply of manifestly true negative claims. It is perhaps for this reason that this potential solution to the problem is historically the most underexplored. But it is not completely without advocates. Most recently, I tried to show how someone serious about truthmaker theory had a reason to treat putative negative truths as falsehoods instead (Mumford 2007a). If one believes that truth requires a truthmaker, and thus in claiming something to be true one is ontologically committing to the existence of such a truthmaker, then a natural response to the apparent failures to find positive truthmakers for negative truths is to say that they are not, after all, true. We would be aided in this project if we could substitute for each putative negative truth something else that is not true. The 2007a proposal was that we can utilize the notion of falsehood, where a statement is false just in case there is no truthmaker for it. Falsehoods require neither truthmakers nor falsemakers, since they are not dependent on what there is but on what there is not (with a soft Parmenidean resolve that what is not is not a part of what is). Saying that something is false, then, gives no direct ontological commitment to anything in the world. Hence, instead of asserting it to be true that (13), <There are no unicorns>, one can instead say that it is false that (13*) <There are unicorns>, and instead of asserting it is true that (8), <There is not a hippopotamus in the room>, one can say it is false that (8*): <There is a hippopotamus in the room>. It requires the existence of nothing at all for it to be false that there are unicorns. For it to be false that there is a hippo in the room, it might require the existence of the room at

most but there need be no hippo-like existent or subsistent of any kind for it to be false and, arguably, it should still be false even if there is no room. Isn't it false that Syd Barrett's fourth solo album contained ten tracks, given that Syd Barrett did not have a fourth solo album?

Cheyne and Pigden call this a 'desperate remedy' (2006: 251) so they might be surprised to learn that the strategy is not quite as rare as one might think. As Horn (2001: 56ff) details, treating negation as falsity, or reducing it into falsity, is found in Wood (1933: 422), Baldwin (1928: 147), Sigwart (1895: 122), Bradley (1883: 118) and even Leibniz, who said that 'If B is a proposition, not-B is the same as that B is false' (Leibniz 1686: 58). Russell too developed this line in one place. Perhaps unhappy with his earlier commitment to negative facts, we see this in his later work: 'When you say "this is not cheese" you mean "the statement 'this is cheese' is false"' (Russell 1940: 74).

However, there is a difficulty with this approach, which has been pointed out by Frege (1919), Austin (1950: 128f), Quine (1951: 27f), and Geach (1980: 76), which is that the identification of *not* and *false* is based upon a conflation of object language and metalanguage (Horn 2001: 58). Affirmation and negation seem both to be about things whereas, when we say that p is false, ($f < p >$), we are making a metalinguistic statement about a proposition. But $< \text{The door is not red} >$ is just as much about the door as is $< \text{The door is red} >$. It is contentious that simply in using the word 'not', one is switching from speaking about the door to speaking about a proposition. The same would be said of affirmative statements, of course, since there is an apparent difference between 'the door is red' and ' $< \text{the door is red} >$ is true.' Use of the angled brackets shows this difference. The first statement is about the door, whereas the second statement is a metalinguistic statement about the first statement. As Geach points out, the distinction can also be shown clearly for the case of non-declaratives, since "Do not open the door!" is a command on the same level as "Open the door!" and does not mean "Let the statement that you open the door be false!" (Geach 1980: 76). The sort of problem is, indeed, acknowledged by Russell, since he admits that someone who can speak only in an object language will be able to tell you everything that is in the larder but will be unable to tell you what is not in the larder. The latter requires, on Russell's view, switching to the metalanguage so that you can say it is false that certain things are in the larder.

This line of objection might not be completely devastating to the prospects of dropping (Miii), however. Someone supporting it could right away declare that they are not attempting a reduction of *not* to *false*, nor even an elimination of *not*. One could concede that negation is logico-linguistically ineliminable. But the aim of this approach is primarily metaphysical and even if we cannot eliminate negation from language, we can still attempt to eliminate negative entities from our ontology. Nor would such an account be an analysis of the meaning or use of the word *not*. And I, along with many others, agree with Armstrong (1978b: 9) that

we should reject the argument from meaning. There need not be a property corresponding to every predicate, so nor need there be some entity corresponding to every negative predicate, term, or statement. Rather, in the spirit of showing the fly the way out of the bottle, we can explain how Molnar's problematic—basically the paradox of negative judgement—has arisen. It may be that negation has taken on a life of its own even if it was originally used to express denial and may allow us to express things that cannot be equated with falsity or denial. But this still does not show that negation is prior to denial nor that there is a negative element in reality that corresponds to negation. I will develop such an account in the chapter 10 and take it that rejection of Molnar's (Miii) is still a live option.

The remaining alternative is to drop (Miv): that every true claim about the world is made true by something that exists. This final proposition is one that has come under much scrutiny, as a result of which it is now common that (Miv) is the one sacrificed in order to retain negative truth. One reason for this is that while many see the attraction of truthmaker theory, not all truthmaker theorists are truthmaker maximalists (for example, Simons 2005). Truthmaker maximalism is the claim that all truths have truthmakers: every truth is made true by something that exists. But a possible position is that there are truthmakers for some of the truths but not all of them. Negative truths might be true by default, for instance. Affirmations, given that they concern what is, need truthmakers, since there has to be something making them true, whereas negative truths, which concern what is not, require nothing in order to be true. Perhaps the only condition would be an absence of falsemakers, as Lewis stated.

The main difficulty with this as a solution is whether there is a principled reason for saying that the problematic statements do not need truthmakers in order to be true. Is the move *ad hoc*? As soon as we find a difficult case, in which it seems the truthmakers for the truths are not to be found, then the 'solution' is to drop the requirement of having a truthmaker. But without a good and principled reason why this class of truths has to be different from the rest of them, then it looks as if dropping the truthmaker requirement is motivated only by the fact that it doesn't seem to work here. If the requirement to have a truthmaker can be dropped so easily, then it probably also undermines the general motivation for having truthmakers in the first place. The move concedes that there can be truth without a truthmaker in which case it might be reasonable to wonder whether we should have truthmakers at all. This is a view that Dodd (2002) has advanced, treating the issue of negative truth as make-or-break for truthmaker theory. Others are more sanguine, for instance Parsons (2006), who still sees a motivation for truthmaker theory even if it fails for the case of negative existentials.

Dodd thinks there are principled reasons to drop truthmaker theory and one can do so while still maintaining that truth is asymmetrically grounded in being (Dodd 2007: 400). One can thereby uphold the original motivation behind truthmaker theory but without adhering to the letter of it, since it has failed in the case

of negative truth. Dodd presents an alternative framework to explain how there can be asymmetric grounding of truth in being without truthmakers: the identity of the truth is determined by the identity of its ground. But he still maintains that truth depends on '*how things stand* with what exists' (Dodd 2002: 74), and explains how negative truths depend on being in the following way:

A negative existential, such as <There are no arctic penguins> is true not because some thing or things exists, but because there are no things of a certain kind (arctic penguins, in this case). Likewise, a negative inessential predication, such as <This liquid is odourless>, is true not because something exists that necessitates its truth, but because the liquid in question lacks the property of having an odour. (Dodd 2007: 400)

This shows that the problem is not really solved but, rather, skirted over in Dodd's approach. *How things stand* might be different from Armstrongian states-of-affairs, but truthmaker theory is surely not just a matter of what, specifically, are the sorts of things that play the truthmaker role. Whether a truthmaker is a fact, a state of affairs, or how things stand, is by the by, really. The key question is whether the truthmaker, whatever it is, has to be of a negative kind in order to account for negative truth. The quotation from Dodd shows that, on his account, the being that grounds the truths in question—being that I think could perfectly well be called the candidate truthmaker—is indeed negative. In the first case it is that there are no things of a certain kind. In the second case it is that something is lacking a property. Both of these are negatives, which might mean that Dodd's account is really one that gives up (Mii) after all.

9.5 A more recent solution

There have been several recent attempts to solve the problem of negative truth. I shall not consider them all, since such a catalogue would soon become tiresome to the reader. Instead, I will present just one more attempt to solve the problem, chosen because it purports to contain a new innovation in the debate but also because it eventually runs aground on some familiar rocks. Griffith (2015) professes to rise to the challenge of finding truthmakers for negative truths, arguing that negative truths are indeed made true to the extent that truth supervenes on being. Use is made of a distinction in Hofmann and Horvath (2008) between existence dependence and variation dependence. The truth that <dogs exist> is dependent on the existence of dogs. For any entity *x* and proposition *p*, *p* is existentially independent of *x* when any of the following four possibilities can occur:

- (a) x exists and p is true
- (b) x exists and p is false
- (c) x does not exist and p is true
- (d) x does not exist and p is false.

Hence, the existence or non-existence of x makes no difference to the truth or falsehood of p. For example, <dogs exist> is existentially independent of there being other planets.

The problem of negative truths can be understood in terms of existence dependence to the extent that the truth <there are no unicorns> does not seem to depend on anything in reality, therefore it *prima facie* violates the principle that truth is dependent upon being.

However, Griffith urges that negative truths do depend on reality in a significant sense because there is also a notion of variation dependence/independence to be considered. That <there are no unicorns> does not depend on reality existentially but does depend on reality in the sense that, were things different, then it might no longer be true; namely, if unicorns existed. Griffith uses this finding to deny the claim that a truth such as <there are no unicorns> ‘has no connection to reality’ or that it’s ‘truth “floats free” of being’ (Griffith 2015: 324). Positive truths will be both existentially and variation-dependent on being, while negative truths will be variation-dependent alone. But variation dependence is enough.

However, while this notion of variation dependence/independence seems coherent and applicable to some things in reality, there are reasons to doubt that it does the exact work that Griffith wants. He states his solution thus:

Even though a truth like (1) [<there are no unicorns>] is not necessitated by the existence of any entity, there are plainly ways in which reality could vary that would change (1)’s truth-value, to wit: if a unicorn came to exist (and the other existing entities came to stand in new relations to the unicorn), then (1) would be false. (Griffith 2015: 325)

And again later, much the same point is made in the same kind of way: ‘that the truth-value of a negative truth would change with certain possible variations or changes in reality (in particular with changes in what exists)’ and ‘*Only* a change in being can render a change in a negative truth p, i.e. a change in what exists *alone*’ (both Griffith 2015: 325). But this way of putting it obscures the fact that it is a very specific change that would be required to make a negative truth false. The change in being has to be a change from what there is not to what there is. It is significant that the change be of this kind since it has surely to be conceded that while <there are no unicorns> is true, there is nothing in being upon which it depends. Certainly, if it became false because unicorns existed, then there would

be something real upon which its falsehood depended. But it seems that variation dependence should not be considered strong enough a ground for saying that negative truths have truthmakers because it does not adequately redeem the principle of truth depending upon being. A contingent negative existential depends upon what there is not. If it no longer was the case that something was not, then of course the contingent negative existential would no longer be true. But that does not justify a conclusion that its truth depends on what there is, even if its falsehood would.

We can see that this point is hidden by statements such as the following:

All the present account says is that the world is what exists and that changes in what exists would be responsible for changes in the truth-values of certain negative existentials. (Griffith 2015: 329)

This sounds reasonable, but it is clear that the specific change we require, one in which a true negative existential will become false, must be a change from something not existing to that thing existing. Hence, the negative existential's truth depends on something not existing; that is, not being part of reality. It depends on non-being. Subsequent claims that only acceptable entities are involved as truthmakers do not wash. We can accept entities such as the world, the Arctic, and a fridge when considering truths such as <there are no unicorns>, <there are no Arctic penguins>, and <the fridge is empty>, but those entities listed are not enough to account for the truths in question. We would also have to invoke the emptiness of the fridge—not just the fridge itself—and absences of unicorns and penguins from the Arctic. To do so looks like it lapses back into reifying absences and nothingnesses to serve as truthmakers. Griffith claims that he is retaining all of Molnar's (Mi)–(Miv). For reasons given above, it is more accurate to say that he is rejecting (Miv): that every true claim about the world is made true by something that exists. He argues that he has shown 'is made true by' can be interpreted in different ways. But, I suggest, it is vital for his view that, when true, negative existentials have to be made true by something that is not. He cannot claim, then, to have found positive truthmakers for negative truths.

9.6 Impasse?

Ideally, an examination of a philosophical topic will set up the problem, explain why previously offered solutions have shortcomings, and then provide a new solution that avoids all the pitfalls. We could then declare that the problem is solved. Unfortunately, we are not yet in a position to complete this task in the present case. The problem of negative truth appears very real and still stands. We have seen that we cannot find positive truthmakers for negative truths, although

Molnar's four propositions entail that they must exist. Nor can we see any obvious ways in which one of the four propositions can be rejected. It looks like we have reached a philosophical impasse, therefore, which was indeed how Molnar himself concluded examination of the topic. Despite arriving at much the same conclusion, I hope in this chapter to have made some further contribution to the discussion, providing additional detail and analysis of the difficulty as well as considering some of the more recent treatments.

However, the matter is not closed, and I am going to consider the key topic of denial in more detail in the next chapter. I see this as a potential help in understanding how we have arrived at the problem of negative truth. It will also be ventured that a correct understanding of the role of denial unlocks not just the problem of negative truth but other problems that we have considered earlier, such as negative properties and causation by absence. What we find when we consider denial, however, is not a straightforward, outright solution of these difficulties. As I have said, we are in the position of the fly trapped in the bottle. There might be no direct route out of it, but what I can do is identify the philosophical positions that led us into the trap. This might simply illustrate the dictum that in philosophy, as in politics, it is possible to create insoluble problems.

Negation and Denial

10.1 Partial success?

Philosophy is at its most satisfying when it articulates a problem clearly and then provides a neat solution: one that is demonstrably better than its rivals and ties up all the loose ends. Unfortunately, it does not look like we can have this sort of unconditional success in the case of negative attributions. In the previous chapter, we looked at the paradox of negative judgement, now called the problem of the truthmakers for negative truths. We could find no such positive truthmakers, nor judge easily which of the premises that produced the problem was mistaken.

I reaffirm here that no easy solution is possible and, although I am no major advocate of the approach to philosophy taken by the late-period Wittgenstein, our ambition might have to be limited to showing the fly the way out of the bottle; or, perhaps even less than that, of explaining how the fly got into the bottle in the first place. Negative truth might be an insoluble problem. This suggestion is not like that of mysterianism regarding the nature of consciousness (e.g. McGinn 1991), though. That is the view that it cannot be explained in principle how consciousness arises from physical matter. The various problems created by negation might be perfectly understandable. One might be able to identify the exact moves that were made that led us to our current difficulties, just as a historian could explain the origins of the political problems in the Middle East. But understanding exactly how a problem arose still does not mean that it is possible to solve it. Such problems may have resulted from irreversible changes. Such is my suggestion concerning the role of negation.

In this chapter I will argue that there is a perfectly natural way to understand negation in terms of denial and, contrary to claims that such an account is untenable or unhelpful, I will argue that this provides a dissolution of the problem of simple negative truth. I argue that the account is not counterintuitive, nor a desperate measure. But I will concede that it does not solve every problem of negative truth since there is an advanced problem, which concerns negations in logical complexes. Negation remains logico-linguistically irreducible, I concede, and in certain uses there is no clear recasting in terms of denial.

Nevertheless, there is some purpose to this limited success within the context of our soft Parmenidean project. We can see that basic negative truths can be understood not as affirmations but as denials, where denial contains no commitment to any truthmaker nor, really, to any truth. The problem, we will see, is when

those denials are converted into negations which then figure in logical complexes. We cannot then be clear what, if anything, is being denied, and the convention of treating denial in terms of truthful affirmation leaves us with an irresolvable problem. I hope that the attraction of this qualified solution will become clear once we examine the detail, starting with an explanation of what I call the equivalence thesis.

10.2 The equivalence thesis

Part of the explanation of how we have a problem of negative truth is related to the view we evidently hold on the connection between affirmation and denial. It is undeniable that we have an inclination to affirm negatives, such as when we say that the ball is not red, that there is not a hippopotamus in the room, that there are no more cakes left, and so on (recall Hommen's claim to this effect, back in 2.2.3, above). Molnar's proposition (Miii) enshrines this inclination to affirm a negative as the thesis that there are negative truths. The basis of the practice can be questioned, however. We can consider whether we *need* to affirm negatives and, separately, whether we *ought* to affirm them.

Our thinking is subject to what I call an affirmation bias. This plays on the notion of confirmation bias: a tendency to look for evidence that confirms one's beliefs and ignore evidence that conflicts with them. Affirmation bias, I say, is a tendency or preference towards affirmation rather than denial, even where it inclines us to affirm a negative rather than deny a positive.

Frege (1879, §2), in the *Begriffsschrift*, had an assertion or judgement stroke, \vdash , that could appear before a proposition to indicate that the proposition was asserted or affirmed. The stroke was needed, he thought, because there is a difference between asserting a proposition and merely entertaining it. One can consider the conjecture that the world will end tomorrow without actually believing or asserting it. Frege's view, at least according to Geach (1980: 304ff), was that the assertion stroke played an extra propositional role. It could not be part of the proposition itself, since you cannot say that $\langle P \rangle$ is true simply by saying $\langle P \text{ is true} \rangle$ because there could always be the further question of whether $\langle P \text{ is true} \rangle$ is itself true. What must be asserted, then, is $\langle P \rangle$; that is *assert*: $\langle P \rangle$ rather than $\langle \text{assert-}P \rangle$. I am again using the convention of angled brackets to isolate the truth-value bearers, which some think are propositions.

If Frege is right, and there is a difference between asserting $\langle P \rangle$ and merely entertaining the proposition that $\langle P \rangle$, it raises the question of why there isn't a corresponding denial stroke, such as \neg , since there would be, presumably, a difference between merely entertaining the proposition that $\langle P \rangle$ and denying that $\langle P \rangle$ (and also, between denying a proposition and affirming one). After all, aren't assertion and denial very different attitudes one can take towards a proposition,

or forces with which one can express a proposition? Frege considers therefore, whether there are two acts of judging: affirmative and the other, which he calls negative (Frege 1919: 129). Partly motivated by the requirement of keeping the primitives of a language or logical system down to a minimum necessary (p. 130), he decides that he does not need a separate denial stroke in addition to an assertion stroke, since 'Judgment is the same act whether the answer to the question is affirmative or negative' (p. 129). There then comes a key move:

Thus for every thought there is a contradictory thought; we acknowledge the falsity of a thought by admitting the truth of its contradictory. The sentence that expresses the contradictory thought is formed from the expression of the original thought by means of a negative word.

(Frege 1919: 131, though with the argument stretching over pp. 129–31)

In short, the view is that we do not need a separate force of denial, since instead of denying $\langle P \rangle$, we can simply assert $\langle \text{not-}P \rangle$. To do so is to accept the equivalence thesis, which is to hold the following two claims:

Equiv.1. Assertion: $\langle \neg P \rangle \equiv$ Denial: $\langle P \rangle$

Equiv.2. Assertion: $\langle P \rangle \equiv$ Denial: $\langle \neg P \rangle$

Commitment to the equivalence thesis seems to be common. It is in play in the standard view expressed by Molnar's (Miii). Modern logic has dropped Frege's assertion stroke, treating it as effectively redundant, but this is not because any separate or irreducible notion of denial has a place. Not only can we say that asserting a negation raises no eyebrows, but that there seems to be a tacit assumption that it is proper to do so rather than have any separate force of denial attaching to a proposition. This will work only by acceptance of the equivalence thesis, which allows a denial of $\langle P \rangle$ to be reduced to assertion of $\langle \text{not-}P \rangle$, supposedly without any cost or loss of meaning at all. Witness, for example, Quine's articulation of the equivalence:

The peculiarity of *statements* which sets them apart from other linguistic forms is that they admit of truth and falsity, and hence may be significantly affirmed and denied. To deny a statement is to affirm another statement, known as the *negation* or *contradictory* of the first. To deny 'The Taj Mahal is white' is to affirm 'The Taj Mahal is not white'. (Quine 1974: 9)

The equivalence thesis tells us that a denial *is* an assertion. It is an assertion of a negation, but an assertion nonetheless. You might object that use of the word 'not' seems *prima facie* to express denial, so how can one ignore that and treat it as an

assertion instead? But Frege dismisses this, saying that the word 'not' is part of the predicate, so saying 'x is not-P' seems, for Frege, like any other assertion, as when we say a man is 'uncelebrated' (Frege 1919: 131).

The equivalence thesis thus allows the possibility of a language consisting entirely of assertion for expression of its declaratives. Perhaps a user-friendly language will contain imperatives, optatives, interrogatives, and exclamations besides, but it need not contain denials in addition to assertions. For anything that you would wish to deny, simply assert the negation instead. If the equivalence thesis is true, then to add another force, of denial, would double the number of force indicators in the language unnecessarily.

Perhaps it might be thought that since the equivalence thesis offers a symmetrical relationship between assertion and denial, the reduction can go both ways, so it should not be assumed automatically that denial is to be reduced to assertion. The equivalence thesis would also allow a declarative language consisting solely of denial: a possibility that Wittgenstein (1921) entertained in the *Tractatus* (§4.062). This would mean that instead of asserting $\langle P \rangle$, one would deny $\langle \text{not-}P \rangle$ (and instead of asserting $\langle \text{not-}P \rangle$, one would deny $\langle P \rangle$). Dummett instantly rejected this suggestion, though: 'This possibility has only to be stated to be recognised as spurious' (1973: 318). But there is no further direct argument as to why a language consisting solely of assertion is acceptable but one consisting solely of denial is unacceptable. This could be understood, therefore, as Dummett, and Frege before him, manifesting affirmation bias.

However, all is not quite so simple because Dummett offered another argument that complicates the matter somewhat. We have been assuming that there is a difference between denial being reduced to assertion and assertion being reduced to denial, such that a language consisting solely of denial and another consisting solely of assertion would be discernible. Considering the hypothetical example of a previously unknown linguistic community, Dummett concluded that this is not the case:

For instance, if 'A' is an abbreviation of a sentence of theirs, and it is noticed that they utter the expression 'fA' principally when it is raining, then if 'f' is taken as a sign of assertion, 'A' will be guessed to mean 'It is raining'; but if 'f' is taken as a sign of denial, 'A' will be guessed to mean 'It is dry'. But it is clear that there is absolutely no way of deciding, by observing what they say and do, between the two interpretations of 'f'. (Dummett 1973: 318)

From this he concludes: 'The two alternative descriptions of their language are completely interchangeable: they do not describe distinct possibilities at all' (Dummett 1973: 319). This throws a sceptical spanner in the works, for if a language consisting solely of assertion is no different from a language consisting solely of denial, then we are not in a position to side with Frege in thinking that

his sign *F* indicates assertion rather than denial. Of course, the same reasoning could also be used in support of Frege's claim that there is only one kind of judgement to be made, if there is no real difference between assertion and denial, which would be why it is better to call *F* simply the judgement stroke.

Dummett's argument might seem like another dead end, but what I want to highlight is that this sceptical conclusion itself rests upon acceptance of the equivalence thesis. Specifically, if we cannot treat every denial as if it is an affirmation of a negative, nor vice versa, then the sceptical conclusion would not go through that the two possible languages were indiscernible with respect to affirmation and denial.

An answer to Dummett's challenge is not our main goal, however. The immediate reason for our interest in denial was that instead of asserting a negative truth, not-*P*, we might instead deny the positive truth, *P*. It is potentially significant, if we can be taken to be doing that, since a denial, presumably, needs no truthmaker: a claim that I aim to vindicate below. Perhaps it does not need a falsemaker, either, since a denial seems to be making no commitment concerning how the world is. At most, it is making a commitment concerning how the world is not, and it is fair to assume that no truthmaker is needed for that. Further, it might also be reasonable to suggest that 'not' is the word we have standardly come to use to indicate denial. When one says 'not-*P*', then one should be taken to be denying *P*, rather than asserting not-*P*. We have seen, however, that Frege and Quine defend an equivalence that allows them to say that a denial is to be treated just as the assertion of a negation. That might make things easier from a logical point of view, but, as has been shown in preceding chapters, it has left us with metaphysical problems of finding truthmakers for negative truths, explaining negative properties and kinds and non-existents objects. If denial is primitive and has a distinctive function, perhaps there should never have been a paradox of negative judgement in the first place since we were never making negative assertions. We were making denials instead. It is clear, then, that we need to examine the equivalence thesis in detail and consider whether the alleged equivalences really hold. If they don't, then perhaps it was always wrong to treat denial as an assertion and, specifically, to take 'not' as something to be asserted rather than as a denial indicator.

10.3 The separate functions of assertion and denial

Our preference for affirmation leads us to accept assertion and attempt to dispense with denial. Admittedly, we don't need two primitive linguistic functions or forces for sentences if one will do. However, the equivalence thesis should be rejected. A denial is not the same as the assertion of a negative, nor vice versa. Assertion and denial play different functions, even with the addition of negation.

Table 10.1 Table of assertion and denial

Assertion	Denial
Aims at truth	Rejects a truth
Right when what is asserted is true	Right when what is denied is false, though not necessarily only then
Requires a truthmaker	Requires no truthmaker
Commits to a way the world is	No commitment to the way the world is
Determinate to a degree	Indeterminate
Error risking	Error avoiding
Expresses a judgement	Withholds a judgement
Observes the norm of epistemic responsibility	Does not observe the norm of responsibility
Reason for action	Reason for inaction
Bound by the law of non-contradiction	Exempt from the law of non-contradiction
Univocal	Equivocal
Conventionally opening	Conventionally responsive
Compatibility-silent	Indicative of incompatibility

It follows that there should not be a language consisting solely of assertion, nor one consisting solely of denial, and this defeats Dummett's scepticism over the possibility of two such languages.

The key question, as Smiley (1996: 4) says, is whether there are things that can be done with rejection or, as I am saying, denial that cannot be done with assertion alone. If there are, then Frege is not entitled to deploy Occam's razor here. It is of the utmost importance, then, that we note the following differences between the functions of assertion and denial, as indicated in Table 10.1.

Before explaining some of these contrasts between assertion and denial, it will first be in order to add a few general notes.

As has been mentioned already, affirmation and denial concern only declarative statements. Imperatives are neither true nor false and cannot be asserted or denied. They could be followed or not followed. Similarly, interrogatives are answered or not answered but are neither true nor false. Of course, the answers to interrogatives can be true or false, but the questions they answer cannot be. One might say additional things about leading or rhetorical questions that carry implications, but we will set those aside too. Assertion and denial concern only statements: the sorts of things that are true or false (I assume bivalence). The division concerns only a part of ordinary language, therefore. A language containing only declaratives might be possible, but it would almost certainly have inadequate expressive power for any reasonably intelligent form of life.

Second, it should be noted, as in chapter 9, that the concern here is only with contingent truths: the truths that depend in some substantial way on how the world is. I am not venturing any theory as to what are the truthmakers for

necessary truths, and nor can my account be assumed to extend to the denial of a necessary truth. This lacuna is to be excused on the grounds that the subject of the book overall concerns how we understand claims of non-existence, and it is contingent falsehoods where this issue is to the fore.

Another point to make clear is that assertion and denial need not exhaust all attitudes it is possible to take towards a statement, all the forces by which one might convey a proposition, or all the doxastic states by which one might regard a proposition. Some statements we will assert, others we will deny. But it seems quite possible to be in a not entirely decided epistemic position, where one feels that *P* is more likely than not, for instance, in which case one might assert *P* with reservation or reticence. On the other hand, one could believe that *Q* with absolute certainty and be willing to assert it with confidence. Given that there are degrees of belief, then we ought to allow some shades of assertion and similarly of denial.

There are still other clarifications to be made but they will be more appropriate later, in discussion of the individual features. Let us now move on to a description and justification of the differences listed in Table 10.1.

10.3.1 Truth and falsehood, correctness and incorrectness

Assertion aims at truth (Marsili 2018). There are various norms that enforce this. You should assert only what you believe to be true, for example. Aim is an intentional-dispositional word that does not entail success. One does not want to say, for instance, that assertion is correctly used only when what is asserted is true (contrary to what Marsili calls the *factive account*). One is still making an assertion when one asserts what is false, and it is just as much of an assertion as one that is true. We must distinguish, therefore, what I am calling *correct assertion* (assertion used correctly) and *true or right assertion* (where what is asserted is true). One is correctly using assertion when one aims at the truth, even where it is unknowingly false. What of people who tell lies then? Are they not making assertions? And, if not, how then do they lie? I agree with Dummett (1973: 298–302) that the way to understand liars is that they are misusing or abusing assertion, exploiting the fact that we all know that the aim of assertion is truth. A liar states something with the intention of conveying an impression of aiming at the truth, but when they are not doing so. That is how the deception occurs, and given that the aim of a lie is deception, then I think of lies as *fake assertions*. This contrasts with someone who asserts unknowingly what is false, who is still using assertion correctly. I find this account more precise than a view in which ‘The liar deceives by false assertion’ (Webber 2013: 651). Perhaps that sounds credible since it is right, as Webber says later (2013: 652), that ‘To lie is to express a proposition one believes to be false.’ The liar certainly expresses the statement in question, but

whether they thereby assert it is a different matter, I contend. If they are knowingly not aiming at truth, I think that is grounds to reject the view that a liar is asserting anything and that they deceive through a false and disingenuous expression that they want us to interpret as assertion.

What of denial? Clearly, one does not aim to deny what is true. The simplest thing to say would be that one aims to deny what is false. Certainly, someone who uses denial this way is using denial correctly in that instance. However, some think that this is not the only correct use of denial since another way in which denial is understood is as rejecting, cancelling, correcting, or contradicting a previous assertion (as with Givón 1979: 139 and Strawson 1952: 7). An aim of cancelling or contradicting a claim is not the same as aiming at a falsehood even if such a contradiction would be correct when what is contradicted is false. I say more about the equivocal use of denial below. It is of course entirely possible to use denial incorrectly; for instance, if one denies something that one knows to be true. One can lie by denial as well as by assertion. One perjures oneself by denying something under oath that one knows is true (Green 2001). Fake denial is possible too, then.

10.3.2 Truthmaking

According to truthmaker theory, truths are true in virtue of having truthmakers. Assertions need truthmakers in order to be true, and truths are thus dependent on *being* in some such way. We saw in chapter 9 how it has proved notoriously difficult to identify the truthmakers for ‘negative’ truths. However, a difference between assertion and denial might be useful here. In denying something, one is not claiming that it, or that anything else, is true. Hence, an expression of denial, since it does not aim at truth, needs no truthmaker and nor, one may say, false-maker. While truth is dependent upon what there is, falsehood is not. There is a clear asymmetry here, since if it is true that there is a hippopotamus in the room, something has to exist: the hippopotamus, as well as the room, and perhaps a relation of being inside. If one denies that there is a hippopotamus in the room, one is not saying that anything exists. What is said might entail the existence of things, like language and people and concepts, but one is not *saying* that they exist and nor are they the truthmakers of what is said.

10.3.3 Commitment

The above shows us that assertion and denial differ in respect to their commitment. The reason assertions need truthmakers is that they commit to a way that

the world is. In making an assertion, one is making such a commitment. Again, in the terms of truthmaker theory, in making an assertion you are committing to the existence of its truthmaker. This commitment could turn out to be mistaken, where the assertion is false, and thinking in terms of commitment shows why it is false: one committed to the existence of something that was not really there. In contrast, denial is relatively commitment free. If one denies that there is a hippo in the room, at most one might be committed to the existence of the room. As suggested above, however, on one use of denial, you might not even commit to this since one possible reason to deny the statement is if you doubt the existence of the room. The focus is usually on the hippo, but suppose someone looks at film of a hippo and suspects the room in which it appears is simply computer-generated imagery and the hippo is really in the wild, where it should be. Details aside, it seems that one reason to deny something is so as to make no commitments. This means, of course, that one need not commit to an alternative to *P*, if one denies *P*. One can be quietist about what there is. In denying that there is a hippo in the room, one says nothing at all about what is in the room, where actually is the hippo, whether there even is a hippo, or a room, and so on. There is a contrast with assertion, then, since one can assert something only by committing to a way the world is, by which I mean the existence of that in virtue of which such a truth would be true. Perhaps the case of negative existentials is even clearer (see Parsons 2006). To say that there are no unicorns commits to nothing at all (again, setting aside things like concepts). Even saying that there are no Arctic penguins makes no commitment. It would be a mistake to assume that it commits to the existence of the Arctic, for instance.

10.3.4 Determinacy

Connected to the above, an assertion is a determination of a way the world is, whereas denial is relatively indeterminate, since it concerns what is not the case rather than what is. When one asserts, for instance, that the pen is blue, one is identifying a particular state of affairs. When one denies that the pen is blue, one is not picking out a determinate state of affairs. Indeed, it is dubious that one is picking out any state of affairs at all. This difference was suggested by Ayer (1952) to distinguish negative statements: they would be less specific than affirmations.

Now it might be objected that determinacy is all a relative matter and that this is not, therefore, any sure way to distinguish assertion and denial. After all, the pen being blue is also indeterminate: indeterminate relative to it being a particular shade of blue (for determinables and determinates, again see Armstrong 1978b: ch. 21). But this sort of objection does no real harm. The point can be taken that there might always be more determinate states of affairs than the one asserted. There is nothing wrong *qua* assertion with asserting something relatively

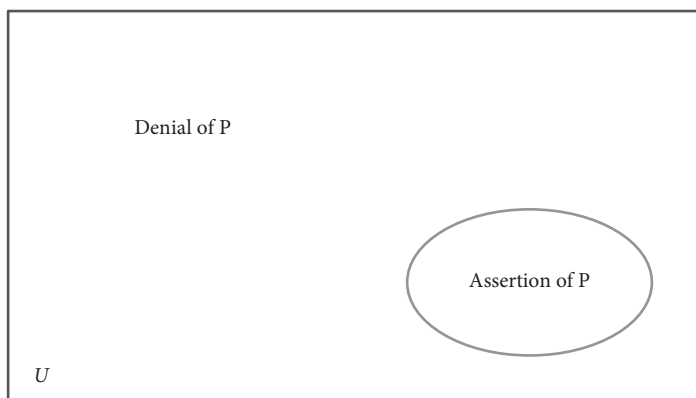


Figure 10.1 Relative determinacy of assertion and denial

indeterminate, for instance, that <there is *something* in the room>. But the indeterminacy that follows from denial is of a different kind. If one denies that the pen is blue, this means that it could be any number of different and incompatible colours at the same level of determinacy as that which is denied of it. The pen could be black or red or green, but it cannot be more than any one of these. In this sense, there is something distinctive about the indeterminacy of denial, since in the case of assertion, the indeterminacy, if there is any such, must occur at a relatively greater level of determinacy. Perhaps this already concedes too much, however. Given what was said under 10.3.3, in denying that the pen is blue, one need not even commit to the existence of the pen, hence one is not committed to it having any colour. One could think of the point this way. Assertion commits to their being some subset of possibilities: the more determinate the assertion, then the smaller the subset of possibilities. No matter what is asserted, however, its denial is consistent with all other possibilities; that is, everything outside of that subset, as shown in Figure 10.1.

There is another objection to the claim that assertion and denial differ in respect of determinacy. This is that it is possible to assert a negation, which is thus an assertion that is indeterminate in the sense outlined immediately above. It is clear that I will discard this objection straight away, however, which is justified given the project of this chapter. The point of considering the differences between assertion and denial is to show that treating a denial of *P* as an assertion of not-*P* is to change what is said in quite a deep way. It is not just a matter of surface grammar. Taken as a whole, the differences between assertion and denial outlined here show why one wishes to assert something that is relatively determinate, consistent with Grice's (1975) maxim of quantity. Our aim is for assertions that are informative, hence determinate. Thus, while technically one could assert a statement as indeterminate as <something exists>, doing so breaks conversational maxims, which is why nobody outside of philosophy would ever assert it.

10.3.5 Error

That assertion is relatively determinate and denial is indeterminate explains why a connection has been drawn with the subject of error. In making an assertion, I am putting myself at the risk of error. The above shows how. An assertion includes a relatively determinate commitment to a way the world is: specifically, to the existence of that in virtue of which the assertion would be true. One risks error in that one can be wrong that the world is that way. But we have seen that, in denial, one makes no such commitment, and it is thus a less risky act. Sparby attributes such a risk-avoidance view to Kant:

Negation in the form of denial for Kant can also have the role of ‘avoiding error’ [Kant: 1781: B97]—when I say that a circle is square-shaped I say something untrue; I avoid this error by a denial, stating that it is not the case that the circle is square-shaped. I do not say, however, “the circle is non-square-shaped,” because square-shaped has no determinate opposite—when I say that something is non-square-shaped, I would not know whether it is triangle-shaped, circle-shaped, or a pentagram, etc. I would know, however, that anything non-square shaped has a shape that is not a square-shape. (Sparby 2012: 28)

10.3.6 Judgement and belief

There is a further reason why assertion risks error while denial can avoid it. Assertion has an accepted connection with the expression of a judgement. If one accepts the above equivalence thesis, one will say that denial is just an assertion of a negation. But it is possible to think instead of denial as the withholding of a judgement or declining to make a judgement and this would be a reason to reject the equivalence thesis. More powerfully, one could think of denial as a way to bar a sentence from your belief register (Price 1990: 225). In contrast, if you have such a thing as a belief register, then you should only assert something if it is in there.

This shows that the equivalence thesis is epistemically misleading. In declining a judgement on *P*, one need consider no other proposition or hold any other epistemic attitude towards anything else. For this to be treated as equivalent to asserting that not-*P* is then entirely wrong because to so treat it does say that, *qua* assertion, one is forming a judgement and adopting a belief (that not-*P*). So what we are initially treating as a withholding of judgement gets incorrectly converted into a judgement. It is obvious that there is a difference between (i) not judging that-*P* and (ii) judging that not-*P*. It is plausible that at least some denials should be considered as cases of (i), though I accept that some cases will be like (ii). But where we have a denial of the first type, it is clearly an error to immediately treat it as an assertion of not-*P*. (i) need not express a judgement, which (ii) does.

10.3.7 Epistemic norms

There are norms that attach to assertion and govern its proper use. These follow from there being epistemic virtues (Zagzebski 1996, Baehr 2011). Denial is not subject to the same norms, though there might be others that are relevant for it. The norm that seems most relevant, in showing the distinctness of assertion and denial, is that you have an epistemic responsibility to exercise due diligence in your assertions. You cannot assert whatever you like, even if you are aiming at truth. After all, if you assert *everything* then you will succeed in asserting all the truths, but you will probably have just as much, or more, falsehood besides. This norm shows us why it is possible to incorrectly assert even when one has aimed at the truth and hit the target.

Here is an example. Suppose you hold your hands behind your back and ask me to say in which hand you hold a reward, which I can have if I pick correctly. We both know that I cannot see in which hand the reward is and nor do I have any other information that will allow me to locate it. I say that it's in your left hand. Lo and behold, it is. But are either of us right to think that my statement was an assertion rather than, for instance, a mere guess? I suggest that the answer is no: I just got lucky. I had no particular reason to say the reward was in the left rather than the right hand. I was aiming to get it right and I got it right, but I was unable to observe a norm that governs assertion. I might articulate the norm something like this: you need some sort of reason for believing P before you go around asserting it. It cannot be a lucky guess. Nor can your assertions be indiscriminate. You need grounds to make an assertion: groundless assertion is irresponsible.

This sort of consideration shows why assertion carries another acknowledged conversational implicature. When I assert something, the default understanding is that I do so with a reason such that you should believe it too if you knew what I knew. Indeed, this shows that assertion has a point. We don't assert for its own sake. Usually we are doing so to convey our view of things, often so that others can update their own beliefs or act according to the new information. Hence, one shouldn't assert something that you know your listener already knows and remembers. For assertion to have this function, it must be well grounded, by which I mean not just that it is true but that it is based on some good reasons. I assume that what makes a reason good is that it has been shown to reliably deliver the truth. I am not simply expressing an allegiance to a reliabilist theory of knowledge, even though that has its attractions. There is a simpler point concerning language and truth: if someone's assertions have proven unreliable then you can have reasons to doubt the utterances of that person even though by chance some of them might be right. Speaking the truth isn't enough, then.

Now consider denial. Compared to assertion, you can be more irresponsible in your denials. You can deny something with relatively less due diligence, simply

where you see no grounds for asserting it. In a case where there is no knowledge either way, as to the truth or falsity of *P*, for instance, the thing to do is to deny *P*. (Could one also deny not-*P* at the same time? I will address this question below, under 10.3.10.) So one does not need any particularly good reason for denying *P*. Not believing *P* is enough, which again seems to show why a denial of *P* is not the same as an assertion of not-*P*. Presumably, following the above epistemic norm, one would need a good reason to assert not-*P*. But if you think of denial as having a function of judgement-withholding, then it is not that one needs a good reason for the opposite of that judgement in order to withhold the judgement. As stated in 10.3.5, denial is the least risky option in cases of uncertainty so should be considered the default.

Perhaps there is an exception to this. The stubborn denier is someone, let us suppose, who denies everything, including cases where there is incontrovertible evidence presented to them for the truth of a claim, and yet still they deny it. Might we say that the stubborn denier is breaking a norm for denial that is exactly parallel to the norm for assertion? You need good reasons to assert what you assert and maybe you need good reasons to deny what you deny. Someone who affirms everything will sometimes have no good reason to do so and, similarly, someone who denies everything will sometimes have no good reason to do so.

It seems, however, that the stubborn denier is someone who fails to observe a different norm. The first epistemic norm, above, was that you should have good evidence in order to assert something. The second norm, the one that the stubborn denier ignores, is that you should accept something as true when there is good evidence for it. Needless to say, one can make many denials without having to pay heed to either the first or second of these norms. A denier might, however, be put in some circumstances in which she should pay heed to the second norm and no longer deny *P* if there is conclusive evidence for its truth. From this, we can see that the norms governing assertion and denial are different.

There are other epistemic norms, of course. If you assert *P*, then you ought to accept what is entailed by *P*, at least if you also accept that it is entailed by *P*. Denial does not have quite this norm. One does not want to say that if you deny *P*, then you should either accept or deny whatever is entailed by the denial of *P*, since it is not clear to what, if anything, a denier of *P* is committed. On the other hand, there might be other norms that attach to denial. For instance, if you deny that *P* then you ought also to deny anything that would entail *P*, at least if you also accept that it entails *P*. This is a significant logical asymmetry between assertion and denial. For example, if I assert that there is a hippo in the room and also accept that all hippos are four-legged, then I ought to accept that there is something four-legged in the room. If I deny that there is a hippo in the room, and accept that all hippos are four-legged, there is no norm that I must deny that there is something four-legged in the room: rightly so, since a cat might be there.

10.3.8 Reasons for action

Assertion can provide a reason for action, if coupled with desire. As suggested above, assertion plays a role in our linguistic exchanges. For instance, it can be used for persuasion. Now suppose one is persuaded through an assertion that there is bread in the cupboard, and one is also hungry and has permission to eat. Acquiring the belief that there is bread in the cupboard can then guide your action. That beliefs are action-guiding is well understood, so it is not surprising that assertions can also guide actions if they lead to acquisitions of beliefs. Assertion can play this role because, as we have seen, it commits to a determinate way that the world is. It tells you what there is, and it is reasonable that our actions take account of what there is. The point need not be laboured.

Denial has no such role, since it is relatively uncommitted and indeterminate. Suppose I am an animal welfare officer, chasing an escaped hippo. If I hear that the hippo is in the room, I have a reason to go in and get it. If, on the contrary, I hear a denial that the hippo is in the room, then I have no reason to go in and get it. I perhaps have to think of something else to do that would be based on assertion or belief. Maybe the hippo is in the garden: look there. One is tempted to say that denial gives a reason for inaction or omission—I can omit to go into the room and look for the hippo if you or I deny that it is there. But omissions are not actions, of course. And even here, the denials might not guide me to make an omission in every case since, being non-committal, the denial might just be an expression of a lack of evidence: the epistemic use of denial. I could still go and look for the hippo even if I hear a denial that it is there, then.

This is one place where it really matters whether someone asserts a negation or not. There is a difference between someone thinking there is inadequate evidence to believe there is a hippo in the room and someone who is convinced that there isn't. If so, can both be treated as cases of denial? How then can we distinguish not believing *P* from believing not-*P*? To consider this further, we need a look at the next dividing feature, which is univocity and its opposite.

10.3.9 Univocity

Table 10.1 indicates that assertion is univocal while denial is complex and equivocal. Suppose there is a spectrum of degrees of belief. Assertion would be appropriate only at one far end of this spectrum, as per the norm for assertion. Assertion should come with belief; that is, where at least some minimal threshold for belief has been passed, even if that is short of absolute certainty. It violates epistemic norms to say that there is a doubt over *P*, but it will nevertheless be asserted.




Figure 10.2 Relative scope of assertion and denial

Denial allows a greater variety of doxastic states, most obviously where one denies *P* because one thinks it false or one is denying *P* because one thinks there is no good reason to assert *P*. Denial is thus equivocal. The difference relates to the Kantian idea that denial can be used as a way of withholding judgement, which can fairly be described as not asserting *P* but without a commitment to not-*P*.

The situation is something along the lines indicated in Figure 10.2, which shows that denial covers a broader range of possible states than does assertion. Denial can be used both to withhold judgement and to indicate incompatibility: the denial of *P* is incompatible with an assertion of *P*. Further, one wants to be able to deny nonsense or category mistakes without saying that they are false. One should deny that a thought is five metres long, for example. To assert that thoughts are not five metres long raises the question of just how long they are. A pay-off of such an account of denial, if perhaps a trivial one, is suggested by Smiley (1996). Suppose we want to reject (1):

- (1) Smith, who committed the murder, was in Berlin.

Should we do so simply through an external negation, as with (2)?

- (2) It is not the case that Smith, who committed the murder, was in Berlin.

Smiley says not, since both (1) and (2) will be false if Smith did not commit the murder. Instead, then, we can reject (1) by denying it. Given what we have said above, it would be clear that the denial of (1) does not commit us to the truth of (2). As that seems the right result, this is a further reason for rejecting the equivalence thesis. There are thus very good independent reasons for saying that a denial of *P* should not commit you to the truth of not-*P*. This is how to answer the ‘have you stopped beating your dog?’-type question. Because the example is distasteful, as it is supposed to be, I feel forced to substitute it with one slightly less so. Assuming there is no truth to the allegation, you still do not want to say:

- (3) I have stopped stealing lunch,

since a statement that something has stopped implies that it was happening. But nor do you want to say the following is true:

(4) It is not the case that I have stopped stealing lunch.

(4) puts you in an even worse position than (3). The healthy attitude, then, is to deny (3), which, for the reasons given, does not commit you to its negation. This shows another relevant asymmetry, which has been so far only implicit: we want to explain negation in terms of denial rather than explain denial in terms of negation. The argument that denial is primary is helped by consideration of cases such as these since they show where denial should absolutely not be treated as negation, nor as assertion of a negation. Negation is a tool we can sometimes use to indicate denial but not always.

10.3.10 The law of non-contradiction

When I assert P , I cannot at the same time assert $\neg P$, since this would violate the law of non-contradiction. This is, of course, another normative matter since there is little doubt that people hold contrary beliefs and from time to time make statements that contradict their other statements, either because they have changed their minds, they are deceitful, or they have simply forgotten that they hold a contrary view. If there is a norm that you should not assert inconsistent things, it might well originate in the more fundamental norm that in assertion you should aim at the truth. You cannot be aiming exclusively at the truth if you deliberately make inconsistent statements since they cannot all be true. There are no contradictory states of affairs in reality, we suppose, that could make contradictory statements true. Since denial has no commitment, however, and can be used to withhold a judgement, there is no such constraint. Hence, one can in principle deny both P and $\neg P$.

This might sound wrong, but only if one still thinks of a denial of P as also an assertion of $\neg P$. If one legitimate understanding of denial is that it indicates not asserting P , then it is clear that one might want to assert neither P nor $\neg P$, and thus that one might deny both.

The following case is then possible. Hercule Poirot is the world's greatest detective because he only allows his beliefs to fit the evidence and he only draws valid inferences from what he knows. Now suppose he hears an assertion of:

(5) The suspect was in the room at the time of the murder.

Evaluating all the evidence, he might decide that there is insufficient to assert (5). There is nothing that conclusively places the suspect there at that time. Now what if someone hears Poirot deny (5) and then concludes:

- (6) The suspect was not in the room at the time of the murder.

Poirot realizes that we should not say this either. There is no evidence to place the suspect in the room at the time but no evidence to place him elsewhere either. Poirot realizes, then, that he should deny both (5) and (6). The law of non-contradiction does not apply to our denials.

The reason why *non-contradiction* constrains our assertions but not our denials follows from what has been said above. Assertion commits to a way the world is, and we cannot commit to contradictory ways that the world is. Denial makes no such commitment and can be used to withhold judgement. There is nothing to stop us withholding judgement on P and withholding judgement on not-P. We have found another way in which assertion and denial differ in their functions.

10.3.11 The responsive convention

Denial seems to differ from assertion in another way. Denial is conventionally responsive while assertion is conventionally opening. Assertion tends to come before denial in a discussion since denial can be used to express disagreement with the assertion. Such a view is supported by the following:

Negative declarative sentences constitute a different speech act than the corresponding affirmatives. Affirmatives are used to convey new information on the background of assuming the hearer's ignorance. Negatives are used to correct misguided belief on the background of assuming the hearer's error. (Givón 1979: 139)

This is, however, a conversational convention only. No one would claim that assertion must always occur before denial. It is more an expression of what tends to be the case and, like any convention, it is up to the participants whether they follow the convention or not. As a good philosophical illustration, however, we can take Plato's work, since his early Socrates seems to specialize in denying philosophical theories. Note how much work Socrates must go through in order to elicit an assertion from his audience first. This provides him with a thesis to be denied in the next part of the work. In *The Sophist*, Socrates repeatedly assures his companions that they should have no fear in offering a definition even though, when one is eventually forthcoming, it then gets systematically demolished.

Perhaps one can make too much of this. Clearly it is possible to start a discussion with a denial, especially one that uses the word 'not'. Two football fans meet, for instance, and one says to the other 'United are not playing very well these days.' Sometimes, it could really matter that something is not the case, and that might be the first thing one wants to say. Perhaps there is not even a convention, then, that denial is responsive and assertion is for opening. One might reply,

however, that in this example there is a perfectly good contextual explanation of why the convention should be flouted and the discussion opens with a denial. It might occur in a context where United have had a long spell of excellent form and this fact constitutes the assumed background knowledge that the participants can take as read: effectively helping themselves to a tacit assertion. With United's loss of form, now is the right time to deny that they play well.

It might look as if there is no conclusive way of deciding whether or not there is a convention along the lines I have suggested, and one can choose as one wishes to open with an assertion or a denial. However, there is some support for the view that denial is conventionally responsive when one considers the final feature of denial because this shows that denial has a function that explains why it is apt to be used responsively. In other words, denial has a function that assertion does not have that gives it a particular responsive capability.

10.3.12 Incompatibility

The responsive function of denial stems, I suggest, from its ability to indicate incompatibility with an assertion. Assertion alone contains no such indication, as Price (1990: 224) says. For example, suppose that John goes to visit Fred and, walking up to Fred's house, sees him in the back garden. When John knocks on the door, Freda answers. John says 'I've come to see Fred' and Freda replies 'He's in the kitchen.' Here we have an assertion. Now John knows that Freda is mistaken, since he saw that Fred was actually in the back garden. John has a choice of how to respond to Freda. He could say:

(7) Fred is in the back garden.

Or he could say:

(8) Fred is not in the kitchen.

John might even say one of these followed by the other, but which would be most important, in this context? As Price suggests, John should respond with (8) since it is clear that (8) is a denial of Freda's statement that John is in the kitchen. The denial is clearly indicated with the word 'not', which contradicts what she says. Were John to, instead, assert (7), then there would have been one assertion followed by another assertion. In asserting something new, it would have to be inferred that John is contradicting Freda. Many assertions can be true together, such as that China has the largest population of any country and that bananas are bent, which are in fact both true. It cannot be simply assumed, then, that a second assertion contradicts a first. And in the example given above, any such inference

depends on an extra premise being added that it is not possible for someone to be both in the kitchen and in the back garden at the same time. Let us put aside any question of whether this additional premise is itself a negative one. The point is that denial registers incompatibility and John can use it to contradict Freda's assertion. Were he to use only an assertion in response to hers, it is not clear cut that he contradicts her, certainly not without an additional step of reasoning.

Lewis Carroll (1895) explained another aspect of this issue. One assertion B is in conflict with another, A, only if one adds an additional claim, C, that A and B cannot be both true together. But then the same difficulty reappears since it might be wondered why all of A and B and C cannot be true together. One might add a further premise, D, then, that tells us that A and B and C cannot all be true together. But the same question can be asked of A and B and C and D. Can't they be all true together? A regress is underway. The way to prevent the danger of a regress is to allow ourselves a direct, non-inferential marker of incompatibility, and this is what denial does to assertion. The denial of P is straightforwardly incompatible with the assertion of P and does not have to go via some, possibly incomplete, inference.

This, then, gives some reason why denial could be conventionally responsive. It is able to do something in response that assertion, unaided, cannot. Denial expresses disagreement directly, which assertion does not. To do this, of course, a denial needs something to disagree with, and that is why it follows assertion. As Bosanquet (1888/1911: 280) puts it, 'negation does presuppose some affirmation.' Starting with a denial is still possible, but, in those cases, typically there is an implicit background assumption that the denial contradicts.

This brings us to the end of our examination of the differences between assertion and denial. The aim was to show that assertion and denial have different functions such that one cannot take a denial of P to be the equivalent of asserting not-P. In denying P, one is doing different things from what one does in asserting anything. If the above account is successful to some degree, then a denial is not to be treated as a kind of assertion.

We might not, however, be home and dry. Someone could acknowledge the above differences between assertion and denial and still think there is no reason to reject the equivalence thesis if, for instance, there would be some major problem in doing so. In the remainder of this chapter, therefore, I will examine possible objections to understanding negation in terms of denial.

10.4 What are you denying?

What is offered here is not an analysis of the meaning of 'not', nor an account of how the word 'not' is used in ordinary language. It is an attempt to show how we

can interpret the word in a metaphysically harmless way that avoids commitment to negative existents.

However, there is a major objection that will come right away to this kind of account, as Horn (2001: 59ff) details. I already touched on this objection in §9.4 but I left the matter partially unresolved there and return to it now. There has been a history of attempts to reduce negation to falsity or denial, as with Leibniz (1686: 58), Bradley (1883: 118), Sigwart (1895: 122), Baldwin (1928: 147), Wood (1933: 421), and Russell (1940: 81), while also Ayer (1936: 17) reduced falsity to denial. But Frege (1919), Austin (1950), Quine (1951), and Geach (1980) think this is essentially confused.

If you say:

(9) The door is red,

It seems clear that you are talking about something in the world: the door. To do this, you are using, following Tarski (1933), an object language, which is a language where the referents are worldly things.

Now suppose one suggests, as in Mumford (2007a), that when one makes a statement of the form not-*P*, one should understand this to mean that:

(10) <*P*> is false.

So let us now consider:

(11) The door is not red.

According to the account under consideration, this is to be understood as <the door is red> is false. But here we can see an unreasonable asymmetry between the two cases of (9) and (10). We said that (9) was expressed in an object language. But if we move to (10), we see that what this is about—what it declares to be false—is not a worldly object but a proposition or truth value-bearer. In other words, attributions of falsehood (and truth) are metalinguistic statements about propositions. But this is wrong, according to the objection. ‘The door is not red’, (11), is just as much about the door as is (9), ‘the door is red.’ Why should we accept a theory that says that just because we are using the word ‘not’ we are switching from an object language to a metalanguage? As Smiley (1996: 2) puts it, “‘*P*’ is false” will not do, because...It is about *P*, not about what *P* is about.’

Some have been prepared to accept this asymmetry between affirmation and negation. Bergson (1911: 287–8) said that negation is only an attitude taken by the mind. Affirmation is a judgement about an object while negation is a judgement on a judgement. And Joseph (1916: 172) said that ‘positive’ statements could still be true in a world without minds while negative statements could not be.

In answering this sort of objection, we need to keep two matters separate. This first is the idea that positive and negative statements are about very different things: one about the world and the other about language. The second matter is whether in affirmation and negation we perform different kinds of act. These must be separated because, in the account of negation as denial, the second point should be accepted but the first should be rejected. We refer to the idea that affirmation and negation are different as asymmetricalism. The asymmetricalist would reject the Frege equivalences, for example. But even if one allows that there is a difference between affirmation and negation, it still matters which difference. We have to get that right.

Let us consider the two issues, then, since I suggest that they are run together in the criticism we are facing. The first issue suggests that we must be referring to different things in assertion and negation, especially where negation is taken to be a statement that a proposition is false. But we are not in the business of trying to reduce negation to falsehood, here. We are not even in the business of reducing negation to denial, as will be explained later. But it is an account of negation in terms of denial, admittedly. This is a significant move, however, since it is reasonable to suggest that denial is just as much about the world as is assertion. One can both assert and deny in an object language. (9) asserts something of the door while (11) denies something of the door. They are both about the door. So while there are many asymmetries between assertion and denial, one of them is not that denial is metalinguistic, in contrast to assertion. Assertion and denial can have something in common, then, and we do not have to defend a view that merely by introducing the word 'not' into a statement one is thereby switching to a whole different way of speaking, in comparison with the corresponding assertion.

Nevertheless, as we have seen, assertion and denial differ in many ways, and it seems perfectly reasonable to say so. What has been detailed above explains why. What it leaves us with is an asymmetry in which assertion concerns what is, and is about that in the world that would make it true if it exists. An assertion that the door is red is true when that which it is about—the door being red—really is a state of affairs. Denial has to be different from this if our account is going to succeed in dissolving the problem of negative truth. Denial is about the world—that it contains a red door or unicorns—but it is a way of saying that there are no such things in the world. If there is any truth in Bergson's asymmetricalism, in which negation is only an attitude of mind, then this is it. Assertion is aiming to say something about being and is true when things are as they are said. But denial can be about what there is not: non-being, where there is nothing corresponding to the denial in reality. So there is no worldly correspondent to a denial, when it is correct, while there will be a worldly correspondent to an assertion that is correct. Asymmetricalism in this regard seems absolutely right.

10.5 Can you deny a denial?

Gale (1976: 21) is sceptical about treating negation as denial. Just as you can affirm a negation, he thinks, so can you deny one. A negative cannot be the same as denial, then, since it can itself be denied. Whether statements are positive or negative and affirmed or denied, he says, are independent matters, not to be conflated.

So what happens when you deny not-P, which it seems that you can, perfectly sensibly? You might deny that there's not a hippopotamus in the room, for instance, just as much as someone could deny that there is one. Is a denial of a negative, on the account offered here, a denial of a denial? What would that be? Can denials be meaningfully iterated (I deny that I deny that I deny that I deny that....)? For that matter, can assertions be meaningfully iterated either? Is an assertion of an assertion any more than the original assertion? Suppose we prohibited iteration of both assertion and denial. Then matters would be simpler. When one wants to affirm P, one simply asserts P by saying 'P'. And when one wants to deny P, one does so by saying 'not-P'. Why, for instance, deny that not-P since we all know, by law of double negation elimination, that not-not-P is the equivalent of P? Instead of denying not-P, then, one should simply assert that P and eliminate any confusion.

The above will not quite do, however. We must acknowledge that while there are improper uses of both iterated denial and iterated assertion, there are proper uses of each as well. To put the difference in the most abstract and brief terms, it is improper to double negate or double affirm when one intends a statement to be in an object language. Iteration has a place, however, where one intends a metalinguistic statement. Now let's consider this difference with some concrete examples.

In some cases, one might double negate P, with not-not-P, when what one really wants to say is P. This is an improper use of double negation. A good copy editor, for example, will spot such double negations and suggest that they should be changed to assertions. But not all double negations should be changed since there are some perfectly valid contexts in which a double negation makes sense, especially once the idea that 'not' is an expression of denial is taken seriously. For one thing, we noted that denial was equivocal. One might deny P because one thinks P is false, but one can also deny it just because there are insufficient grounds to assert P. And in a double negation, it is possible that both of these types of denial are in play. If the first negation (not-P) is made in the belief that P is false, then a denial of that negation (not-not-P) need not entail that P is true if it is a denial that there are adequate grounds to say that P is false. For similar reasons, Price (1990) argues that a law comparable to the law of double negation elimination does not apply for the case of denial. A proper denial of a denial is not an assertion. And that is a result that we will welcome if the differences between assertion and denial in Table 10.1 are taken seriously, since it gives us

various reasons why we should not be able to convert automatically a denial into an assertion. Converting a denial of a denial into an assertion would be a special case of what we are prohibiting. Instead, an assertion unambiguously corrects a denial. Therefore, if this is what one intends when one denies a denial of *P*, then one should indeed assert *P*. But, as I suggested, there are at least some contexts in which one can legitimately double deny, and other contexts in which one can legitimately double assert.

Smiley (1996) notes correctly that while $\neg\neg P$ might be logically equivalent to *P*, just as $\neg(\neg P \vee \neg Q)$ is logically equivalent to $(P \& Q)$, it does not mean that they are exactly the same in all non-logical ways, since there is some difference that is not a logical difference. We can envisage a context in which *A* says:

(12) I did not say that *B* was not in the room.

A is using denial appropriately here, and it should not be treated as an assertion that *B* was in the room, since what *A* says is consistent either with *B* being in the room or *B* not being in the room. Suppose *B* comes looking to find out who reported his alleged absence and asks *A*, 'Did you tell them I was not in the room?' Giving the answer (12) clearly denies having reported *B*'s absence while remaining neutral on whether *B* was actually absent. *A* is simply denying that she performed a prior act of denial. It is not essential to the example that it is made in the past tense or that the second denial always must refer to a temporally prior denial. Someone might shake their head while asserting *P* and then immediately go on to say, still shaking, that they are not denying *P* by shaking their head; it's just that their head shakes.

Why iterate a force at all, though? In one of the most famous philosophical jokes of recent times, J. L. Austin was lecturing in the United States and uttered words to the effect that 'it is a peculiarity of the English language that while there is a law of double negation, which says that two negatives make a positive, there is not a corresponding law of double positive that makes them into a negative,' and at this point Sydney Morgenbesser uttered, from the audience, 'Yeah, yeah.' Notwithstanding this lightning-fast quip, Austin seemed to be on to something. When we iterate negation, it switches the truth value of the proposition in question, whereas when we iterate affirmation, it seems to leave the truth value of the proposition unchanged (assuming Morgenbesser's rejoinder can be explained away pragmatically). There seems at least some possible point to iterating negations, then, but with no corresponding purpose to iterating assertions. Why bother saying, 'I assert that I assert that I assert that...'? Isn't that just the same as the original assertion?

However, given that assertion is an act, then there is of course a point in metalinguistically asserting that one has asserted something, comparable to the above example in the case of denial. For example, where a reward is at stake, there might be

some dispute over who (first) asserted that the suspect was hiding in the cellar. With others seeking to take the credit, C might assert that it was she who asserted that the suspect was hiding in the cellar. Just as double denial was a denial of an act of denial, so too, then, can double assertion be used as an assertion of an act of assertion. There was an improper use of double negation, though, where what one should really have done is assert; and, similarly, there is an improper use of double assertion. If one simply is asserting P in the object language, one should just assert P and no more. Likewise, saying that it's true that it's true that P is just for it to be true that P.

It is worthy of note that the account of double negation will be relative to the particular language in question. It is more of a pressing problem in English than in Latvian, for instance. In English, some people say things such as 'I haven't done nothing bad,' to the amusement of those who know about double negations, since it might be taken to imply literally that the speaker has done something bad. The equivalent in Latvian would be:

(13) Es neko sliktu neesmu izdarījis.

Transliterated, (13) would be 'I nothing bad haven't done,' but, unlike English, the two negations do not cancel each other out. Additional negations can be added fairly liberally in Latvian, with the effect of strengthening the negative claim. Hence, in:

(14) Viņš nemaz nav stulbs,

('He not at all is not stupid') there is a clear and strengthened denial of stupidity.

(15) Es nekad neko nedaru nepareizi,

contains four negations but remains a denial ('I never nothing don't do incorrectly.') This simply demonstrates that different natural languages can indicate their denials in different ways. I do not rule out similar emphatic uses of double denials even in English, where the law of double negation elimination is suspended to make a point. I've yet to hear of a natural language without a means of denial (see Jespersen 1917).

10.6 Logic and complex denials

Returning now to use of the English language, once we reach triple denials and triple negations, then, unlike Latvian, we really seem to be losing a grip on what is meant and intended. One might just be able to make sense of a triple denial but with a quadruple one we start losing the meaning altogether.

Does this show a flaw in the proposal? Is it further evidence that negation cannot be treated as denial since a logician can perfectly well understand $\neg\neg\neg P$,

while the meaning of quadruple denial seems less clear? All is not as it seems, however. I do not claim, in particular, that the account of denial that I have offered can explain away every use of negation, especially in the contexts of logical systems. This is not to disparage interest in rejective interpretations of negation in formal logic (e.g. Humberstone 2000), about which I express no judgement here. There is a view about the priority of denial and negation that is relevant, however. Following Price (1990), I take it that denial is prior to negation in the sense that it explains what we are doing when we negate, and this seems satisfactory for simple or basic negations. When we say not-*P*, we are denying *P*. As I argued above, we can also make sense of some slightly more complex uses of negation in terms of denial, again where it concerns its use in natural English. But once we had a word 'not' or a symbol, \neg , to express negation, then it was able to take on a life of its own especially in logical systems. $(\neg(P \vee \neg Q) \ \& \ (Q \rightarrow (R \vee \neg S))) \rightarrow (\neg Q \ \& \ \neg S)$ is a well-formed formula of classical propositional logic. It includes five negations and the prospects of cashing them all out as denials look grim. But quintuply negated formulae are an artefact of such a logic. It would have to be truly exceptional circumstances for five negations to be required in a natural English sentence, other than for pure emphasis, and, were that to occur—perhaps in a legal document—we would have to ponder hard how we should interpret it. Until then, I am reasonably satisfied that accounting for negations as denials will suffice for all practical purposes. And, besides, it is quite likely that even though 'not' is introduced for the purposes of expressing denial, then it can be used in all sorts of ways that go beyond that original purpose. We will see this in the case of distinguishing external from internal negations shortly, where it is possible to make this distinction grammatically by changing the place of the 'not'. There are different things that we can do in language and logic with not and \neg . But we do not have to read the structure of the world from the structure of either logic or language. Yet those who think there must be positive truthmakers for negative truths seem to be drawing this inference: from an expression in language to what there must be in the world.

One can choose to ignore logic then, as a construct designed for particular purposes. But, even so, there are still things that one wants to say naturally, and inferences one wants to draw, that seem to have a logical structure. Some of those will involve negation and be complex. What do we say of those? Is there any response to the so-called Frege-Geach problem, for instance, which seems to count against treating negation as a falsehood or a rejection?

Let us consider a simple example, which involves a negation within a *modus ponens*:

1. $\neg P \rightarrow Q$
2. $\neg P$

Therefore, *Q*

With premise 1, one is not necessarily denying P when one says $\neg P$ since this premise is in material conditional form. It is only that *if* $\neg P$, *then* Q , so the negation here most certainly cannot be taken as a denial. But premise 2 does say that $\neg P$, which, if we apply our account, will be taken as a denial that P . But now we have the problem. An inference of *modus ponens* can be drawn only if the $\neg P$ in premise 2 is the same thing as the $\neg P$ in the antecedent of premise 1. However, we have already said that while $\neg P$ in premise 2 is a denial of P , in premise 1 it cannot be. *If* $\neg P$, is not the same as $\neg P$, since *if* $\neg P$ says nothing about whether P actually is negated or not and, thus, nothing about whether it should be asserted or denied, under my account.

Nevertheless, sense can be made of an inference such as *modus ponens* in terms of denial. As Smiley (1996: 4–6) shows, this depends upon acceptance of certain rules of rejection: inference rules for handling arguments that contain a rejective negation. Treating negation as falsehood, which is what Frege (1919) discussed in raising the original problem, might be more difficult. But now we can interpret premise 1 as saying that if P is denied, then Q is affirmed. Again, *if* is used here, hence premise 1 is neutral on whether P actually is denied. But premise 2 tells us that it is denied, and thus the inference to Q is warranted. Q is conditional upon the denial of P , in premise 1, and that is what justifies its inference once we know that P actually is denied (this is similar to Smiley's (1996: 3) response). I would seek to offer similar treatments for other basic arguments involving negation, such as $P \vee \neg Q$, $\neg P$, therefore $\neg Q$. This could be interpreted naturally to mean that either P is affirmed or Q is denied, but P is denied, therefore Q is denied. The opposition we have outlined between affirmation or assertion and denial allows this.

Understanding negation as denial for metaphysical purposes need not cause havoc with our reasoning. Further, I am happy to leave formal systems untouched since they have a different purpose from arguments in natural languages. I do not accept, therefore, that anything in logic rules out this account of denial.

10.7 Internal and external negations

A distinction is often drawn between internal and external negations. We can understand internal negation as a predicate denial, as when we might say that a is unclothed. But later came external negation, which was not a part of Aristotle's original account and was introduced in subsequent theories of negation (see Horn 2001: 21f.). Here, we would say that it is not the case that a is clothed. The topic of privations is also relevant to this distinction, as covered in §3.4, above. A privation is understood as a deficiency, where something could, should, or ought to have a certain power or capacity which it does not.

Now what of denial in relation to internal and external negations and privations? I am content to treat all these cases as denials. For instance, both (16) and (17) deny that something can see, even if they do so in different ways:

(16) It is not the case that a rock can see

(17) The man is unseeing/blind.

Both uses of negation are appropriate, given that the man has a privation but the rock doesn't. They use their internal and external negations correctly. But, the objection might come, in treating both these as instances of denial, isn't there some loss of meaning? The seemingly useful and significant distinction between internal and external negation is lost when we move to talk of denial, instead, so does this show that denial is an inadequate tool for treating the subtle differences in types of negation? There are answers, however. First, we must remind ourselves that the account of denial offered here is not and never was intended as a translation of the ordinary usage of 'not' and its cognates in all their different colours. Rather, it is used to show how we can use negation without ontological commitment to any truthmakers. The fact that in making a denial, we could do so on the grounds of an internal or external negation does not impinge on this success. Second, however, there are, of course, ways to indicate differences in the scope of denial that correspond to differences of scope of negation. In (17) one denies seeing of a man, where the narrow scope consequently admits an acceptance that the man exists. (17) does, therefore, have some ontological commitment: to that man of whom seeing is denied. The denial in (16) is wide scope and would thus still be correct even if there were no rocks.

10.8 Conclusion

We should reject the equivalence thesis. According to that thesis, a denial of P is equivalent to an assertion of not-P. We have seen why this is not the case. Assertion and denial perform different functions.

A less obvious conclusion can be drawn from the above analysis. Suppose that someone makes a statement involving a negative, such as that there is not a hippopotamus in the room and, furthermore, insists that what they are doing is *asserting* this. An account has been offered in which negation is understood in terms of denial, but someone who is not impressed by the necessity of such an account might insist that they are nevertheless able to assert negatives. I could try to claim that what they are actually doing, in such cases, is denying a positive. But is this, then, just a stalemate between two rival interpretations of the utterance? Is there any principled way of deciding which interpretation is right? Surely a statement is not an assertion or a denial simply because its speaker thinks or says it is. Cannot

speakers be mistaken? However, the argument is not really in stalemate. What was provided in Table 10.1—the table of assertion and denial—and then explained in detail, can effectively be used as criteria for distinguishing assertions from denials. Someone might claim to be asserting something, just because they expressed it. But if what is said is uncommitted and indeterminate, if it avoids the risk of error and carries a low epistemic responsibility, if it is not necessarily bound to the law of non-contradiction, then, on what was argued above, it is a denial whether the speaker understands it as such or not. Merely insisting that one's statement is an assertion is not enough. What we have offered allows us to tell assertion and denial apart on the basis of principled criteria. Of course, there could be philosophical quibbles with the criteria I have offered. But I do insist that this is the right way to proceed. Let an opponent offer better criteria, if they can, and justify them. We must move away from merely insisting or assuming without argument when an utterance is being asserted. Our criteria allow any such assumption to be assessed and overruled, if it is appropriate to do so.

A lot of detail has been provided in this chapter. Perhaps it has been too much. But I accept that the claims made here will challenge conventional thinking and so it was perhaps appropriate to provide a weight of argument that corresponded to the weight of assumption on the other side. Having established our conclusion, we are now in a position to apply it to the remaining problem concerning our negative beliefs.

Negative Belief

11.1 The final reckoning

We are now in a position to pass judgement on our soft Parmenidean project or, rather, two judgements, since we noted that there were ontological and methodological strands to it. The preceding chapters have provided enough resources for us to make an informed assessment in both instances. We need to consider whether some but not all of Parmenides' claims are sustainable. This will allow us to see whether we should support any form of Parmenideanism at all and, if so, whether that should be a soft or the hard version. To form that judgement, we should also consider how well the soft Parmenidean methodology has served us and whether it has identified any ineliminably negative entities that are indispensable in our best theories.

It might be recalled (§1.4, above) that the aim was to approach the problems before us with a preference for avoiding negative entities or any such negative elements of reality while also accepting that there would be no choice but to do so if certain conditions were met, namely, that such things were indispensable to our best theories and could not be reduced into other non-negative terms. The approach was Parmenidean in that it sought to avoid negatives but soft in that it would countenance them if these conditions were met. The approach was to an extent conciliatory, then, since it would not refuse negatives in every circumstance, irrespective of the weight of argument. With that reminder of our purposes, we can look again over the territory that we have covered.

11.2 Inventories of being and non-being

We considered putative negatives within a number of categories. First, we looked at alleged negative properties (in chapter 2), such as being non-red. We then considered a host of apparent nonentities or negative particulars (chapter 3), including limits and boundaries, privations, shadows, holes, omissions, negative mathematical and logical entities, and negative epistemic states. These two chapters covered negatives from many of the traditional ontological categories: negative particulars and properties, but at least some more besides. As well as negative entities, however, there are temptations to accept other negative aspects of reality. Two cases were considered in detail. The first was causation by absence

(chapter 4), which coupled with a certain view of what is real would seem to imply that absences are real. The second was possibilities (chapter 5). We have a sense that some mere possibilities are real and some are not, and the question then is whether the sense in which mere possibilities are real commits us to saying that something that is not also is, which would violate the Parmenidean strictures. It might be recalled that it was not possible to resolve all of these matters entirely using metaphysical arguments alone, since some of the discussions were inconclusive or were found to contain circular reasoning. Some reminders will help, during which we can look at the strengths and weaknesses of the Parmenidean approach.

11.2.1 Possibilities

In respect of mere possibilities, we saw that a number of theories were in danger of reifying the mere possibilities that are not actual. Both modal realism and dispositionalist approaches have versions in which what is merely possible nevertheless has some sort of being, either elsewhere in the plurality of worlds or in the dispositions or powers of things. As an alternative, I offered a fictionalist view of mere possibilities in which they had no reality at all but ‘existed’ only as fictions. This view would be assisted if we had a better understanding of what a fiction is and how it is possible to talk about something that is fictional only, but we did, of course, return to that in chapter 8. Now there might be an objection at this point that the soft Parmenidean approach was to consider whether there were irreducibly negative entities that were indispensable parts of our best theories and—so goes the contention—fictionalism is not the best theory of possibility. Perhaps, someone could argue, the leading theory of possibility comes as part of David Lewis’s (1986a) modal realism. What would it matter if there was another theory—one that did not reify mere possibilities—if it was not the best theory?

Such an objection is not decisive. It is worth stating why it is not since the same considerations apply to a number of the topics we have covered. A first consideration is that whether any particular theory of modality is obviously the best possible is a matter that is still contested. It is to be noted that we are not seeking merely the best available theories, since in some areas they might be far from fully developed and ideal. We are asking, rather, whether the best possible theory of some important phenomenon would contain a commitment to negative entities. This matter might seem complicated in the present case since we are looking for the best possible theory of what is possible. Hopefully, we have enough about us to comprehend that notion. The fictionalist theory of possibility I have offered here is manifestly not as developed as other theories have been. The question is whether it could be. How we should answer that leads to the second consideration.

The second consideration is how we would make our judgements about what is the best theory within some domain. It seems entirely legitimate, for instance, to bring in concerns from outside that domain since there are reasons to accept that theory choice is often, and ought to be, a holistic matter to at least some extent (such as in Kuhn 1962). One does not accept a theory as best merely because it handles well its own internal problems and respects and accommodates the phenomena within its delimited field of study. One should also consider how well its commitments would cohere with our wider theories. For example, a theory in which the motions of bodies were controlled by animal spirits (for example, in Newton 1687) might be internally coherent, but the reason it would be rejected as a theory of motion is that it conflicts with our broader naturalistic commitments about the nature of the world. We do not have any reason to believe in animal spirits other than the gap they might fill within that theory. Returning to questions of modality, then, it is likely that we will judge which is the best theory of possibility only once we have introduced wider considerations. This gives us a further opportunity for our soft Parmenidean commitments to have influence in the final theory choice. If one looks at the best theory of the world as a totality, including all its subdomains, then a theory that avoids reification of negative entities will be preferable, other things being equal, to a theory that has them. This would allow us to say that even if a theory of *X* looks most suitable for the *X* subdomain, were it to be taken in isolation, it is not the best theory of *X* in the final analysis because it does not cohere with our other commitments. It might, of course, seem prejudicial, in making an assessment of the overall best theory of the world, to introduce the Parmenidean aversion to negatives since we are trying to apply a test to see whether negatives should be admitted: a test in terms of best possible theories. However, soft Parmenideanism does not hide its bias against the negative, in this respect. There is certainly no compunction against using the Parmenidean consideration as a tiebreaker, for example. If two theories are equally good and one admits negative entities while the other does not, then we prefer the theory that does not.

There is a final consideration in respect of the mere possibilities that really does change the terms of this discussion, however. This is that the rival theories to fictionalism effectively dissolve the problem of mere possibilities, from a Parmenidean viewpoint, by taking them into what is. What the Parmenidean position abhors is allowing that what-is-not also is (Pi, of §1.3). The Lewisian and dispositionalist views that we considered effectively declare that what was thought to be negative—that is, not any part of reality—is part of reality after all. We have to be careful how we state this, but it does seem a tenable interpretation of those theories. As is well known, ‘actual’ is no more than an indexical in the theory of modal realism, hence the reason we do not say that a mere possibility is actual is that it has no reality in our world. As is equally well known, however, that same mere possibility is indeed actual at all the other worlds in which it is true. Those

other worlds are just as real as ours: it is only that we do not inhabit them. If one takes reality to be constituted by the whole plurality of worlds, then what we think of as mere possibilities come out decidedly as part of that total reality. Possibilities have made a transition from non-being to being, within that account, then. Similarly for the dispositionalist theory that we considered, if what is possible is nested within what is actual, then we do not have a case where something that is not also is. Within these theories, what is possible has become part of what is. Of course, one might criticize such theories precisely on these grounds: namely that they are no longer about the target phenomena, which was something that was possible without being actual. In that respect, they have not explained mere possibilities but have merely reconceived them as part of what is actual by broadening the scope of the actual. If there is anything in that criticism, we will of course take that as a point in favour of our fictionalist approach. The import of this final consideration, however, is that those rival accounts of possibility still do not treat what is not as something that is, hence they do not really violate Parmenideanism. We might then have to reassert the opening intuition that mere possibilities are not a part of what is, even if they can be accounted for in terms of what is. The response of our rivals has been to abandon the view that mere possibilities are not.

The examination of possibility, then, provides us with no compelling reason to accept negative entities and our soft Parmenideanism remains unscathed. We can now apply it to the next problem.

11.2.2 Causation

Causation by absence was seen as a problem because there appear to be causes—again in our best explanation of what occurs—that are irreducibly negative. They concern what is not the case: what is not there. If so, then it looks like we will have to reify those absences and, thus, what is not will be part of what is.

At the conclusion of chapter 4, we found that absences can be explanatorily useful but that we need not invoke them as actual causes since absences are powerless. Absences can figure in non-causal explanations where such explanations gain traction through their selection of a pertinent counterfactual. The absence of *Y* explains *E* where there is a true counterfactual that had *Y* occurred then *E* would not have occurred. For example, the absent umbrella does not cause you to get wet, nor causally explain why you get wet, but it non-causally explains you getting wet because of the truth of the counterfactual that if you'd had an umbrella, you would have stayed dry.

This account is metaphysically satisfactory. It does, however, still invoke absences in its explaining away of causation by absence. There is no causation by absence, but there is explanation by absence. The use of absences in explanation are adequate grounds for a hard Parmenidean to object since it violates the

Parmenidean principle Piv. We have to be able, at least, to think about and name an absence for it to play a role in explanation. Whether an absence can really hold explanatory value remains dubious, however. Although we may have made some progress on the question of causation by absence, then, we cannot say that it is complete until we have resolved these further matters. We will return to them in §11.3, below.

11.2.3 Nonentities

In chapter 3, a collection of apparent nonentities were gathered. Many of them were dealt with there in a satisfactory way. Fictional and mythical objects need have no being, for example. It had to be admitted, however, that there seems to be reference to such things, at least, so it was conceded that an account of empty terms was due. This will be also the case with other nonentities, such as the empty set and zero. We were content to conclude that there is nothing that they name, strictly speaking, but this leaves us with the same problem of how there can be a name for nothing. Related to this is the question of how words can name different things if they all name nothing. It sounds as if zero, the empty set, mermaids, and *Oliver Twist* would all have the same referent, since there is only one nothing. An explanation was thus a pressing matter.

We found that there were some nonentities that raised a different kind of issue. There was no need to claim that holes were negative entities since they seem real enough. An absent hole would be a different matter, however, and there is no reason to believe that there are such things. This does not stop me from sometimes wishing that there were a hole somewhere that there isn't, for instance, where there should have been a hole punched in some paper, for filing purposes, but there wasn't. But real holes, which it seems there are, raise the issue that they at least require the absence of something. Holes have hosts, and they are real, but they also involve a discontinuity. This discontinuity usually involves empty space, though not necessarily as some holes are filled. Empty space means that the hole's host must somewhere cease to be. Likewise, other things have boundaries, edges, and limits, which seem to be points beyond which they are not. Armstrong's totality facts were also a 'no more'. Does this mean that what-is-not is part of what-is, as a precondition of what-is having any boundaries? Similarly, when we considered shadows, we found that they essentially involve an absence of light since for there to be a shadow there has to be the partial blocking of a light source.

As Parmenideans, we can stick to the claim that there is only what is. There is an ambiguity here, however. We might say that everything exists, but because the only way for something to be a thing is for it to exist. This does not mean that unicorns exist, or square circles. When we accept that reality is a single, complete plenum (Pv), then, we might say that there is nothing that does not exist, but this

is only meant trivially. If we had a way of overcoming Parmenides' Pvi claim that we cannot name or think about nothing, then it would allow us to name something that did not exist. This would allow us too to apply such an account to many of the above problems. We could name or think about zero, the empty set, an absence of light or a hole's host, and so on, without having to say that those things thereby existed as part of the plenum. For this reason, our study moved on to such matters of reference and aboutness.

11.2.4 Negative properties

We embarked on our soft Parmenidean investigations with the subject of negative properties. There we encountered a host of direct arguments for why there were no such things, but we also found that many of those arguments were question begging. For instance, Armstrong alleged that particulars that had the same property should thereby be similar: but two things could be non-red because one of them was blue and the other was yellow. Negative properties would violate the principle that a property is a One that runs through Many. These two non-reds would not be a One since they would be different colours. This was question begging, we saw, since only if we had already ruled out non-red as a property could we say that these two things had no colour in common. The way in which they would be similar is that they are non-red and non-red could be the One running through these two things, shared by them.

Chapter 2 was left more inconclusive than others, so now we have a duty to add something. Here, I hope the importance of what followed can be appreciated, particularly in the second half of the book. Methodological soft Parmenideanism requires that we consider whether the alleged negative entities are an indispensable part of our best theories, and in this case the theory in question is the theory of properties. We might have found no entirely convincing way of ruling out negative properties, but that need not matter if we can find a way of dispensing with them that does not create havoc with our remaining commitments. There is now a good prospect that precisely this can be done. To see this, however, we should gather together the evidence from the remaining chapters that featured in this book's second half.

11.3 How to believe that something is not

We need not allow that negative properties are real if we can get by without them, and it seems that there are indeed ways of doing so. Properties are real, we can grant, but only some properties. There need not be a property for every predicate, as Armstrong said. This allows us to jettison the putative negative properties if we

are able to provide an explanation of why such predicates are correctly applied. This we can do in terms of negative belief.

We have an account of how we can see that something is absent (from chapter 7). It is possible to see that something is not red, for example, if one is looking for something that is red. This could give rise to an expectation that is disappointed, for instance, and a feeling of something not being a certain colour. One person without that interest might just see a blue ball whereas another looks and sees a ball that is not red. This explains one simple kind of negative belief and why someone might use a negative predicate not-red. It is conceivable too that someone might want to separate all the red balls from those of every other colour, such as when a snooker referee gathers together all the reds to place in a triangle when setting up a new game. It is possible to look at the seven other balls and see them specifically as not-red when engaged in that activity.

That is not all we have at our disposal, however. We also considered, in some detail, the thorny problem of the truthmakers for negative truths. This problem too applies in the case of negative properties since it is commonly taken to be a possible truth that a particular thing is not-red, which seems like a negative truth. One matter is how we can come to know that something is not-red, and that is something that an account of perception of absence can provide. It is another matter, however, to consider how it can be true that something is not red, or that there is not a hippopotamus in the room, or that my pockets are empty and so on. It was suggested that a good explanation of the perception of absence was that there is an evolutionarily useful non-rational mechanism in the mind that draws such conclusions non-deductively and presents them to the mind's user in a phenomenological way. This does not require something in reality that is a property of non-red, nor a fact, nor any other sort of negative existent. How, then, can it be true that this thing is not-red if there is no such existent? What would be the truthmaker of such a claim?

In chapter 10, however, we have seen that truthmakers for the negative truth are effectively dispensable from that theory of truth. We struggled to find solutions proper to the Molnar problematic: that is, we could find no positive truthmakers for the negative truths. Others, having noted this failing already, seemed willing to countenance negative entities, in which case the soft Parmenidean programme will have hit a major obstacle. Such a 'solution' would have been a rejection of Molnar's second claim, (Mii). There are other possibilities, however, and these include a rejection of (Miii), that some negative claims about the world are true. We need not rest all our hopes on a theory of falsehood, as has been explored previously (such as Mumford 2007a). Instead, I suggested, we should consider denial as a separate act from assertion. The work of differentiating assertion and denial (mainly ventured in chapter 10) allows us to say that the assertion of negative truths ought to be reconsidered instead as denial. Hence, instead of asserting that the ball is not-red, one can deny that it is red, where denial requires no

ontological commitment to a truthmaker or falsemaker. Denials are correct when there is nothing that would make the statement in question true, but this is not an ontological commitment to any negative entity, such as the absence of a truthmaker. We could deny the presence of a truthmaker instead.

It would be a problem for such an account if there were no plausible theory of empty reference, since how can one say of a hippopotamus that it is not there in the room since there is no such hippopotamus. This is why an account of empty reference, or an alternative to empty reference, was necessary and we have good grounds for thinking that there is an acceptable account for the soft Parmenidean: an account that lets us speak about things that are not there. As detailed, in chapter 8, the preference here is for a form of fictionalism in which we make fictional, pretended references to things that we know do not exist, or we just fail to refer in cases where we are mistaken in believing that things exist when they don't. In either case, however, our words can be about things, since one can think about things that do not exist even if one cannot refer to them.

Armed with denial, the capacity for perception of absence, and an alternative to empty reference, we are in a position then to believe that something is not without ontologically committing to the existence of what is not. We can say, then, that this epistemological and linguistic investigation is what allows us to clear up the remaining matters for which there was no immediate metaphysical solution. We see now that there are no irreducibly negative indispensable entities within the theories that we have considered, since we can indeed dispense with them by dealing with them in terms of negative beliefs. There is no property of being non-red, but it is possible to believe that something is not red since we can and should at times deny that something is red, taking negation, as argued at the end of chapter 10, as an indicator of denial. Similarly, absences are not causes, but they can figure in explanations. How so? One can say that there is no oxygen in a space to explain a suffocation or no screw in a chair to explain why it collapsed. Such explanations are good, but they do not commit to parts of reality such as absences of oxygen and missing screws. We are tempted to say that there exists only a belief that the oxygen and the screw are absent, but we now see that these beliefs themselves permit further analysis: as a denial that oxygen is present and that a screw is present. Such an account sits well with the other component in our explaining-away of causation by absence. We deny that oxygen is present while also accepting that had oxygen instead been present, then the victim would have lived. Likewise, the absence of the host, at the place that is its hole, is not a part of reality. Instead of asserting that the host is absent, at its boundary, we deny that it is present there, better reflecting the metaphysical reality that there is only what is. We can offer a similar treatment of shadows, and so on. We can then also preserve the idea that there are different nothings; that is, the absence of light is a different thing from the absence of a hippo. The denial of one does not exclude the presence of the other because they are about, and thus deny, different things.

The availability of denial bears also on the topics of universal quantification and totality facts. We need not assert that there are no more (first-order) facts when we want to total some set of facts. We can deny that there are more facts, which needs nothing in the world to make it true. This would reverse the common judgement that universal quantification contains a negative component. Each person in the room could be a philosopher, and if we deny that there are more people in the room than these, then we are in a position to say that everyone in the room is a philosopher.

It might at this stage be objected that the analysis works only by committing some kind of egregious category mistake. Our initial Parmenidean concern was to deny negative entities or elements as part of reality. It looked as if there were absences, for example, and that they played a role in our explanations of various phenomena, such as causation and shadows. The account offered of them here has been not as things but as a different category altogether: beliefs, denials, and so on. What legitimizes the shift from an ontological discussion to an epistemological and logico-linguistic one?

In reply, it should be pointed out that this sort of shift is entirely the one to be expected of a soft Parmenidean position. It would not be a move open to a hard Parmenidean. However, our conclusion simply reflects the soft Parmenidean view that negativity is (A) not a feature of the world itself but (B) a feature of the way we think and talk about the world. The soft Parmenidean can accept both the components, (A) and (B), of this claim. A hard Parmenidean can accept (A) but cannot accept (B), since it is inconsistent with their commitment to Pvi (from §1.3). This shows us a major divergence between the soft and hard varieties of Parmenideanism as a position held, rather than just a methodology. Further, it seems that the soft version has the advantage and good sense on its side. What-is-not has no existence. It is not a part of what is. We said in chapter 1 that this was almost, if not quite, a trivial truth. But it is clear that we can and do talk about what there is not. This also seems a manifest truth, and we saw that Parmenides' claim to the contrary, Pvi, could only really be defended if it was taken narrowly to be about the real relation of reference specifically. Now, given that it seems we can talk and think about what-is-not, but what-is-not is not any part of what-is, then it seems that the best thing to say about what-is-not is that we can only talk and think about it, which gives us our conclusions (A) and (B) that what-is-not is only a part of what we think and say and not a part of what-is. Negatives are a part of language and thought and not the world.

Seeing it this way may allow us to explain one other quandary. We noted in discussing Molnar's (Mii)—everything that exists is positive (see §9.2, above)—that 'positive' seems pleonastic if attached to 'existent' in a Parmenidean framework. There is no cleavage in reality that corresponds to a positive/negative division so everything that exists is 'positive' trivially. Clearly, however, there is such a division in language even if it is not always entirely transparent. There are

polarities among many terms (a point exploited in Beall 2000) where one is understood as negative and the other positive. Since neither the word 'not', nor giveaway suffixes (-less) and prefixes (un-), need always be used, it might require some metaphysical judgement to know which is the positive and which is the negative, as with the pair 'empty' and 'full'. And there are some oppositions where, although one is the negation of the other, and polar opposites in that sense, it is not clear that one is a negative and the other a positive at all, as with 'crooked' and 'straight'. We can thereby say positive or negative things about the world and, as argued in chapter 10, credibly venture that these reflect the two differing acts of assertion and denial. Doing so does not require applying the same division to the world. We can thus understand the Parmenidean view that says that what exists is everything: and nothing is not. To say that all that exists is positive, as with Molnar's (Mii), does seem redundant, then, which is not to deny that it is useful to articulate the issue this way in problematizing negative truth.

Such matters considered, we are in a position to reassert that there are no negatives in the world. We can have negative beliefs about the world and they can help us to explain away what look to be otherwise recalcitrant metaphysical difficulties.

11.4 Types of negative belief

It would be remiss at this point not to acknowledge that negative beliefs admit of variety. They do not form a monolith, as Apostel (1972) and many others demonstrate. For example, we have already seen that there is a difference between believing not-P and not believing P. Even further distinctions can be made here.

Among cases of not believing that-P, I can think of scepticism and ignorance. A sceptic chooses not to believe P, presumably on the basis of some principle, either general or specifically concerning the case of P. The ignorant can fall into two further categories. It is standard to think of ignorance as involuntary since in a rational tradition it is assumed that everyone wants to know and thus no one would purposefully not know. However, it also seems credible that there is a category of wilful ignorance where someone chooses to not know, such as an exam candidate who refuses to look at his results.

Among cases of believing not-P we can distinguish external and internal negations. One can believe that it's not the case that P or one can believe that it's the case that not-P. Both of these I take to be forms of denial, but the distinction shows that once we use negation instead of denial we may be able to express more clearly with it. Negation allows us the agility to deny a predicate of a particular or to deny a whole statement or fact of what is the case. How to exactly divide up the possibilities of denial and negation are themselves historically contested matters, of course, as Horn (2001: ch. 1) demonstrates. Other types of negative belief that

we might distinguish include negative existentials—denying that something exists—and general truths, which involve the denial that there is more than there is.

One further issue, that I will bracket and leave alone here, is the issue of believing falsely, since this is not the same as believing negatively. The issue, which occupied Plato in the *Sophist*, for instance, need not detain us, now that we are at liberty to speak of what is not. We can follow pretty much an Aristotelian account (*Metaphysics* 1011b25) in which to believe falsely is either to believe of what is not that it is or of what is that it is not.

11.5 Ontological Parmenideanism

Our final task is to review the fate of our soft ontological Parmenideanism. The ontological version of the theory was a position in which some though not all of the following claims were held:

- Pi. Nothing is not (A)
- Pii. There are no degrees of being (A)
- Piii. Nothing comes from nothing (A/R)
- Piv. Non-being is unknowable (A/R)
- Pv. Reality is a single plenum (A/R)
- Pvi. Non-being is unthinkable and unnameable (R)
- Pvii. Neither motion nor change is possible (R)

Having conducted our review of several of these issues, the decisions on whether to accept (A), reject (R), or either accept or reject (A/R), depending on the interpretation, can now be understood. Equal consideration has not been given to all these matters, of course. Pi has effectively had the most scrutiny, since our soft methodological approach has largely been an enquiry into whether there are grounds on which to relinquish Pi. As reported above, it seems that we have found none that are compelling. Of course, not every single matter under the sun has been considered within this purview and there remains the possibility that some irreducibly negative and indispensable entity comes to light, in which case we will have to revise our assessment. Until then, however, we are in a position to declare soft ontological Parmenideanism validated. In the preceding chapters, I have explained specifically why we should uphold Pi and Pii, why Piii, Piv, and Pv should be upheld under certain interpretations, though rejected under others, and why Pvi and Pvii should be rejected. The careful reader will have noted by now that the rejection of Pvi has been crucial to us in upholding Pi since it has allowed us to show how one can talk about nothing and what is not without being committed to its being or existence. This is overall more credible than a hard

Parmenidean line that we cannot even think about what is not. Pvii has not really been taken seriously in this book, I will admit. It has been dismissed, though I hope not completely out of hand. By allowing ourselves to talk about nothing, when we reject Pvi, we can worry less about the possibility of change. Parmenides' concern was that change involved either what was not coming-to-be, or what was coming not-to-be. We are now able to understand better how this could and should be interpreted in a way that makes no ontological commitment to absences, hence no violation of Pi. A soft Parmenidean can reject Pvii, therefore, if the foregoing account of negative belief is accepted.

11.6 Final words

These conclusions ought to be comforting, even though they should not come as a surprise. It is apparent that they leave things unchanged, since no revolutionary finding has been our reward. Nor was it our goal. The ideal outcome was to leave things unchanged. Had our soft Parmenidean project failed, that's when we would have been looking at a radical reappraisal of some deeply held beliefs. Our preservation of the status quo did not come as a result of fakery or prejudice, however. It was conceded at the outset that negative entities would be admitted if certain specific conditions were met. Having examined the problems and the accompanying arguments, we are in a position to say that those conditions are not met and there is no need for negatives as part of the world.

I have sought to preserve the spirit of Parmenides in respect of his primary commitments. In doing so, I hold no particular brief for Eleatic or other ancient philosophy. Unless we are historians of philosophy, such ideas deserve our attention only insofar as they are relevant to our present concerns. What should have been clear in the main body of the current study is that Parmenides' main thoughts on these issues do indeed remain relevant to contemporary debates. Before, during, and after twentieth-century philosophy, continuing into present days, respectable philosophers have thought that they have identified an inescapable need to accept negative entities of some kind in order to preserve their favoured theories. Still now we see negative facts, negative instantiations, absences as causes, and negative properties offered as a way of resolving recalcitrant problems, filling the theoretical gaps where something is wanting. While Parmenides is still relevant, however, he is also still largely right. There is nothing that is negative and also exists in non-linguistic, non-intentional reality. What-is-not cannot also be. It has been necessary, then, to revive our old Parmenidean philosophy and reassert that there is only what exists. In particular, our various more detailed discussions have shown that there are no compelling reasons to accept absences or any other sweet nothings as part of the world there is.

References

- Anjum, R. L. and Mumford, S. (2018) *What Tends to Be: The Philosophy of Dispositional Modality*, London: Routledge.
- Anjum, R. L., Lie, S. A., and Mumford, S. (2012) 'Dispositions and Ethics', in R. Groff and J. Greco (eds), *Powers and Capacities in Philosophy: The New Aristotelianism*, London: Routledge, pp. 231–47.
- Apostel, L. (1972) 'Negation', *Logique et Analyse*, 15: 209–317.
- Aristotle, *Categories*, J. M. Ackrill (ed.), Oxford: Oxford University Press, 1975.
- Aristotle, *Metaphysics*, trans. H. Lawson-Tancred, London: Penguin, 1998.
- Armstrong, D. (1968) *A Materialist Theory of the Mind*, rev. edn, London: Routledge, 1993.
- Armstrong, D. (1978a) *Nominalism and Realism*, Cambridge: Cambridge University Press.
- Armstrong, D. (1978b) *A Theory of Universals*, Cambridge: Cambridge University Press.
- Armstrong, D. (1980) 'Identity Through Time', in P. van Inwagen (ed.), *Time and Cause*, Dordrecht: Reidel, pp. 67–78.
- Armstrong, D. (1983) *What is a Law of Nature?*, Cambridge: Cambridge University Press.
- Armstrong, D. (1989) *A Combinatorial Theory of Possibility*, Cambridge: Cambridge University Press.
- Armstrong, D. (1997) *A World of States of Affairs*, Cambridge: Cambridge University Press.
- Armstrong, D. (1999) 'The Open Door: Counterfactual versus Singularist Theories of Causation', in H. Sankey (ed.), *Causation and Laws of Nature*, Dordrecht: Kluwer, pp. 175–86.
- Armstrong, D. (2004) *Truth and Truthmakers*, Cambridge: Cambridge University Press.
- Armstrong, D. (2010) *Sketch for a Systematic Metaphysics*, Cambridge: Cambridge University Press.
- Austin, J. L. (1950) 'Truth', in *Philosophical Papers*, 2nd edn, Oxford: Oxford University Press, 1970, pp. 117–33.
- Austin, J. L. (1962) *Sense and Sensibilia*, Oxford: Oxford University Press.
- Ayer, A. J. (1936) *Language, Truth, and Logic*, London: Victor Gollancz.
- Ayer, A. J. (1952) 'Negation', *Journal of Philosophy*, 49: 797–815.
- Azzouni, J. (2010) *Talking About Nothing*, Oxford: Oxford University Press.
- Bacon, F. (1620) *The New Organon*, L. Jardine and M. Silverthorne (eds), Cambridge: Cambridge University Press 2000.
- Baehr, J. (2011) *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*, Oxford: Oxford University Press.
- Balashov, Y. (1999) 'Zero-Value Physical Quantities', *Synthese*, 119: 253–86.
- Balashov, J. M. (1928) 'Negation' and 'Negative', *Dictionary of Philosophy and Psychology*, vol. 2, New York: Macmillan, pp. 146–9.
- Baldwin, T. (1996) 'There Might Be Nothing', *Analysis*, 56: 231–8.
- Barker, S. and Jago, M. (2012) 'Being Positive about Negative Facts', *Philosophy and Phenomenological Research*, 85: 117–38.
- Barnes, J. (ed.) (1987) *Early Greek Philosophy*, London: Penguin.
- Barrow, J. (2000) *The Book of Nothing*, London: Jonathan Cape.

- Beall, J. C. (2000) 'On Truthmakers for Negative Truths', *Australasian Journal of Philosophy*, 78: 264–8.
- Beebe, H. (2004) 'Causing and Nothingness', in J. Collins, N. Hall, and L. Paul (eds), *Causation and Counterfactuals*, Cambridge, MA: MIT Press, pp. 291–308.
- Bergson, H. (1911) *Creative Evolution*, trans. A. Mitchell, New York: Modern Library.
- Bernstein, S. (2015) 'The Metaphysics of Omissions', *Philosophy Compass*, 10: 208–18.
- Black, R. (2000) 'Nothing Matters Too Much, or Wright is Wrong', *Analysis*, 60: 229–37.
- Borghini, A. and Williams, N. (2008) 'A Dispositional Theory of Possibility', *Dialectica*, 62: 21–41.
- Bosanquet, B. (1988/1911) *Logic*, I, 2nd edn, Oxford: Clarendon Press.
- Bradley, F. H. (1883) *Principles of Logic*, London: Kegan Paul and Trench.
- Braun, D. (1993) 'Empty Names', *Noûs*, 27: 449–69.
- Brentano, F. (1874) *Psychology from an Empirical Standpoint*, L. L. McAlister (ed.), Atlantic Highlands, NJ: Humanities Press, 1973.
- Buchdahl, G. (1961) 'The Problem of Negation', *Philosophy and Phenomenological Research*, 22: 163–78.
- Carroll, L. (1871) *Through the Looking-Glass and What Alice Found There*, in M. Gardner (ed.), *The Annotated Alice: The Definitive Edition*, Harmondsworth: Penguin, 2000.
- Carroll, L. (1895) 'What the Tortoise Said to Achilles', *Mind*, 4: 278–80.
- Casati, R. (2004) *Shadows*, New York: Vintage.
- Casati, R. and Varzi, A. (1994) *Holes and Other Superficialities*, Cambridge, MA: MIT Press.
- Cavedon-Taylor, D. (2017) 'Touching Voids: On the Varieties of Absence Perception', *Review of Philosophy and Psychology*, 8: 355–66.
- Chalmers, D. (1996) *The Conscious Mind*, Oxford: Oxford University Press.
- Cheyne, C. and Pigden, C. (2006) 'Negative Truths from Positive Facts', *Australasian Journal of Philosophy*, 84: 249–65.
- Clarke, R. (2018) 'Absence Causation for Causal Dispositionalists', *Journal of the American Philosophical Association*, 4: 323–31.
- Coggins, G. (2010) *Could There Have Been Nothing?*, Basingstoke: Palgrave Macmillan.
- Crane, T. (2013) *The Objects of Thought*, Oxford: Oxford University Press.
- Crivelli, P. (2012) *Plato's Theory of Falsehood*, Cambridge: Cambridge University Press.
- Davies, P. (2013) 'The Day Time Began', in J. Webb (ed.), *Nothing*, London: Profile Books, pp. 44–55.
- Della Rocca, M. (2020) *The Parmenidean Ascent*, Oxford: Oxford University Press.
- Demos, R. (1917) 'A Discussion of Certain Types of Negative Propositions', *Mind*, 26: 188–96.
- Dennett, D. (1991) *Consciousness Explained*, London: Penguin.
- Dennett, D. (2017) *From Bacteria to Bach and Back*, London: Penguin.
- Diogenes Laertius, *Lives of the Eminent Philosophers*, trans. P. Mensch, ed. J. Miller, New York: Oxford University Press, 2018.
- Divers, J. (1999) 'A Genuine Realist Theory of Advanced Modalizing', *Mind*, 108: 217–39.
- Divers, J. (2002) *Possible Worlds*, London: Routledge.
- Dodd, J. (2002) 'Is Truth Supervenient on Being?', *Proceedings of the Aristotelian Society*, 102: 69–86.
- Dodd, J. (2007) 'Negative Truths and Truthmaker Principles', *Synthese*, 156: 383–401.
- Donnellan, K. (1966) 'Reference and Definite Descriptions', *The Philosophical Review*, 75: 281–304.
- Donnellan, K. (1974) 'Speaking of Nothing', *Philosophical Review*, 83: 3–31.
- Dowe, P. (2001) 'A Counterfactual Theory of Prevention and "Causation" by Omission', *Australasian Journal of Philosophy*, 79: 216–26.

- Dretske, F. (1969) *Seeing and Knowing*, Chicago: University of Chicago Press.
- Dummett, M. (1973) *Frege: Philosophy of Language*, London: Duckworth.
- Dyke, H. (2008) *Metaphysics and the Representational Fallacy*, New York: Routledge.
- Efrid, D. and Stoneham, T. (2005) 'Genuine Modal Realism and the Empty World', *European Journal of Analytic Philosophy*, 1: 21–36.
- Farennikova, A. (2013) 'Seeing Absence', *Philosophical Studies*, 166: 429–54.
- Farennikova, A. (2015) 'Perception of Absence and Penetration from Expectation', *Review of Philosophy and Psychology*, 6: 621–40.
- Field, H. (1980) *Science Without Numbers: A Defence of Nominalism*, Oxford: Blackwell.
- Frege, G. (1879) *Begriffsschrift*, in P. T. Geach and M. Black (eds), *Translations from the Philosophical Writings of Gottlob Frege*, 3rd edn, Oxford: Blackwell, pp. 1–20.
- Frege, G. (1884) *The Foundations of Arithmetic*, trans. J. L. Austin, Oxford: Basil Blackwell, 1950.
- Frege, G. (1892) 'On Sense and Meaning', in P. Geach and M. Black (eds), *The Philosophical Writings of Gottlob Frege*, 3rd edn, Totowa, NJ: Barnes and Noble, 1980.
- Frege, G. (1895) 'Critical Elucidation of Some Points in E. Schröder's *Lectures on the Algebra of Logic*', *Translations from the Philosophical Writings of Gottlob Frege*, 3rd edn, Oxford: Blackwell, pp. 86–106.
- Frege, G. (1919) 'Negation', in P. T. Geach and M. Black (eds), *Translations from the Philosophical Writings of Gottlob Frege*, 1st edn, Oxford: Blackwell, pp. 117–35.
- Gale, R. (1970) 'Negative Statements', *American Philosophical Quarterly*, 7: 206–17.
- Gale, R. (1976) *Negation and Non-Being*, Oxford: Blackwell.
- Gallo, D. (1984) 'Introduction', in Parmenides, *Fragments*, trans. D. Gallop, Toronto: University of Toronto Press, pp. 3–28.
- Geach, P. T. (1980) *Logic Matters*, 2nd edn, Berkeley: University of California Press.
- Gibb, S. (2013) 'Mental Causation and Double Prevention', in S. Gibb, E. J. Lowe, and R. D. Ingthorsson (eds), *Mental Causation and Ontology*, Oxford: Oxford University Press, pp. 193–213.
- Gibson, J. J. (1966) *The Senses Considered as Perceptual Systems*, London: George Allen and Unwin.
- Givón, T. (1979) *On Understanding Grammar*, New York: Academic Press.
- Gorgias, *On What-is-Not*, in D. W. Graham (ed.), *The Texts of Early Greek Philosophy*, Cambridge: Cambridge University Press, 2010, pp. 725–88.
- Graham, D. W. (2010) 'Introduction', in *The Texts of Early Greek Philosophy*, Cambridge: Cambridge University Press, pp. 1–14.
- Green, S. (2001) 'Lying, Misleading, and Falsely Denying: How Moral Concepts Inform the Law of Perjury, Fraud, and False Statements', *Hastings Law Journal*, 53: 157–212.
- Grice, P. (1961) 'The Causal Theory of Perception', *Proceedings of the Aristotelian Society*, supp. vol. 35: 121–68.
- Grice, P. (1975) 'Logic and conversation', in P. Cole and J. Morgan (eds), *Syntax and Semantics*, 3: *Speech Acts*, New York: Academic Press, pp. 41–58.
- Griffith, A. (2015) 'How Negative Truths are Made True', *Synthese*, 192: 317–35.
- Grossmann, R. (1974) *Meinong*, London: Routledge & Kegan Paul.
- Haldane, J. (2011) 'Identifying Privative Causes', *Analysis*, 71, pp. 611–19.
- Hanson, N. R. (1958) *Patterns of Discovery*, Cambridge: Cambridge University Press.
- Hart, H. L. A. and Honore, A. (1985) *Causation in the Law*, 2nd edn, Oxford: Oxford University Press.
- Hawking, S. and Mlodinow, L. (2010) *The Grand Design*, New York: Random House.
- Hawley, K. (2001) *How Things Persist*, Oxford: Clarendon Press.
- Hegel, G. (1812) *Science of Logic*, trans. A. V. Miller, London: George Allen & Unwin, 1969.

- Heidegger, M. (1929) 'What is Metaphysics?', in D. Farrell Krell (ed.), *Basic Writings: Martin Heidegger*, London: Routledge, 1978, pp. 93–110.
- Heinemann, F. H. (1944) 'The Meaning of Negation', *Proceedings of the Aristotelian Society*, 44: 127–52.
- Heraclitus, *Fragments*, trans. T. Robinson, Toronto: University of Toronto Press, 1987.
- Hesiod, *Theogony in Theogony and Works and Days*, trans. M. L. West, Oxford: Oxford University Press, 1988.
- Hicks, M. (2015) 'Pretense and fiction-directed thought', *Philosophical Studies*, 172: 1549–73.
- Hitchcock, C. (2006) 'What's Wrong with Neuron Diagrams?', in J. K. Campbell, M. O'Rourke, and H. Silverstein (eds), *Causation and Explanation*, Cambridge, MA: MIT Press, 2006, pp. 69–92.
- Hofmann, F. and Horvath, J. (2008) 'In Defence of Metaphysical Analyticity', *Ratio*, 21: 300–13.
- Homer, *The Odyssey*, trans. A. Verity, Oxford: Oxford University Press, 2016.
- Hommen, D. (2013) 'Negative Properties, Real and Irreducible', *Philosophia Naturalis*, 50: 383–406.
- Hommen, D. (2016) 'Absences as Latent Potentialities', *Philosophical Papers*, 45: 401–35.
- Hommen, D. (2018) 'Making Sense of Negative Properties', *Axiomathes*, 28: 81–106.
- Horgan, T. and Potrč, M. (2008) *Austere Realism: Contextual Semantics Meets Minimal Ontology*, Cambridge, MA: MIT Press.
- Horn, L. (2001) *A Natural History of Negation*, Stanford: CSLI Publications.
- Humberstone, L. (2000) 'The Revival of Rejective Negation', *Journal of Philosophical Logic*, 29: 331–81.
- Hume, D. (1739) *A Treatise of Human Nature*, L. A. Selby-Bigge (ed.), Oxford: Clarendon Press, 1888.
- Hume, D. (1748) *An Enquiry Concerning Human Understanding*, P. Millican (ed.), Oxford: Oxford University Press, 2007.
- Hunt, I. (2005) 'Omissions and Preventions as Cases of Genuine Causation', *Philosophical Papers*, 34: 209–33.
- Jacobs, J. (2010) 'A Powers Theory of Modality, or How I Learned to Stop Worrying and Reject Possible Worlds', *Philosophical Studies*, 151: 227–48.
- Jacquette, D. (2015) *Alexius Meinong: The Shepherd of Non-Being*, Cham: Springer.
- Jenkins, C. and Nolan, D. (2012) 'Disposition Impossible', *Noûs*, 46: 732–53.
- Jespersen, O. (1917) *Negation in English and Other Languages*, Copenhagen: A. F. Høst.
- Joseph, H. (1916) *An Introduction to Logic*, 2nd edn, Oxford: Clarendon Press.
- Kachi, D. (2011) 'The Power of Holes', in M. Okada (ed.), *Ontology Meeting* (supplementary volume for 2011), Japan Open Research Center for Logic and Formal Ontology, pp. 7–12.
- Kalderon, M. E. (ed.) (2005) *Fictionalism in Metaphysics*, Oxford: Clarendon Press.
- Kant, I. (1781) *Critique of Pure Reason*, trans. N. Kemp-Smith, London: Macmillan, 1929.
- Kaplan, D. (1989) 'Demonstratives', in J. Almog, J. Perry, and H. Wettstein (eds), *Themes from Kaplan*, Oxford: Oxford University Press, pp. 481–565.
- Kripke, S. (1980) *Naming and Necessity*, Oxford: Basil Blackwell.
- Kroon, F. (2005) 'Belief about Nothing in Particular', in M. E. Kalderon (ed.), *Fictionalism in Metaphysics*, Oxford: Clarendon Press, pp. 178–203.
- Kuhn, T. (1962) *The Structure of Scientific Revolutions*, 2nd edn, Chicago, MA: University of Chicago Press, 1970.
- Lavelle, J. S., Botterill, G., and Lock, S. (2013) 'Contrastive Explanation and the Many Absences Problem', *Synthese*, 190: 3495–510.

- Leibniz, G. (1686) 'General Inquiries about the Analysis of Concepts and of Truths', in *Logical Papers*, G. H. R. Parkinson (ed.), Oxford: Clarendon Press, 1966, pp. 47–87.
- Lewis, D. (1973) *Counterfactuals*, Oxford: Basil Blackwell.
- Lewis, D. (1983) 'New Work for a Theory of Universals', in *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press, 1999, pp. 8–55.
- Lewis, D. (1986a) *On the Plurality of Worlds*, Oxford: Blackwell.
- Lewis, D. (1986b) *Philosophical Papers*, II, Oxford: Oxford University Press.
- Lewis, D. (1992) 'Critical Notice of D. M. Armstrong, *A Combinatorial Theory of Possibility*', *Australasian Journal of Philosophy*, 70: 211–24.
- Lewis, D. (2000) 'Causation as Influence', *Journal of Philosophy*, 97: 182–97.
- Lewis, D. (2004) 'Void and Object', in J. Collins, N. Hall, and L. Paul (eds), *Causation and Counterfactuals*, Cambridge, MA: MIT Press, pp. 277–90.
- Locke, J. (1690) *An Essay Concerning Human Understanding*, P. H. Nidditch (ed.), Oxford: Clarendon Press, 1975.
- Lucretius, *De Rerum Natura*, trans. R. E. Latham, *On the Nature of the Universe*, London: Penguin, 1951.
- Mabbott, J. D. (1929) 'Negation', *Proceedings of the Aristotelian Society*, supp. vol., 9: 67–79.
- Maddy, P. (1992) 'Indispensability and Practice', *Journal of Philosophy*, 89: 275–89.
- Maddy, P. (1995) 'Naturalism and Ontology', *Philosophia Mathematica*, 3: 248–70.
- Marmodoro, A. (2017) 'Aristotelian Powers at Work', in J. Jacobs (ed.), *Causal Powers*, Oxford: Oxford University Press, pp. 57–76.
- Marsili, N. (2018) 'Truth and Assertion: Rules versus Aims', *Analysis*, 78: 638–48.
- Martin, C. B. (1996) 'How It Is: Entities, Absences and Voids', *Australasian Journal of Philosophy*, 74: 57–65.
- Martin, C. B. (2008) *The Mind in Nature*, Oxford: Oxford University Press.
- Martin, J.-R. and Dokic, J. (2013) 'Seeing Absence or Absence of Seeing?', *Thought: A Journal of Philosophy*, 2: 117–25.
- Martinich, A. P. and Stroll, A. (2007) *Much Ado about Nonexistence*, Lanham, MD: Rowman and Littlefield.
- McGinn, C. (1991) *The Problem of Consciousness*, Oxford: Blackwell.
- McGrath, S. (2005) 'Causation by Omission: a Dilemma', *Philosophical Studies*, 123: 125–48.
- McGuinness, B. (1988) *Wittgenstein: a Life*, London: Duckworth.
- McTaggart, J. (1921) *The Nature of Existence*, Cambridge: Cambridge University Press.
- Meinong, A. (1904) 'Über Gegenstandstheorie', trans. as 'The Theory of Objects', in R. M. Chisholm (ed.), *Realism and the Background of Phenomenology*, Glencoe, IL: Free Press, 1960, pp. 76–117.
- Melamed, Y. (2012) "'Omnis determinatio est negatio": Determination, Negation and Self-Negation in Spinoza, Kant, and Hegel', in E. Förster and Y. Melamed (eds), *Spinoza and German Idealism*, Cambridge: Cambridge University Press, pp. 175–96.
- Menzies, P. (2004) 'Difference-making in Context', in J. Collins, N. Hall, and L. Paul (eds), *Causation and Counterfactuals*, Cambridge, MA: MIT Press, pp. 139–80.
- Molnar, G. (2000) 'Truthmakers for Negative Truths', *Australasian Journal of Philosophy*, 78: 72–86.
- Molnar, G. (2003) *Powers: a Study in Metaphysics*, S. Mumford (ed.), Oxford: Oxford University Press.
- Moore, M. S. (2009) *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics*, Oxford: Oxford University Press.
- Mulligan, K., Simons, P., and Smith, B. (1984) 'Truth-makers', *Philosophy and Phenomenological Research*, 44: 287–321.
- Mumford, S. (2004) *Laws in Nature*, London: Routledge.

- Mumford, S. (2007a) 'Negative Truth and Falsehood', *Proceedings of the Aristotelian Society*, 107: 45–71.
- Mumford, S. (2007b) *David Armstrong*, Chesham: Acumen.
- Mumford, S. (2008) 'Power, Dispositions, Properties: or A Causal realist Manifesto', in R. Groff (ed.), *Revitalizing Causality: Realism about Causality in Philosophy and Social Science*, London: Routledge, pp. 139–51.
- Mumford, S. (2011) *Watching Sport: Aesthetics, Ethics and Emotions*, London: Routledge.
- Mumford, S. and Anjum, R. L. (2011) *Getting Causes from Powers*, Oxford: Oxford University Press.
- Mumford, S. and Anjum, R. L. (2013) 'With Great Power Comes Great Responsibility', in B. Kahmen and M. Stepanians (ed.), *Causation and Responsibility—Critical Essays*, Berlin: de Gruyter, pp. 219–37.
- Newton, I. (1687) *Philosophiae Naturalis Principia Mathematica*, trans. A. Motte, *Mathematical Principles of Natural Philosophy*, 1729, revised by F. Cajori, Cambridge: Cambridge University Press, 1934.
- Nietzsche, F. (1883) *Thus Spoke Zarathustra*, trans. R. J. Hollingdale, London: Penguin, 1961.
- Parmenides, *Fragments*, trans. D. Gallop, Toronto: University of Toronto Press, 1984.
- Parsons, J. (2006) 'Negative Truths from Positive Facts', *Australasian Journal of Philosophy*, 84: 591–602.
- Parsons, T. (1980) *Nonexistent Objects*, New Haven: Yale University Press.
- Pears, D. (1976) 'The Causal Conditions of Perception', *Synthese*, 33: 25–40.
- Perszyk, K. (1993) *Nonexistent Objects: Meinong and Contemporary Philosophy*, Dordrecht: Springer.
- Plato, *Euthyphro*, in H. Tredennick (ed.), *The Last Days of Socrates*, London: Penguin, 1959, pp. 17–41.
- Plato, *Sophist*, in *Theaetetus and Sophist*, Christopher Rowe (ed.), Cambridge, Cambridge University Press, 2015, pp. 99–177.
- Price, H. (1990) 'Why "Not"?', *Mind*, 99: 221–38.
- Price, H. H. (1929) 'Negation', *Proceedings of the Aristotelian Society*, supp. vol., 9: 97–111.
- Priest, G. (1999) 'Perceiving Contradictions', *Australasian Journal of Philosophy*, 77: 439–46.
- Priest, G. (2005) *Towards Non-Being*, Oxford: Clarendon Press.
- Prior, A. N. (1967) 'Negation', in P. Edwards (ed.), *Encyclopedia of Philosophy*, vol. 5, New York: Crowell, Collier and MacMillan, p. 461.
- Putnam, H. (1979) 'What is Mathematical Truth', in *Mathematics Matter and Method: Philosophical Papers*, vol. 1, 2nd edn, Cambridge: Cambridge University Press, pp. 60–78.
- Quine, W. V. O. (1948) 'On What There is', in *From a Logical Point of View*, 2nd edn, Cambridge, MA: Harvard University Press, pp. 1–19.
- Quine, W. V. O. (1951) *Mathematical Logic*, New York: Harper and Row.
- Quine, W. V. O. (1969) 'Natural Kinds', in *Ontological Relativity and Other Essays*, New York: Colombia University Press, pp. 114–38.
- Quine, W. V. O. (1974) *Methods of Logic*, 3rd edn, London: Routledge & Kegan Paul.
- Quine, W. V. O. (1976) 'Carnap and Logical Truth', in *The Ways of Paradox and Other Essays*, rev. edn, Cambridge, MA: Harvard University Press, pp. 107–32.
- Quine, W. V. O. and Ullian, J. S. (1970) *The Web of Belief*, New York: Random House.
- Quinton, A. (1957–58) 'Properties and Classes', *Proceedings of the Aristotelian Society*, 58: 33–58.
- Raju, P. T. (1941) 'The Reality of Negation', *The Philosophical Review*, 50: 585–601.
- Richards, T. (1975) 'The Worlds of David Lewis', *Australasian Journal of Philosophy*, 53: 105–18.

- Roberts, T. (2016) 'A Breath of Fresh Air: Absence and the Structure of Olfactory Perception', *Pacific Philosophical Quarterly*, 97: 400–20.
- Rodriguez-Pereyra, G. (1997) 'There Might Be Nothing: The Subtraction Argument Improved', *Analysis*, 57: 159–66.
- Rosen, G. (1990) 'Modal Fictionalism', *Mind*, 99: 327–54.
- Routley, R. (1980) *Exploring Meinong's Jungle and Beyond. An Investigation of Noneism and the Theory of Items*, Canberra: Research School of Social Sciences, Australian National University.
- Russell, B. (1900) *A Critical Exposition of the Philosophy of Leibniz*, Cambridge: Cambridge University Press.
- Russell, B. (1903) *The Principles of Mathematics*, Cambridge: Cambridge University Press.
- Russell, B. (1905) 'On Denoting', *Mind*, 14: 479–93.
- Russell, B. (1918) 'The Philosophy of Logical Atomism', in *The Collected Papers of Bertrand Russell*, vol. 8, London: Routledge, 1986, pp. 157–244.
- Russell, B. (1940) *An Inquiry into Meaning and Truth*, London: George Allen and Unwin.
- Russell, B. (1948) *Human Knowledge: Its Scope and Limits*, London: George Allen and Unwin.
- Russell, B. (1951) 'Ludwig Wittgenstein', *Mind*, 60: 297–8.
- Russell, B. (1959) *My Philosophical Development*, London: George Allen and Unwin.
- Russell, B. and Whitehead, A. N. (1910–13) *Principia Mathematica*, Cambridge: Cambridge University Press.
- Sartre, J.-P. (1943) *Being and Nothingness: An Essay on Phenomenological Ontology*, trans. H. E. Barnes, London: Methuen, 1958.
- Schaffer, J. (2000) 'Causation by Disconnection', *Philosophy of Science*, 67: 285–300.
- Schaffer, J. (2004) 'Causes Need Not be Physically Connected to Their Effects: The Case for Negative Causation', in C. Hitchcock (ed.), *Contemporary Debates in Philosophy of Science*, Oxford: Blackwell, pp. 197–216.
- Schaffer, J. (2005) 'Contrastive Causation', *The Philosophical Review*, 114: 327–58.
- Schaffer, J. (2010a) 'The Internal Relatedness of All Things', *Mind*, 119: 341–76.
- Schaffer, J. (2010b) 'Monism: The Priority of the Whole', *Philosophical Review*, 119: 31–76.
- Searle, J. (1969) *Speech Acts*, Cambridge: Cambridge University Press.
- Searle, J. (1981) 'Minds, Brains and Programs', in D. Hofstadter and D. Dennett (eds), *The Mind's I*, London: Penguin, pp. 351–73.
- Shalkowski, S. and Bueno, O. (2000) 'A Plea for a Modal Realist Epistemology', *Acta Analytica*, 15: 175–93.
- Shoemaker, S. (1980) 'Causality and Properties', in *Identity, Cause and Mind*, expanded edition, Oxford: Oxford University Press, 2003: 206–33.
- Sigwart, C. (1895) *Logic*, vol. 1, 2nd ed., trans. H. Dendy, New York: Macmillan.
- Simons, P. (2005) 'Negatives, Numbers, and Necessity: Some Worries about Armstrong's Version of Truthmaking', *Australasian Journal of Philosophy*, 83: 253–61.
- Singh, S. (1997) *Fermat's Last Theorem*, London: Fourth Estate.
- Smiley, T. (1996) 'Rejection', *Analysis*, 56: 1–9.
- Sorensen, R. (2008) *Seeing Dark Things*, New York: Oxford University Press.
- Sparby, T. (2012) *Hegel's Conception of the Determinate Negation*, Leiden: Brill.
- Spinoza, Baruch (1674) *Opera*, C. Gebhardt (ed.), Heidelberg: Carl Winter, 1925.
- Stapleton, J. (1994) *Product Liability*, London: Butterworths.
- Stern, R. (2016) "'Determination is Negation": The Adventures of a Doctrine from Spinoza to Hegel to the British Idealists', *Hegel Bulletin*, 37: 29–52.
- Strawson, P. F. (1950) 'On Referring', *Mind*, 59: 320–44.
- Strawson, P. F. (1952) *Introduction to Logical Theory*, London: Methuen.

- Strawson, P. F. (1974) 'Causation in Perception', in *Freedom and Resentment*, rev. edn, London: Routledge, 2008, pp. 73–93.
- Tarski, A. (1933) 'The Concept of Truth in the Languages of the Deductive Sciences', in J. Corcoran (ed.), *Logic, Semantics, Metamathematics: Papers from 1923 to 1938*, 2nd edn, Indianapolis: Hackett Publishing Company, 1983, pp. 142–78.
- Taylor, R. (1952) 'Negative Things', *Journal of Philosophy*, 49: 433–49.
- Todes, S. and Daniels, C. (1975) 'Beyond the Doubt of a Shadow: A Phenomenological and Linguistic Analysis of Shadows', in D. Ihde and R. M. Zaner (eds), *Selected Studies in Phenomenological and Existential Philosophy*, The Hague: Martinus Nijhoff, pp. 203–16.
- Vaassen, B. (2019) *Causal After All: An Interventionist Model for Dualist Mental Causation*, PhD Thesis, Umeå University.
- Varzi, A. (2007) 'Omissions and Causal Explanations', in F. Castellani and J. Quitterer (eds), *Agency and Causation in the Human Sciences*, Paderborn: Mentis Verlag, pp. 155–67.
- Vetter, B. (2015) *Potentiality: From Dispositions to Modality*, Oxford: Oxford University Press.
- Walton, K. (1978) 'Fearing Fictions', *Journal of Philosophy*, 75: 5–27.
- Walton, K. (1990) *Mimesis as Make-Believe*, Cambridge, MA: Harvard University Press.
- Waterfield, R. (2000) *The First Philosophers: The Presocratics and the Sophists*, Oxford: Oxford University Press.
- Webber, J. (2013) 'Liar!', *Analysis*, 73: 651–9.
- Wertheimer, M. (1912) 'Experimentelle Studien über das Sehen von Bewegung', *Zeitschrift für Psychologie*, 61: 161–265.
- Wittgenstein, L. (1921) *Tractatus Logico-Philosophicus*, trans. D. F. Pears and B. F. McGuinness, London: Routledge, 1961.
- Wittgenstein, L. (1953) *Philosophical Investigations*, Oxford: Basil Blackwell.
- Wood, L. (1933) 'The Paradox of Negative Judgment', *Philosophical Review*, 42: 412–23.
- Woodward, J. (2006) 'Sensitive and Insensitive Causation', *The Philosophical Review*, 115: 1–50.
- Zagzebski, L. (1996) *Virtues of the Mind*, New York: Cambridge University Press.
- Zangwill, N. (2011) 'Negative Properties', *Noûs*, 45: 528–56.

Index

For the benefit of digital users, indexed terms that span two pages (e.g., 52–53) may, on occasion, appear on only one of those pages.

- abnormality 73–4
- aboutness 130, 137–8, 141, 145–6
- absence of seeing 103
- abstract entities 62
- act 12–13
- action 54, 65, 181
- affirmation 162–3, 168–72
- affirmation bias 169
- affirming the negative 33–6
- agnosticism 55–6
- alternating disc experiment 123–5
- Anderson, J. 7–8
- animals 125
- Anjum, R. 53–4, 71, 77, 90, 120
- anti-instantiation 15, 44, 157
- appropriateness 73
- argument from meaning 21, 35–6, 162–3
- Aristotelianism 12–13
- Aristotle 48–9, 193, 206
- Armstrong, D. 2–3, 7–8, 10, 17, 27–9, 35–7, 43–4, 47–8, 61–2, 72, 91–2, 97–9, 145, 147, 156–9, 162–3, 176–7, 200–2
- assertion 169–91, 194–5
 - groundless 179–80
 - iterated 189–91
- asymmetricalism 187–8
- atheism 55–6
- Austin, J. 119–20, 162, 187, 190
- axiom of existence 130–3
- Ayer, A. 176, 187
- Azzouni, J. 139, 142
- Bacon, F. 121
- Baehr, J. 179
- Balashov, Y. 58
- Baldwin, J. 162, 187
- Baldwin, T. 61, 96–7
- bans 53–4
- Barker, S. 2–3, 15, 44, 46, 64–5, 157
- Barrow, J. 57, 60–1
- Beall, J. 204–5
- because 82
- Beebe, H. 65, 69–70, 72, 74, 76, 82–3
- being so (sosein) 135–6
- being vs existence 7–8, 134–5
- belief 114–15, 178
 - false 206
- Bergson, H. 187
- Bernstein, S. 78–80
- Big Bang 66–8
- bindu 57
- Black, R. 57
- blame 76
- Borghini, A. 89
- Bosanquet, B. 186
- Botterill, G. 72, 74
- boundaries 41–2, 45–8, 105, 112–13
- Bradley, F. 162, 187
- Bradley's regress 15
- Braun, D. 129
- Brentano, F. 137
- Buchdahl, G. 155–6
- Bueno, O. 96
- Carroll, L. 5, 186
- Cartesian inner theatre 118, 123, 145
- Casati, R. 40–3, 49–51, 112–13
- category mistake 48, 204
- causal contrastivism 74
- causal explanation 82–3
- causal history 82–3
- causal relevance 81–2
- causal selection problem 72
- causation 8–9, 69, 71
 - by absence 25–6, 32, 64–84, 111–12, 199–200
 - counterfactual dependence theory 71
 - normative conception 83–4
 - of absence 27, 64–5
 - process theory 71
- Cavedon-Taylor, D. 105, 115–16
- certainty 181–3
- Chalmers, D. 119–20
- change 8, 12–13
- Cheyne, C. 154–5, 157–8, 162
- Clarke, R. 64–5
- classes 60

- Coggins, G. 41, 62–3
 combinatorialism 99
 commitment 179–80
 communication chord 65
 complex ideas 145–6
 conceivability 134–6
 conserved quantity 79–80
 context-dependence 76
 contextual cue 113–14
 contextual salience 74–5
 contingent truth 148–9, 173–4
 conversational implicature 179
 counterfactuals 32, 71, 73, 79–80, 83–4, 199
 counterpart 89–90
 Crane, T. 130–8, 143–6
 creation ex nihilo 12–13, 66–9
 Crivelli, P. 22, 45
 cross-categorical relation 148–9
- Daniels, C. 51
 Davies, P. 67–8
 declaratives 173
 deduction 114–18
 degrees of being 7–8, 30–1, 39
 Della Rocca, M. 3–4, 10, 68–9
 Demos, R. 153–4
 denial 47–9, 54–5, 168, 169–96, 203–4
 - complex 191–3
 - iterated 189–91
 - priority of 191–2
 - responsive 184–5
 - stubborn 180
- Dennett, D. 102–3, 118–25
 Descartes, R. 23
 determinables 156–7, 176–7
 determinacy 176–8
 determinates 156–7, 176–7
 determination 30–3
 deviant causal chains 111
 Diogenes Laertius 67–9
 directedness 146
 discontinuity 41
 Divers, J. 62–3, 92–3
 Dodd, J. 163–4
 Dokic, J. 112, 116, 122–3
 Donnellan, K. 130, 133–4, 143–4
 Dowe, P. 64–5, 71–2, 79–80, 82
 Dretske, F. 113
 Dummett, M. 171–5
 Dyke, H. 64
- edges 45
 Efird, D. 63
 Einstein, A. 143
 Eleatic reality test 8, 26, 32, 53, 70
- empiricism 102
 empty set 60–1, 129–30
 empty terms 11, 58–9, 128–30, 133, 135–6, 200–1, 203
 empty world 61–3, 96–7
 English 191
 epistemic norms 179–80
 - entailment 180
- epistemology 204
 equivalence thesis 169–72, 188, 194
 error 130, 143–4, 178
 ersatz realism 87
 escalation problem 72–7, 83
 Euthyphro question 94–7
 evolution 120–1
 exclusion (metaphysical) 153–7, 160
 exclusion (normative) 53–4
 existence dependence 164
 expectation 108–9
 experiential knowledge 105–6, 119
 explanation 78, 80, 82–4, 199
 external negation 48–9
- facts 43–4, 147
 falsehood 11, 161–2, 174–5, 186–8
 falsemakers 161–3, 175
 Farennikova, A. 101, 107–10, 112–13, 120–3
 feeling of surprise 112, 116, 122–3
 Fermat's last theorem 56
 fiction 92–3, 130, 139, 143–4
 fictionalism 91–3, 99, 138–40, 197–9, 203
 fictional objects 140
 Field, H. 15
 force 174
 forgetting 55
 Frege, G. 60–1, 131–2, 162, 169–73, 187, 193
 Frege-Geach problem 192–3
- Gale, R. 15–16, 33–4, 149, 155–6, 189
 Gallop, D. 3, 5–6, 11–12, 68
 gaps 10–11
 Geach, P. 162, 169, 187
 general facts 17, 155
 genesis 8, 12
 Gibb, S. 71
 Gibson, J. 113
 Givón, T. 175, 184
 Gorgias 2–3
 Graham, D. 2
 Green, S. 175
 Grice, P. 111, 177
 Griffith, A. 164–6
 Grossmann, R. 39
 grounding 94, 97–9
 guessing 179

- Haldane, J. 64–5
- hallucination 9
- Hanson, N. 120
- hard Parmenideanism 4
- Hart, H. 73–4
- Hawking, S. 67
- Hawley, K. 10
- Hegel, G. 2–3, 45
- Heidegger, M. 2–3, 5
- Heinemann, F. 135
- Heraclitus 12–13
- Hesiod 67
- Hicks, M. 92
- Hitchcock, C. 120
- Hofmann, F. 164
- holes 10–11, 40–3, 200
- Homer 5
- Hommen, D. 2–3, 17, 33–6, 78–9, 169
- Honore, A. 73–4
- Horgan, T. 10
- Horn, L. 48–9, 149, 162, 187, 193, 204–6
- Horvath, J. 164
- how things stand 163–4
- Humberstone, L. 191–2
- Hume, D. 79, 87–8, 97–8, 102
- Humphrey objection 79, 95–6
- Hunt, I. 64–5

- ignorance 56, 205
- illusion 104, 119, 122–5
- immaterial objects 41, 43, 51–2, 65–6
- imperatives 173
- impossible objects 106
- inattentional blindness 103
- incompatibility 155–6, 160, 170–1
- indexicality 87
- India 57
- indispensability 14–15
 - argument 15, 32
- inegalitarianism 7–8, 30–1
- infants 125
- intentional inexistence 137, 145
- intentionality 137, 145
- intentional object 39–40
- internal negation 48–9
- interrogatives 173
- intra-world modality 88
- Irish principle 23
- irreducibility 14–16

- Jacobs, J. 89
- Jacquette, D. 39
- Jago, M. 2–3, 15, 44, 64–5, 157
- Jenkins, C. 92–3, 139–40
- Joseph, H. 187

- judgement 178
- judgement stroke 169
- judgement withholding 178–80

- Kachi, D. 40–1
- Kalderon, M. 138–9
- Kanizsa triangle 104
- Kant, I. 178
- Kaplan, D. 132–4
- Kripke, S. 79, 95–6, 133–4
- Kroon, F. 133, 139, 144
- Kuhn, T. 198

- Latvian 191
- Lavelle, J. 72, 74
- laws of nature 97–8
- Leibniz, G. 68–9, 162, 187
- Le Verrier, U. 142–5
- Lewis, D. 20, 28–9, 61–3, 65–6, 71, 73, 77, 80, 82–3, 87–90, 92–4, 97–9, 160–1, 197
- Lie, S. 53–4
- lies 174–5
- limits 45
- Lock, S. 72, 74
- Locke, J. 102, 136–7, 145
- Lucretius 8

- Mabbott, J. 155–6
- Maddy, P. 15
- manifestation 90–1
- Marmodoro, A. 90–1
- Marsili, N. 174–5
- Martin, C. 41, 66, 81–2, 120, 159–61
- Martin, J.-R. 112, 116, 122–3
- Martinich, A. 131–3, 143
- mathematical entities 15
- maxim of quantity 177
- McGinn, C. 168
- McGrath, S. 83–4
- McGuinness, B. 117
- McTaggart, J. 23–5
- Meinong, A. 7–8, 39–40, 135–6, 151, 160
- Meinongianism 39, 160
- ‘mentation’ 145
- Menzies, P. 72
- metalanguage 162, 187–8
- metaphysical nihilism 61–3
- mismatch model 122–5
- misperception 103–4
- Mlodinow, I. 67
- modal fictionalism 91–2
- modality 100
- modal logic 86
- modal realism 61, 87–8, 94–7, 197
- Molnar, G. 146, 150–67, 170, 202–5

- monism 10
- Moore, M. 53
- Morgenbesser, S. 190
- motion 12
- Mulligan, K. 147
- multiple drafts model 119
- mutual manifestation 120, 124–5
- mysterianism 168

- naturalism 17, 96–9, 145
- naturalness 28–9
 - natural classes 29
- necessitarianism 7, 11–12
- necessity 100
- 'neerg' 33–4
- negation 162, 168, 205–6
 - double 189–90
 - external 193–4, 205–6
 - internal 193–4, 205–6
- negative belief 113, 115, 201–6
- negative causation 64–6
- negative entities 16, 18, 38
- negative events 52
- negative existential 128–9, 132, 136–8, 141, 155, 163, 175–6
- negative facts 2–3, 43–4
- negative holes 42
- negative judgement 149
- negative knowledge 9
- negative norms 53–4
- negative numbers 59
- negative predicates 15–16, 19, 54–5
- negative properties 2–3, 15–16, 19–37, 106, 196–7, 201–2
- negative statements 188
- negative truth 54–5, 147, 148–68, 202
 - analytic 152
 - a priori 152
 - necessary 152
- Newton, I. 198
- Nietzsche F. 56
- nihilism 56
- Nolan, D. 92–3, 139–40
- nominalism 19–20, 22–3
- non-contradiction, law 183–4
- noneism 138
- normativity 121
- not 140, 185–8, 191–2, 204–5
- 'nothing comes from nothing' 8–9
- null individual 63

- object language 187–8
- Occam test 16, 173
- O'Hear, A. 85
- omissions 52–3, 181

- One through Many 21–2, 91
- overdetermination 51

- paraconsistency 99
- Parmenides 2–13, 66, 68–70, 128, 206–7
- Parsons, J. 150–2, 155, 163, 175–6
- Parsons, T. 39
- particulars 19–21
- parts of difference 22
- Pears, D. 111
- perception 9–10, 101–27
 - causal theory 111–13
 - cognitive theory 113–18
 - cognitively laden 119–20
 - mechanism 120–7, 202
 - naïve realism 120
 - transparency 123
- perishing 8, 12
- Perszyk, K. 39
- phenomenal collapse 107–8
- phenomenology 101–2, 107–8, 114–16, 118–19, 122–7
- Pigden, C. 154–5, 157–8, 162
- Plato 2, 8, 22, 24–6, 30–1, 45, 70, 94–5, 184, 206
- Platonism 140
- plenum 10–11, 46–7
- plurality of worlds 29
- Poirot, H. 183–4
- positional number system 58
- positive statements 188
- possibilia 78–9, 85–100
- possibilities 85, 197–9
 - counter-logical 92–3, 99
 - counter-nomic 92–3, 99
 - legal 96
 - logical 85–6, 96
 - natural 85–6, 96, 98
- possible worlds 87–8, 94–7
 - closest 99
- potency 12–13
- Potrc, M. 10
- powers 8, 12–13, 25–7, 89–91, 97–9
 - clusters of 26
- pretence 92–3, 140–3
- Price, H. 178, 185–6, 189–92
- Price, H. H. 155–6
- Priest, G. 2, 106, 138
- primitivism 29
- Prior, A. 155–6
- privations 48–52, 193–4
- prohibitions 53–4
- properties 19–20, 91
 - categorical 97–8
 - conjunctive 27–8
 - determinable 33

- determinate 33
- disjunctive 27–8
- number of 23–5
- sets of 133–4
- see also* negative properties
- propositions 147, 169, 174
- proprioception 105
- Putnam, H. 15

- quantification 46–7, 138
- quasi-causation 79–80
- Quine, W. 6–7, 11, 15, 28, 46, 135, 138, 162, 170, 187
- Quinton, A. 21, 28–9, 33–4

- Raju, P. 155–6
- real relation 131, 138
- reference 11, 40, 130–3, 136–8
 - descriptive theory 133–4
 - direct 133–4
 - empty 128–47
 - extra-fictive 142–3
 - failed 130, 134, 143–6
 - general 129
 - pretend 141–3
 - proxy 133–7, 141
 - representationalism 135–6
 - singular 129
- relevance 53
- reliability 179
- resemblance 21–3
- Richards, T. 95
- Roberts, T. 105
- Rodriguez-Pereyra, G. 61, 96–7
- Roman numerals 58
- Rosen, G. 91–2
- Routley, R. 39, 135, 138
- rules of rejection 193
- Russell, B. 2–3, 7–8, 17, 39, 43, 60–1, 68–9, 117, 134–5, 151, 155–6, 162, 187

- Sartre, J. P. 2–3, 101
- scepticism 55–6, 205
- Schaffer, J. 2–3, 10–11, 17, 64–5, 69, 71, 74–5
- Searle, J. 130, 145
- seeing that 101–3, 107
- seeing what is not there 104
- sensations 114–15
- sense data 119–20
- Shalkowski, S. 96
- shadows 49–52, 65–6
- Shoemaker S. 26
- Sigwart, C. 162, 187
- silence 105, 113–14
- similarity 89–90

- Simons, P. 147, 163
- smell 105
- Smiley, T. 173, 182, 187, 190, 193
- Smith, B. 147
- Socrates 184
- soft Parmenideanism 3, 160, 196
 - methodological 13–17
 - ontological 3–13, 206–7
- Sorensen, R. 2–3, 49–51, 103, 105, 107, 111
- Sparby, T. 178
- Spinoza, B. 45, 47
- spinning disc argument 10
- spontaneity 9, 68–9
- square circles 92–3, 132–3, 139–41, 145–6
- stability 75–7
- Stapleton, J. 74
- Stern, R. 45
- Stoneham, T. 63
- stops 45
- Strawson, P. 111, 132, 175
- Stroll, A. 131–3, 143
- strong Parmenideanism 13, 204
- subjectivism 120
- subsistence 39–40
- substitutability 129
- subtraction argument 61–3, 96–7
- success term 102, 131, 144
- supervenience 153
- surface 41

- Tarski, A. 187
- taste 105
- Taylor, R. 107, 155–6
- template 109–10
- Thom, R. 105
- time 148–9
- Todes, S. 51
- totality facts 17, 43, 47–8, 118, 157–9, 200
- touch 105
- truth 147–50, 161–6, 174–5
 - correspondence theory 149, 153
 - see also* negative truth
- truthmakers 147–53, 161–4, 175, 202–3
 - truthmaker maximalism 153, 163
 - truthmaker necessitarianism 157–8
- truth-value bearer 64

- Ullian, J. 6–7
- universals 20
- univocity 181–3

- Vaassen, B. 72–7
- validity 116–18
- variation dependence 164–6

Varzi, A. 40–3, 82, 112–13
Vetter, B. 89, 97–8
victory of particularity 43–4
void 66, 69–70, 77
Vulcan 142–5

Walton, K. 139–40
Waterfield, R. 5–6, 8
Webber, J. 174–5
web of belief 6–7, 14
Wertheimer, M. 123–5

Whitehead, A. 60–1
Williams, N. 89
Wittgenstein, L. 117, 136–7, 145, 151, 159,
168, 171
Wood, L. 162, 187
Woodward, J. 75

Zagzebski, L. 179
Zangwill, N. 2–3, 7–8, 30–4
zero 57–9
zero–value determinates 58