THE LIMITS OF EPISTEMOLOGY

MARKUS GABRIEL

The Limits of Epistemology

The Limits of Epistemology

Markus Gabriel

Translated by Alex Englander
and Markus Gabriel

polity

First published in German as *An den Grenzen der Erkenntnistheorie*

Polity Press
65 Bridge Street
Cambridge CB2 1UR, UK

Polity Press
350 Main Street
Malden, MA 02148, USA

The translation of this work was funded by Geisteswissenschaften International – Translation Funding for Humanities and Social Sciences from Germany, a joint initiative of the Fritz Thyssen Foundation, the German Federal Foreign Office, the collecting society VG WORT and the Börsenverein des Deutschen Buchhandels (German Publishers & Booksellers Association).

ISBN-13: 978-1-5095-2566-9
ISBN-13: 978-1-5095-2567-6 (pb)
ISBN-13: 978-1-5095-2570-6 (epub)

A catalogue record for this book is available from the British Library.

Library of Congress Cataloging-in-Publication Data
Names: Gabriel, Markus, 1980- author.
Title: The limits of epistemology / Markus Gabriel.
Other titles: An den Grenzen der Erkenntnistheorie. English
Description: English edition. | Medford, MA : Polity, 2019. | Includes bibliographical references and index.
Identifiers: LCCN 2019005170| ISBN 9781509525669 (hardback) | ISBN 9781509525676 (pbk.)
Subjects: LCSH: Skepticism. | Knowledge, Theory of.
Classification: LCC BD201 .G3313 2019 | DDC 121--dc23 LC record available at https://lccn.loc.gov/2019005170

Typeset in 10 on 11.5pt Palatino by
Servis Filmsetting Ltd, Stockport, Cheshire
Printed and bound in Great Britain by CPI Group (UK) Ltd, Croydon

Contents

Preface to the English Edition

I am particularly pleased that this book, which was published more than a decade ago in German, is now available to an anglophone audience, because I wrote it during my first period as a DAAD (Deutscher Akademische Austauschdienst) postdoctoral researcher at New York University. As Alex Englander began to work on the translation, I decided against updating the book by addressing the (secondary) literature that has been published in the intervening period, as this would drastically change both the course of its arguments and its formulations of central problems. In addition, the work that has appeared over the last decade on issues relevant to the book – on objectivity, the metaphysics of idealism and realism, scepticism, and so on – is so interesting and far-reaching in its own right that it would take at least another book to do it full justice and to adjust some of my arguments to this altered landscape. Hence, I present the book to the anglophone audience in more or less its original form.

Its point of departure is the simple but all too often ignored fact that epistemology claims knowledge. However, its knowledge claims are not quite ordinary, as its object is knowledge as such (its nature and its limits, as the saying goes). This obvious feature of epistemology's knowledge claims tends to be neglected in contemporary epistemology in favour of an investigation into first-order knowledge. Therefore, epistemology is threatened by a blind spot at the level of its own articulation and justification. As long as epistemology remains unaware of the contexts within which it operates as an intelligible enterprise, we are entitled to entertain some degree of doubt concerning its actual success. Accordingly, the book sets out from the question of what it would take to justify the knowability of the kinds of theses advanced by epistemologists. In other words, it is an exercise in transcendental reflection – reflection on the conditions of knowledge-acquisition about knowledge.

Within this framework, it deals with two major issues, which it identifies as hinges of epistemology.

The first issue is *the problem of the world as such*. This problem arises from the erroneous notion that the world is an object of knowledge. In

this context, I endorse a specific version of Kant's insight that 'the world' designates at most a regulative idea. In order to make sense of this claim, I distinguish between first-order theories and 'the metatheory'. The role of the metatheory is to make explicit the decisions that guide first-order research, decisions that do not pick out objects in their target domain. The distinction between theory levels serves as an anti-sceptical tool, since it undermines Cartesian-style scepticism about the so-called external world. As a matter of fact, the Cartesian sceptic is on the right track insofar as she implicitly realises that there is no such thing as an external world, but she misconstrues her insight as a commitment to an epistemic asymmetry between an internal world and a domain of external objects potentially beyond our cognitive grasp. To be sure, she is not explicitly aware of the real problem of the world, as is clear from her attempt to treat it as an issue about access: how can we so much as know anything about the external world given that we first need to represent it internally?

The second issue is *the problem of fallibility*. If 'knowledge' is our term for the good case, there has to be something that can go right or wrong. Otherwise, there would be no conceptual space left for the very idea of a bad case. I suggest that there actually is an element common to both the good and the bad case, namely knowledge claims. Knowledge claims are subject to normative evaluation in light of the question of whether they amount to knowledge. At the same time, I argue that this does not mean that knowledge is, as it were, an aperspectival grasp of facts. I wholeheartedly reject the notion that there could be such a thing as 'the view from nowhere'. To know anything whatsoever presupposes the stability of discursive rationality – i.e. of an operation that relies on given parameters that form its context. It simply does not make sense to reboot knowledge claims and to base them on any kind of foundation that (supposedly) transcends the variability of contexts. Much of the book is dedicated to formulating this delicate point without succumbing to the paradoxical, self-undermining expressions associated with the tradition of Pyrrhonian scepticism.

For this reason, I recommend a methodological use of scepticism. Scepticism is integrated into the theory construction of epistemology, which moves from a Cartesian to a Pyrrhonian mode of reflection. In this way, it abandons a model according to which knowledge is a relation between mind and world, or subject and object for that matter. The problem of knowledge should not be couched in terms of a distinction between someone's mere attempt to latch onto a mind-independent reality (the world) 'out there' and the concept of the success conditions of said attempt. We need to overcome the assumption that there is a largely inanimate, material universe 'out there', which serves as the metaphysical foundation of the epistemic objectivity possessed by our mental states. This does not, of course, amount to a denial of specific facts of the matter, including the trivial acknowledgement that we do not make reality up by somehow mentally constituting or socially constructing it. Such a version

of overcoming the world (rightly or wrongly associated with Rorty and some time-slices of Putnam's work) remains mired in the mind–world problem insofar as it takes the form of a reduction of the world to the mind, of facts to practices, of conditions of truth to conditions of warranted assertability, or what have you.

Like any other work of philosophy, this book has an autobiographical context. It reflects a range of conversations whose role for philosophical progress lies in putting up rational resistance against our pre-rational, impulsive biases. The view outlined in what follows results from my attendance at Crispin Wright's mind-blowing Heidelberg seminars on scepticism (2004) and the realism issue (2006), which spurred my interest in epistemology. With this inspiration, I decided to spend a year as a postdoc at NYU to work with Crispin while Paul Boghossian was writing his *Fear of Knowledge* and Thomas Nagel was working on the project which became *Mind and Cosmos*. During this time, I profited from constant exposure to sophisticated forms of realism in various areas of theoretical philosophy, and these led me to reassess my earlier idealist inclinations. To be sure, '(absolute) idealism', as I understand it here, has nothing specifically to do with the notion that reality is somehow generally related to the mental; rather, it is a syndrome of epistemological, semantic and metaphysical considerations underpinning the assertion that nothing outreaches 'human inquisitiveness', in Crispin's phrase. According to idealism thus conceived, reality is somehow or other essentially tied to its knowability.

This book is my first systematic attempt to escape the one-way street that leads to (absolute) idealism. I am still convinced that it is not sufficient just to kick a stone and insist that reality obviously outruns knowability, in that it resists conceptualisation. Stamping one's feet on the ground of reality is just not a good move in the attempt to articulate what it means to say that the facts outrun any specific evidence available for warranted knowledge claims. Both realism and (absolute) idealism have to struggle with associated versions of scepticism as they begin to recover their motivation in the form of explicit and reflexive theory construction.

The combination of philosophical styles and sensibilities to be found in what follows will be unusual to some readers. This is largely due to the various historical contingencies that separate different groups of philosophers from one another. In my view, the kind of isolation euphemistically called 'specialisation' does not always serve philosophy's best interests when it comes to dealing with some of the 'big questions' of humanity, including the question of the extent to which reality is knowable and of how to conceive the role of the thinking subjects it contains.

An important characteristic of the book is that it deviates from naturalism in epistemology. It does so on the grounds that there is no epistemically successful way of outsourcing the problem of human finitude to any other science. To put it bluntly, naturalised epistemology in the wake of Quine is a theory failure which stands no real, intellectually honest chance

against the most minimal sceptical pressure. Naturalised epistemology is an ostrich-like policy, which explains its 'taste for desert landscapes'.

Unfortunately, the reality of human knowledge acquisition and knowledge is much more multifarious than suggested by the simplified formula 'S knows that p'. In actuality, we only know anything whatsoever in a constellation of contexts whose totality remains forever out of epistemic reach. Downstream, reality contains too much data for it ever to be fully transformed into the kind of information we can handle with a given epistemological model (a first-order case of knowing together with its metatheory). Upstream, the interaction of contexts within which we claim knowledge includes the language we use in coding our thoughts, the socio-economic conditions under which a topic seems worthy of interest to us, the neurobiological preconditions of the human sensory systems, the historical time we live in, our unconscious biases, etc. There is simply no way for us to transcend all of this and to stabilise, once and for all, the ground on which philosophical reasoning unfolds.

Bonn and Paris, November 2018

Introduction

We refer to the objective world *as if* it were essentially independent of the fact *that* we refer to it. This attitude towards the world is described by the concept of objective knowledge, which simply picks out those of our mental states that put us in touch with how things are. By the 'world', we usually understand the totality comprising everything that pre-exists our two most basic epistemic activities: acquiring knowledge of what is the case and securing that knowledge against any possible objections we might encounter. The concept of the world is therefore indispensable to our understanding of whatever it is that we know. For when we know something, such that we can both express what we know in the form of an explicit knowledge claim and defend it in the face of critical questioning, what we usually know is simply how the world is.

When, in our more reflective moments, we try to arrive at a more articulated understanding of this idea, we hit upon a concept of the world as the unified horizon of everything that is the case. On this conception, the world is the object of each successful representation of what is the case. Or, to be more precise: the *states of the world* and precisely not the *world itself* – that is, the world *simply as the world* – are the objects of each successful representation of what is the case. This is the intuitive basis of what Bernard Williams called 'the absolute concept of reality':[1] from the standpoint of our knowledge claims, the world itself assumes the status of the absolute, of that which is independent of and prior to knowledge. Our knowledge claims, by contrast, are subject to the condition that they must either describe how the world is or undergo revisions when confronted with critical objections. Indeed, our reference to the objective world makes us fallible precisely because the world, in whatever way it ultimately exists, is independent of our acts of referring. Take away this independence, and our reference would no longer be fallible. But that just means it would no longer constitute reference to an objective world: to cognise is just to be in a fallible position.

Cashing out these general reflections in the form of a developed theory is a much more problematic enterprise than it may seem at first glance.

For one thing, it is vital to note how we are already operating on two levels. *On the one hand*, we need to understand first and foremost that it is the *world* that we know when we possess empirical knowledge. *On the other hand*, this very assertion already oversteps the boundaries of objective knowledge, of empirical cognition. Indeed, it does so in a twofold sense, because knowing what knowledge is is not itself a piece of empirical knowledge about how the world is, any more than the world simply as the world can ever become an object of empirical knowledge. Otherwise, the putative piece of knowledge whose content is the proposition that the world precedes our knowledge claims would be just as fallible as our knowledge of a determinate state of the world. Yet this seems to be impossible: knowledge of the conditions of knowledge's fallibility operates (at least *prima facie*) at a different theoretical level to the fallible knowledge it thematises. And so, we might think, it must enjoy an exemption from fallibility, else we would be fallible in relation to the very question of whether we are fallible.

It seems, therefore, that the standpoint from which we go about trying to render empirical – and that means fallible – knowledge comprehensible has to differ from the standpoint from which we make empirical knowledge claims. Knowledge of empirical knowledge is not itself empirical. So, since minimal insight into the relation between empirical knowledge and the world cannot itself be subject to inductive verification or falsification, we need to draw a distinction between two theoretical levels: between the level of objective knowledge on the one hand and the meta-level, of knowledge of what objective knowledge is, on the other. However unremarkable this distinction may at first appear, the entirety of the following study is devoted to spelling out its far-reaching implications and to making them fruitful for contemporary epistemology – in particular, for the debate surrounding scepticism. As a theory that investigates truth claims, epistemology itself lays claim to knowledge as soon as it pronounces on how the concept of knowledge should be understood.

However, as an engagement with the problem of scepticism soon teaches us, the knowledge claims of epistemology are far from unproblematic. If the sceptic convinces us that we cannot, after all, know what objective knowledge is, let alone how it is possible, then objective knowledge itself threatens to collapse. For without any epistemic grasp of the very concept of objective knowledge, we can hardly ascertain whether there really *is* any objective knowledge. There is therefore no way around the question of what it really means to know something. The possibility of knowledge is a methodological presupposition of every instance of empirical knowledge ascription, without itself being empirically knowable.

Perhaps the most important methodological insight of modern epistemology since Descartes is that the theoretical stance of epistemology itself is motivated by scepticism. An engagement with scepticism belongs to the very conditions of epistemology: it first makes space for and then sustains the question of what it means to know something. Generally

speaking, philosophical questions of the form 'What is X?' acquire their determinacy through being distinguished from what they are not (in this respect, they are just like anything else).[2] And the question of the nature of knowledge is no different: it is comprehensible only given the possibility that there is something we *do not know*, or perhaps even *cannot* know. Knowledge, in other words, acquires its particular conceptual contours precisely in virtue of how we come to distinguish it from its constant companion, ignorance.

This point is echoed in the fundamental epistemological insight recently introduced into the debate by Jonathan Schaffer's *contrastivism*.[3] According to this position, the content of all cases of (propositional) knowledge is respectively determined by its belonging to a class of propositions that is distinguished from a contrast class. Instead of 'S knows that p', we should in fact always read 'S knows that p rather than q'. The contrast class (q) thereby contains all those propositions whose truth implies the falsity of the propositions of the contrast class (p).[4] So, the proposition that I, Markus Gabriel, can now see my notebook in front of me belongs to the class of perceptual propositions. If it were the case that all perceptual propositions were false, such that no one ever perceived anything (because, say, we are only ever dreaming, or we are brains in a vat, or . . .), then the class of dream propositions would be a contrast class of the perceptual propositions.

Knowledge generally contrasts with ignorance. Indeed, the nature of error (ψεῦδος) or ignorance has been one of the central questions of epistemology ever since its origins. It was raised by Plato, especially in the *Theaetetus* and the *Sophist*, and it can hardly be a coincidence that the question of knowledge arose in the context of pre-Socratic metaphysics, which paradigmatically marked a distinction between being and seeming that was at once ontological and epistemological.[5] The development from pre-Socratic metaphysics to the Sophists, against whom Plato developed his theory of ignorance, his 'pseudology', is an historical expression of the fact that, in relying on the determinacy of knowledge, the epistemologist thereby owes a debt to the possibility of ignorance. *Omnis determinatio est negatio* thus applies just as much to the determinacy of the concept as to the objects of knowledge. Whoever claims to know what knowledge is generates an ontological space of opposition within which knowledge contrasts with ignorance, and the continual possibility of the latter becomes the wandering shadow of the former.

I want to elaborate and deepen this connection by pursuing the following thought: scepticism is a *condition of intelligibility* of epistemology. In its continual confrontation with ignorance, epistemology occupies the specific standpoint of a metatheory. From this metatheoretical standpoint, it investigates the question of what (first-order) knowledge is and, in doing so, claims (second-order) knowledge for itself. We have already begun to distinguish first-order knowledge (empirical cognition) from epistemological knowledge by designating the object of the former as

states of the world. Since epistemology itself, by contrast, refers both to the successful unity of knowledge claims and the world (knowledge) and to their potential difference (ignorance), it is continually beyond all empirical knowledge. For the world *simply as the world* is no more an ordinary *object* than objective knowledge itself. This, as we shall see, is one of the central insights to be won through our confrontation with scepticism.

The entire first part (Part I) of this study is therefore concerned with the *function of scepticism in the dialectical economy of epistemology*. By 'dialectic', I understand a form of reflection on the connection between the motivation of a theory and its execution. The task of this first section will be to work out the foundational methodological function of scepticism in the construction of (modern) epistemology. Chapters 1 to 4 take up the question of how the concept of the world relates to scepticism. In this context, we will need to draw a distinction between three different conceptions of scepticism, which will be crucial for the arguments to follow. Specifically, we must distinguish between *negative dogmatism*, *Cartesian scepticism* and *methodological scepticism*. The concept of negative dogmatism, I will argue, can be won through an engagement with the sceptical foundations of Kant's transcendental idealism. For, insofar as he attempts to draw boundaries to knowledge, he puts forward negative dogmatic claims: claims to know that there are some things we cannot know.

Specifically, we will have to critically examine Kant's conception of the world against the background of his Refutation of Idealism. Kant's drawing of epistemic boundaries, his negative dogmatism, clearly involves a distinction between two theoretical levels: empirical and transcendental. Nevertheless, the way in which he carries through on the sceptical motivation of his theoretical standpoint does not leave empirical knowledge unaffected: in Kant, the world threatens to vanish into cognition, objectivity into subjectivity. In light of this threat, he seeks to defend his transcendental idealism with the help of the Refutation of Idealism and to show that the objectivity of knowledge, far from being undermined by transcendental idealism, is in fact secured by it. In contrast to a subjective idealism *à la* Berkeley, transcendental idealism is compatible with the assumed existence of things or objects extended in space and time. Nevertheless, Kant's own idealism leads him to over-extend the thesis of his Transcendental Aesthetic. As we shall see, he is ultimately unable (within the context of the Refutation of Idealism at any rate) to mark a distinction between a *spatial representation* and the *representation of something spatial*. Drawing this distinction requires the introduction of a criterion of publicity, and therefore of other subjects in space and time; these subjects must be able both to refer to the same spatial items and to communicate the fact that they can do so.

Throughout the book, I shall be discussing the problematic of idealism in terms of a distinction, due originally to Robert Brandom, between a thesis of the *sense dependence* of objectivity on subjectivity and the thesis of a *reference dependence* of objects on subjects. The distinction guides my

arguments right through to the final paragraphs of this study. A concept P is *sense-dependent* on a concept Q, according to Brandom, if and only if we would not understand P if we did not understand Q. Understanding P presupposes understanding Q. By contrast, Brandom explains, a concept P is *reference-dependent* on a concept Q if and only if there would be nothing which fell under P if there were nothing which fell under Q.[6] The concept of 'idealism' can now clearly be understood in at least two different ways: as a thesis of sense dependence on the one hand and as a thesis of reference dependence on the other.[7] Sense dependence idealism merely asserts that we would have no concept of objectivity if we did not distinguish it from our subjectivity. This thesis is a second-order assertion about a condition of our access to the world (and so belongs to the metatheory). Reference dependence idealism, by contrast, asserts that there would be no objects if there were no subjects and is therefore a first-order thesis about what there is – or, rather, about the way in which what there is exists. Kant oscillates between the two assertions. It is certainly true that nobody has operated with a differentiation of theory levels as clearly as Kant does – a fact to which the distinction between the empirical and the transcendental bears ample testimony. Yet his handling of this differentiation is ultimately inconsistent, and he consequently ends up a victim of his negative dogmatism (see chapters 1–2). As we shall see, this reveals itself as the Achilles' heel of the Refutation of Idealism.

Moore's overreaction to Kant's negative dogmatism – his naïve ontology of individuals – forms the next topic of discussion. Moore (quite intentionally) undercuts Kant's reflections and manages to formulate a decisive objection to the Refutation of Idealism, which I shall elaborate in my own discussion of Kant. Yet Moore neglects both Kant's conception of the world and his distinction between theory levels. Moreover, he is unable to defend the fundamental category of his naïve ontology, the so-called physical object, against objections stemming from the conceptual relativity of our access to the world. This clears the path for a remobilisation of Kant's conception of the world without violating the common-sense intuition that all true judgements describe the world as it is in itself – that is, how it is independently of the beings who describe it. Specifically, adopting the Transcendental Dialectic's conception of the world in combination with a rigorous differentiation of theory levels can help dissolve a confusion lying in the concept of the world. In a word: Kant here shows us how we might distinguish between the world as unified horizon and the world as object of our knowledge. It is unfortunate, therefore, that the Transcendental Dialectic's conception of the world plays no role in the Refutation of Idealism, the latter occurring at a systematically unfruitful point in the *Critique*.

At this stage, Cartesian scepticism enters the picture (chapters 5–6). In this context, 'Cartesian scepticism' admittedly stands neither for a sceptical position that Descartes himself advocated nor for the concrete form assumed by his sceptical reflections in the *Meditations*. Nevertheless, as

the first to recognise the outlines of its logical structure and to deploy it in a methodologically controlled manner, Descartes bequeathed his name to this variety of scepticism. Before Cartesian scepticism can be developed as a general paradox attending the concept of knowledge or the concept of a justifying reason (chapter 6), we need to set out certain methodological provisos (chapter 5). To this end, I introduce a distinction between *logical* and *dialectical* analyses of sceptical arguments. Both methods can then be put to use in order to better appreciate the impetus behind Cartesian scepticism. Furthermore, we shall see how Cartesian scepticism can best be analysed as a *paradox* – that is, as a set containing seemingly acceptable (and well-motivated) premises, seemingly acceptable (and well-motivated) inference rules, and an obviously unacceptable conclusion.

Our confrontation with the paradox or paradoxes of Cartesian scepticism, however, cannot be restricted to attempts at solving (or dissolving) it. It won't do, say, just to question or replace one of the paradox's premises in an effort to circumvent it. Merely *logical* repair mechanisms will not serve our ultimate purpose here. Rather, we must always ask ourselves questions of another sort: in which theoretical context does the paradox arise? Under which theory conditions can it be introduced – i.e. motivated? In contemporary epistemology, such considerations trade under the name of a 'theoretical diagnosis'.[8] They concern the function of Cartesian sceptical premises in the context of determinate theories and make up the focus of what I am designating *dialectical analysis*. We will test out both methods, logical and dialectical analysis, in the course of a critical engagement with Crispin Wright's ingenious anti-sceptical strategy of *implosion*. However, it will turn out in chapter 6 that, in those cases where the strategy is successful, implosion not only disables Cartesian scepticism but also threatens the very epistemological standpoint from which the strategy is itself developed. For its own standpoint is motivated by premises that lead to the formulation of Cartesian scepticism. Accordingly, the first part of the study then ends by constructing a general paradox of Cartesian scepticism. Since we need draw on nothing more than certain foundational premises of discursive rationality to generate this paradox, it represents nothing less than an intrinsic threat to our rational cognition quite generally.

However, the threat can be averted. By reaching for a *dialectical* manoeuvre, we can restabilise discursive rationality by grounding it in human practice, which is essentially tied to fallible, finite knowledge acquisition. Part II, on *contextualism* and *finitude*, is occupied with the prospects for this manoeuvre. Specifically, it endeavours to rebuild epistemology on a contextualist discourse theory by drawing upon the attempts by Wittgenstein and Sextus Empiricus to draw boundaries to our knowledge. It transpires that there is a common denominator shared by Pyrrhonism and the strategy pursued in Wittgenstein's late philosophy, even though the latter is ultimately *anti*-sceptical and directed against hyperbolic

doubt. This common denominator is contextualism, which I understand as a lesson concerning the necessary finitude of objective knowledge. I begin Part II with a systematic outline of Pyrrhonian scepticism (chapter 7), without, however, entering into Sextus' late Pyrrhonism in all of its historical details. My sole concern is to provide a systematic reconstruction of the fundamental operation of self-application, the περιτροπή, the ploy of 'self-refutation' or of 'turning the tables'. This operation consists in applying our insight into the finitude of knowledge to itself and thereby, however paradoxically, in drawing boundaries to our very insight into the boundaries of knowledge.

Chapters 8 to 10 then reconstruct Wittgenstein's contextualism as a theory of assertoric content. It will turn out that the possibility of all discursive determinacy – that is, of all assertoric content – presupposes discursive operating conditions, conditions which discourses cannot have in their self-conscious possession *in ipso actu operandi*. To justify this thesis, I focus especially on the problem of rule-following, while at the same time connecting Wittgenstein's position with certain basic concepts of system theory, primarily of a Luhmannian provenance. Wittgenstein and Luhmann are in agreement insofar as they argue for the necessary finitude of all operations of observation – that is, of all operations of determinacy: whatever can be something determinate for a discursive community counts *as* something determinate only under the presupposition of historically variable parameters, which respectively fix what a community can register. The community constitutes a discourse precisely by fixing hinges upon which all individual moves in the discourse turn. In my reading, this is the most important thesis of Wittgenstein's *On Certainty*, and I will interpret it as an insistence on the necessary finitude of discourse. This insistence, it should be noted, can already be found in Sextus, albeit in a different context.

Chapters 9 to 12 discuss Wittgenstein's private language argument in order to provide a more precise evaluation of the finitude of all discursively communicable knowledge. On my interpretation, which takes its cue above all from Crispin Wright, Saul Kripke and Meredith Williams, the argument tries to prove that all discursively communicable knowledge is finite; knowledge depends upon the stability of a discourse, which has to be fixed as the context of knowledge ascription. Nevertheless, the participants in a discourse cannot determine its foundations *from within* the discourse itself – not without initiating a meta-discourse, which brings along its own presuppositions or operating conditions in turn. One result of this line of thought is that it is impossible to determine when and whether someone knows something from an absolute standpoint. Yet Wittgenstein shows us how we might put this partially sceptical thesis to use in order to develop a contextualism that manages to reconstruct scepticism as a harmless lesson about our discursive finitude.

Chapter 11 then places John McDowell's disjunctivism in this argumentative context, discussing it as an anti-sceptical strategy. Disjunctivism

attempts to develop knowledge or cognition under the conditions of a theory of intentionality, understanding the latter as a cognitive relation of mind and world rather than, as in Wittgenstein, as a necessarily social and, in this sense, discursive phenomenon. Following Wittgenstein, I attempt to show that McDowell's disjunctivism faces difficulties as an anti-sceptical strategy, since he fails to integrate the socio-semantic dimension of subjectivity into his approach from the outset. He determines the subject as cognitive intentionality and not as a person in space and time (who is always already socially implicated).

Chapter 12 contains a response to the rule-following problem developed by Kripke in his interpretation of Wittgenstein. It uses some of the conceptual cornerstones of Kripke's suggestion that we might need a 'sceptical solution' in the guise of a community view of rule-following in order to understand how anyone can ever actually follow a rule. I turn this slightly desperate diagnosis into a more positive account of the social nature of thought and theorising. In other words, we need to give up methodological solipsism in epistemology and stop thinking of the problem of knowledge in terms of an attempt of a single knower (typically called 'S') to put herself in touch with objective matters of fact (typically called 'p') by claiming that they obtain. Knowers turn into such individuals only by being singled out from a community of successful knowers. Epistemology should thus return from its self-inflicted Robinson Crusoe-like isolation to the actual context of knowledge acquisition. I call this move 'stage setting'.

Chapter 13 then enquires into the relationship between the concept of representation, which traditionally leads to the familiar sceptical aporia of mental representationalism and Cartesian scepticism. Since Wittgenstein ultimately uses his contextualism to undermine the methodological solipsism of modern epistemology – a position accepted by so many post-Cartesian (as well as ancient) epistemologists – contextualism would seem to offer an escape route from the general paradox of Cartesian scepticism. By uncovering the problem of rule-following and its communitarian solution (which I discuss extensively in chapter 10), the solipsistic I can be shown its way out of the 'fly bottle'.[9] The concept of representation thus disappears – but at a price: our subjectivity and our metatheoretical standpoint, both of which have to be invoked when we talk *about* discourses as such, come under serious threat.

This difficulty becomes especially clear when we take a closer look at Wittgenstein's liberal naturalism (chapter 14). In order to safeguard the possibility of mutual understanding between different discourses (language games), Wittgenstein introduces 'second-nature', which functions as a kind of unifying horizon of all discourses. Humans can communicate with one another because of their shared natural history, and 'very general facts of nature'[10] ensure that there is ultimately nothing alien within the domain of the human. Indeed, Wittgenstein is also explicit that all knowledge comes about only 'by the grace of nature'.[11] This position, however,

is incompatible both with Wittgenstein's own motivating contextualism and, as we shall see, with the vital sceptical lesson of finitude. Specifically, it is a position that fails to enquire into the operating conditions of the very discourse in which the claim that our nature is constituted in such and such a way enjoys the validity it does. In other words, Wittgenstein does not think to apply his contextualism to his own reflections (περιτροπή).

As a response to this shortcoming, I endeavour to push Pyrrhonian scepticism through to its logical endpoint while still conceiving of the discourse *about* finitude as a *finite* discourse. This attempt throws up the obvious problem that the resultant metatheory has to concede its own revisability and contingency. To be sure, this does not mean that it has to be false or self-undermining (in the manner of a performative contradiction). But it does mean that it has to be distinguished from an alternative of which it necessarily remains ignorant; it cannot know its own contrast class. Given the principle that all determinacy in logical space generates binary oppositions, applying our insight into the finitude of objective knowledge to itself means erecting a boundary between *this* knowledge and an ignorance that, however indeterminable, is an indispensable assumption.

Lastly, by examining Brandom's interpretation of Hegel, chapter 15 ventures a final attempt to rescue the concept of the world from the contingency of the metatheory. We will see that Brandom's interpretation does not take adequate account of the absolute idealism adumbrated in the *Science of Logic* and, instead, operates with a concept of the world that Hegel wants to supersede. To bring out this point, it will pay to play off some of the basic moves of Hegel's logic of reflection against Brandom's version of an objective idealism. This will finally bring into view the possibility of using Pyrrhonian scepticism to develop a discourse theory whose task consists 'merely' in examining the dialectical consistency of given theories. Once absolute idealism has been liberated from certain absurd prejudices (a task to which this book can unfortunately make only a small contribution), one can understand such an approach as delivering on a Hegelian thesis. Taking our cue from Hegel helps provide a basis for a systematic project: of setting out a methodology for epistemology that does not so much seek to solve supposedly grave sceptical problems as merely to delineate the dialectical topography of their possible solutions. With such a method, we might then evaluate their respective prospects for success.

I see this model of philosophical methodology, generated via reflection on the finitude of knowledge, as a means of inheriting Hegel's talk of 'absolute knowledge', particularly in his *Phenomenology of Spirit*. I therefore conclude with an attempt to show, at least in outline, that the paradigm of philosophy as the unity of method and object – that is, of self-thinking activity in the medium of pure thought – makes an unexpected return within the contemporary scepticism debate in a way that lends renewed impetus to Hegel's programme of absolute knowledge. In

this connection, it is already worth noting, even at this early stage, that absolute knowledge is not infinite knowledge. Absolute knowing too has to fail on account of its higher-order immediacy, its own finitude – which is why, of course, absolute knowing does not represent the last word even for Hegel.[12]

Given the length and breadth of what is to follow, the above will have to suffice as a promissory overview. Yet, as an aid to the reader's orientation, it is worth adding that Part I contains a largely *destructive* discussion of the concept of knowledge and possible anti-sceptical strategies, while Part II pursues a more *constructive* project, building contextualism on the ruins of Part I, but without raising objections against the truth of Pyrrhonian scepticism and its lesson regarding our necessary finitude. Epistemology confronts us with our discursive finitude, a finitude we cannot transcend. But, as I try to show elsewhere, this does not rule out metaphysical theories of infinity, provided that these theories conceive their task aright.[13]

A final word on my approach: I work on the understanding that it is legitimate to combine strategies from 'analytic' philosophy – primarily within epistemology, though also from philosophy of mind – with traditional questions of metaphysics and the theory of knowledge. In any case, it is hardly self-evident that the still widespread practice of opposing analytic to continental philosophy is worth sustaining. It would be wrong to classify the arguments that follow by trying to ascertain their respective debts to particular schools; indeed, they do not even presuppose that overgeneralising titles such as 'analytic' and 'continental' allow us to draw anything like adequate distinctions between different methods and schools of thought. The excessive professionalisation of the practice of philosophy in our time does not mean that philosophy itself is a professionalised science that divides neatly into different disciplines, each marked out by its given characteristic methods. For the purposes of classifying the business of modern academic institutions, a strict organisation of philosophical discourse may be acceptable, even advantageous, to a certain extent. But it has the potential to become orthogonal to philosophy as a discursive practice of freedom. For this reason, in trying to determine the function of scepticism in epistemology, I have tried to develop my own position by taking up the approaches of highly diverse thinkers, without worrying about whether I was philosophising within the accepted house style of any particular school.

Part I

The Function of Scepticism

1

Negative Dogmatism and Methodological Scepticism

One of the chief lessons of the sciences is that the world diverges substantially from how we tend to view it as naïve observers. Knowledge acquisition and doubt are therefore not merely compatible – they are not even independent phenomena. We need only consider the all too familiar fact that, with every increase in knowledge, we also acquire knowledge of what we *do not* yet know, of our own ignorance. Indeed, strictly speaking, knowledge in the demanding sense of 'scientific' knowledge is unthinkable in the absence of doubts as to whether the world is exactly as it appears to us to be. Departing from our everyday assumptions is a condition of the very possibility of knowledge *as such* coming into view. Just like anything else human thinkers might investigate, knowledge becomes an object of study only once we have begun to wonder whether and to what extent it actually resembles what we take it to be on the basis of our pre-theoretical assumptions.

As soon as we take even the smallest bite from the fruit of the tree of knowledge, we face the challenge of scepticism: the possibility of knowledge implies the possibility of its impossibility. Knowledge, in short, implies doubt, and the task of alleviating this doubt falls to epistemology.

While scepticism is intimately bound up with the very *possibility* of each and every attempt to achieve theoretical distance from the world,[1] it represents a particular problem for epistemology. By distancing us from our everyday knowledge ascriptions, epistemology too introduces the possibility of scepticism. But by the same token, it implies the possibility of its own impossibility. After all, epistemology is an inherently self-referential enterprise: it strives for knowledge of knowledge. This makes it peculiarly liable to paradoxes. For if it should turn out that such fundamental epistemic concepts as 'knowledge', 'cognition', 'justification', and so on, imply the possibility of scepticism, then epistemology itself comes under threat. And since it depends upon the viability of these concepts no less than any other theory, epistemology has no choice but to face up to the threat to its own possibility from within.

Equipping ourselves with an explicit understanding of knowledge

and doubt is evidently an exercise in reflection, an attempt to seek out stable ground in the face of a recognisable threat. Hence epistemology is always a theory of reflection, a theory which has to give an account of the presuppositions governing its own construction. So much follows from its reflexive character, from the fact that it issues in knowledge claims – knowledge claims about knowledge. And since it is a reaction to scepticism, we can understand scepticism in turn as an enterprise that passes from the destruction of individual knowledge claims to the destruction of knowledge claims as such.

It is therefore not merely an historical coincidence that philosophers in the modern period have deployed Cartesian scepticism as a *motivational* theory for epistemology, as an ultimately *anti*-sceptical strategy. This anti-sceptical strategy, first introduced by Descartes himself, amounts to a *methodological* scepticism, the idea being that we can thematise scepticism as a condition of the possibility of reflection. At the same time, this strategy provokes new, second-order sceptical attacks against which epistemology also has to arm itself.

One of the common moves in the defence of knowledge against the inherent possibility of scepticism is the invocation of immediacy, be it nature (Hume), common sense (Reid, Moore), the everyday (Heidegger), the ordinary (Cavell), and so on. This immediacy, however, is thrown into question not only by specifically epistemological scepticism but by the scepticism implied by the sciences too. The latter continually remind us that some part of the world, or the world as a whole, is not *actually* as it appears to be from the (alleged) pre-theoretical standpoint of immediacy.[2] The sciences, that is, exploit the difference between being and appearance no less than philosophy. So if we want to adopt a scientific attitude towards the world, we have to reckon with a potentially considerable difference between how it is and how it appears.

Certain words can convey the impression of picking out a unified phenomenon although, on closer inspection, they in fact do no such thing. Philosophical vocabulary is notoriously guilty in this respect, and positions such as idealism, realism, relativism, and so on, have taken on an often confusing multiplicity of forms throughout the history of philosophy. They can often designate basic options within a certain philosophical subject area, or even fundamental systematic approaches. Indeed, in the eyes of at least some of their representatives, they can sometimes function as descriptions of nothing less than the totality of existence. 'Scepticism' provides another example of a word that promises more unity than it in fact delivers. The history of attempts to provide constructive theoretical solutions to all kinds of philosophical problems runs in parallel to the history of attempts to develop corresponding, destructive counter-programmes aimed at demonstrating their impossibility, and the variety of the latter reflects the variety of the former. The conflict between dogmatism and scepticism, which plays out upon philosophy's 'battlefield of endless controversies',[3] does not even begin

with Plato's arguments with the Sophists; it already looms large in pre-Socratic philosophy.

Historically, the phenomena that have been assembled under the label 'scepticism' have depended primarily upon whichever constructive proposals were on offer at any given time, with the result that scepticism has often been 'parasitical' upon dogmatism.[4] Without any further specification, therefore, 'scepticism' is thus just as unclear a term as 'philosophy' or 'science'. Abstractly formulated, 'scepticism' can be regarded as a destructive system of assertions formulated with the intention of systematically dislodging some given piece of constructive theorising. Accordingly, the sceptic philosophises from a position of opposition, following a negative programme that presupposes the existence of a positive programme to be used as a foil. This is why the ancient master sceptic Sextus Empiricus, whom we shall encounter again and again throughout this study, determined the 'activity' (ἀγωγή) of sceptical philosophising as 'an ability to set out oppositions among things which appear and are thought of in any way at all' (δύναμις ἀντιθετικὴ φαινομένων τε καὶ νοουμένων καθ' οἱονδήποτε τρόπον).[5] The sceptic, therefore, avowedly pursues a primarily practical (and thus no longer merely destructive) aim insofar as she attempts to finally make good on the salvific promise of 'tranquillity' (ἀταραξία), just like the adherents of rival Hellenistic schools. Yet she does so by seeking eudemonia not in contemplation of the eternal, as did Plato and Aristotle, but in the life and customs (νόμοι) of the community. These customs do not admit of a *philosophical* legitimation but stand for 'what has to be accepted, the given',[6] as they would later for Wittgenstein.

Although one could cite many contemporary philosophers, such as Richard Rorty, Robert Fogelin or Michael Williams, who self-consciously theorise in the tradition of ancient scepticism, none of these authors seriously take up its soteriological dimension. Yet a much more marked and important difference between ancient scepticism and the contemporary scepticism debate (especially as conducted within analytic epistemology) stems from a feature of post-Cartesian philosophy that I have already mentioned, namely that scepticism came to adopt a systematic function within epistemology.

Since Descartes, that is, it has become customary to incorporate scepticism into epistemology's motivation, a tactic which led Descartes himself to introduce the idea of a constructive scepticism. By a theory's 'motivation', I understand the set of reflections that result in the theory's execution but cannot themselves be justified through the theory's (yet to be established) theoretical resources. *Motivation* is accordingly an operation that conditions a theory, while *justification* – i.e. the giving of reasons – is already conditioned by a theory. Justification always comes after the fact of motivation.

Descartes made a purely methodological use of scepticism in a way that would prove decisive for modern epistemology. He thereby became (among other things) the precursor of what I shall from now on, following

Dietmar Heidemann, label *integrative anti-scepticism*.[7] By this term, I mean to pick out any anti-sceptical strategy that regards scepticism as the *condition of the intelligibility* of the basic question of epistemology: the question of the nature of knowledge. Integrative anti-scepticism sets out from the assumption that the project of modern epistemology can be made comprehensible (intelligible) in the first place – in other words, can be motivated – only given a confrontation with the problem of Cartesian scepticism.

By *Cartesian scepticism*, I understand the project of formulating sceptical scenarios that have the potential to trigger hyperbolic doubt. The relevant scenarios exercise this potential by showing how the world as a whole could be utterly other than it appears, such that most, or even all, of our beliefs about the world would stand revealed as false.[8] Clearly, this sense of 'Cartesian scepticism' does not designate Descartes' own ego-logical or theological anti-sceptical strategy, which instead reverses the pattern described above and attempts to deploy sceptical scenarios as a foil for its own *constructive* programme.

When exploited by an integrative anti-scepticism, Cartesian scepticism functions as a condition of modern epistemological theorising: the strategy in question integrates scepticism into the project of epistemology in the sense that, by making a case for the impossibility of knowledge as conceived by a given philosophical theory, it is scepticism which first opens up the space for epistemology's basic questions. However, the business of highlighting the precariousness of knowledge serves only as a spur to secure it against the spectre of its impossibility, and thus to overcome scepticism. When it plays this role, the problem of Cartesian scepticism is therefore invoked purely as something to be overcome, specifically by casting it in the form of a *methodological* scepticism. The latter arises through the confrontation with the possible impossibility of knowledge and goes on to clarify how knowledge is possible after all. In this way, the possibility of knowledge is to be rendered intelligible in and through the thematisation of its potential impossibility.[9]

This broad anti-sceptical strategy makes room for reflection on the conditions of epistemological theorising by assuming from the start that Cartesian scepticism is a condition of the intelligibility of epistemology itself. It thereby leads to the insight that epistemology is a second-order theory, a theory tasked with thematising the conditions of possibility of first-order cognition. Hence epistemology claims to be reflexive insight into the structure of knowledge and, as such, to constitute knowledge itself – specifically, second-order knowledge. The content of this second-order knowledge is then first-order knowledge, while the content of first-order knowledge, at least in paradigmatic cases of empirical knowledge, is everything that is the case independently of its being known.

This does not mean that empirical knowledge has to be flanked by an explicit epistemology: one can know all manner of things without knowing, in addition, how it is possible to know anything at all. This

is why epistemology requires a motivation: there are conditions of its introduction, namely, whichever conditions lead first-order knowledge to reflect on itself. In other words, the transition from a first-order to a second-order theory – i.e. the operation of reflection – always has to be motivated. After all, knowledge is primarily intentional and thus oriented to objects that it does not necessarily have to thematise *as* objects of knowledge. When I know that there is a glass in front of me, I do not *ipso facto* already know that I know that there is a glass in front of me. Taking this further step requires a change in theoretical attitude, and thus a certain provocation.

Since reflection on the very structure of knowledge is not a constitutive ingredient of first-order knowledge, distinctions between different levels of knowledge are drawn only *in* and *for* epistemology. In fact, our everyday first-order knowledge presupposes the absence of epistemological reflection, a feature that Myles Burnyeat captures with the term 'isolation'.[10] As the scepticism debate of recent decades has made clear, scepticism properly belongs to the theoretical conditions of *epistemological* reflection. This, to be sure, does not automatically imply that the epistemologist has to positively advocate sceptical theses (be they of a local variety concerning certain forms of knowledge or of a global variety concerning knowledge as such). When an epistemology is self-consciously motivated as an integrative anti-scepticism, it need not necessarily have sceptical consequences.[11] Nevertheless, we cannot exclude a priori that methodological scepticism will not lead to insights into the boundaries of knowledge, boundaries that first-order knowledge occasionally oversteps. If it turns out that some, or even all, first-order cognition oversteps the boundaries erected by an epistemological metatheory, then that metatheory is entitled to conclusions that will appear to be sceptical from the viewpoint of first-order thinking. At the same time, provided these conclusions do not apply to the metatheory, it need not count as sceptical in its own eyes.

In what follows, our first task will be to draw a fundamental distinction between *methodological scepticism* and another form of scepticism: *negative dogmatism*. As an initial, intuitive approximation, we can say that *negative dogmatism* consists in the thesis that we are not entitled to *assert* propositions of a certain class, and that we are thus not entitled to a set of knowledge claims.[12] It is merely negative in that it attempts to give a systematic proof that there are certain things we cannot know. Accordingly, negative dogmatic propositions always have the following form: that we know that we can know nothing about some X.

By contrast, *methodological scepticism* takes a set of knowledge claims that we cannot renounce without renouncing *all* knowledge or knowledge as such and introduces the spectre of its possible impossibility. It does not put just any old optional class of our beliefs into question, beliefs we may or may not happen to possess, but instead attacks the very foundations of our capacity of belief formation. Crucially, methodological scepticism

burrows so deep that we cannot assert it dogmatically without entailing drastic revisions to our epistemic self-understanding.

Whenever an epistemological theory is explicitly constructed in light of an insight into the function of methodological scepticism, I shall, as above, speak of *integrative anti-scepticism*. Integrative anti-scepticism does not attempt to refute methodological scepticism directly. Yet, unlike negative dogmatism, it does not accept local sceptical conclusions either. Moreover, it is important to stress that methodological scepticism does not always have to function as an integrative anti-scepticism: in order to do so, it requires the further methodological insight that, because the project of epistemology requires methodological scepticism as a motivation, it has to be integrated into the theory itself. Given this insight, it is simply impossible to eliminate methodological scepticism without at the same time doing away with the very epistemological reflection that puts it to use.

Unlike methodological scepticism, negative dogmatism develops arguments aimed at obtaining specific theoretical results. It commits itself to conclusions to the effect that one can no longer affirm some set or category of knowledge claims to which we previously took ourselves to have an unproblematic and unqualified entitlement. By contrast, methodological scepticism consists of paradoxes – that is, of arguments constructed out of apparently acceptable premises and rules of inference that together support evidently untenable conclusions. Paradoxes constitute a peculiar class of argument: although they conform to the conditions of rationality, for a variety of reasons we usually cannot accept their conclusions. Instead, we find ourselves compelled to dissolve them. While *arguments* typically serve to convince us of a particular conclusion, *paradoxes* have conclusions of which we cannot be convinced – or of which we ought to be convinced only in utterly extreme circumstances. Zeno's paradoxes, for example, are paradoxes because they lack the power to convince us (or ought to convince us only in the most extreme circumstances) that nothing moves. Ultimately, it is clear enough that at least some things do indeed move, and, therefore, the way to dissolve paradoxes has to be by asking which properties of the concepts or rules of inference they deploy need to be revised.

Negative dogmatism might assert, for example, that we cannot know anything about God, because the conditions of the possibility of our cognition are incompatible with positive theological knowledge. Pursuing this tack, it can call into question all positive theological knowledge claims – a process which might even necessitate concrete institutional reforms. Yet negative dogmatism does not trade in paradoxes, pursuing instead other paths of reflection to convince us of its conclusions.

Methodological scepticism, on the other hand, offers a lesson about the finitude of our knowledge. It shows, for example, that, since we have no reflective access to the set comprising the totality of the conditions of our knowledge, all knowledge continually depends upon conditions being fulfilled that lie beyond our rational control. It achieves this effect by