University of Groningen

Progress in philosophy

Geerdink, Leonard Martinus

Progress in Philosophy

a defense of Philosophical Skepticism

PhD thesis

to obtain the degree of PhD at the
University of Groningen
on the authority of the
Rector Magnificus prof. E. Sterken
and in accordance with
the decision by the College of Deans.

This thesis will be defended in public on

Thursday 4 July 2019 at 16.15 hours

by

Leonard Martinus Geerdink

born on 9 June 1981
in Gouda

Contents

Acknowledgments

As so many PhD students, I too struggled to finish this dissertation. That I managed to finish this project at all is due to the unwavering support, friendship and love of the wonderful people who I wish to thank here.

First, I'd like to thank Catarina Dutilh Novaes for the opportunity to work on her *the Roots of Deduction* research project. Her support has been invaluable, as has been the generous access to her vast network. Our weekly reading group has shaped my philosophical thinking to an immeasurable extent and I am proud of the joint paper that we wrote together. Catarina, I know that I may have strained your patience beyond any reasonable expectations, but I hope that our friendship has survived the ordeal now that it is over.

I'd like to thank Barteld Kooi for being my second supervisor. Barteld, your comments on my manuscript were invaluable and I think back fondly to the informative talks that we had in your office during a hot summer. Your keen eye for detail has saved me from numerous embarrassments.

I'd like to thank Lodi Nauta for his assistance and encouragements when

I had given up. Lodi, you cannot know how much I appreciate the time and effort that you took to get me back on track. Without you, this dissertation would never have been completed.

I'd also like to thank all of my other colleagues at the University of Groningen for providing me with a scientific community that stimulated my thinking. Without all of you, I would have never become sufficiently sensitive to many issues that I now deem to be of the utmost philosophical importance. The list of people that I want to mention here personally would be unreasonably long, so let me limit myself to the four people that I worked with daily during my time in Groningen, and apologize to those who deserve to be mentioned here but aren't.

I'd like to thank Matthew Duncombe, who was the first post-doc in Catarina's project. Matthew, thank you for your friendship. We were both new to the city of Groningen, and academic life can be lonely. I was sad that you left us so soon, but happy that you have found a great position in the UK and that life is turning out as you dreamed instead of as you feared.

I'd like to thank Rohan French, who was the second post-doc in Catarina's project. Rohan, you managed to become one of my best friends in the short time that we had together, and I still miss you and Laura daily. I greatly admire your dedication to academic life even when it comes at great personal sacrifice. It is great to see that you have now carved your own career path through sheer determination.

I'd like to thank Job de Grefte, who was one of my fellow PhD students

under Catarina's supervision. Job, thank you for all the passionate discussions and for letting me sleep at your place when I was teaching, but no longer living, in Groningen.

I'd like to thank Bianca Bosman, who was the other PhD student under Catarina's supervision. Bianca, thank you for your friendship and a place to stay.

Furthermore, I'd like to thank the people at the MCMP who welcomed me when I visited Munich in the winter semester of 2015. Being part of this world-class research community has turned out to be one of the happiest moments of my life and a turning point in my personal development. Thanking everyone would, again, be impossible, so let me limit myself, again, to a few and apologize to the rest.

I'd like to thank Hannes Leitgeb for allowing me to visit the MCMP and for his interest in and insightful comments on the work that I was allowed to present there.

I'd like to thank Karolina Krzyżanowska, Catrin Campbell-Moore, Johannes Stern, Martin Fischer, and Paula Quinon. Thank you all for making me feel like I belonged somewhere at a time when life was falling apart. More importantly, thank you for infecting me with your love for bouldering.

Furthermore, I'd like to thank my "Munich" friends: Piotr Kaźmierczak, Tom Sterkenburg and Marta Sznajder. It is difficult to start a life in a foreign country where everything is unfamiliar. It is good to have some friends of one's own.

I'd also like to thank my "Dutch" friends: Lars Tump, my paranymphs Peter Sperber and Timmy de Goeij, Bob Coenraats, and Harmen Ghijsen. What they say about making life long friends at university has turned out to be true. Peter and Timmy, thank you especially for standing with me on the day of my defense. We started this journey together, let's also end it together. Lars, thank you for showing me a way forward now that my academic journey has come to an end and a new journey has to begin.

I'd like to thank my family, both old and new. I'd like to thank my parents, for their support and encouragements. I hope I make you proud. I'd like to thank my brothers. Nothing is more annoying to a PhD student than being asked again and again during family meetings when one expects to be finished. Luckily, it is in the essence of a brother to be annoying,[1] and your subtle (and not so subtle) reminders to just get it over with have certainly contributed greatly to the final result. I'd also like to thank my new family for supplying me with a home away from home, and for the love and support that they have given me.

Finally, but most importantly, I'd like to thank my wife and daughter. My beloveds, after all is said and done, I wrote this dissertation for the both of you.

---

[1]Or, for those who reject essences as I do, it is, at the very least, a prominent family characteristic.

Chapter 1

Introduction

Many philosophers think of philosophy as an epistemic practice.[1] The aim of philosophy is to produce or acquire knowledge about philosophical matters, such as the nature of knowledge or the nature of the good, or at least justified beliefs about such issues, they would say. As such, philosophy is conceived of as continuous with the sciences by these philosophers. What distinguishes philosophy from other disciplines is its subject-matter and methodology; it has objective knowledge of the world in common as its aim. Indeed, according to the recently published *An Introduction to Metaphilosophy*, "many philosophers look very much as if they are working within something akin to what Thomas Kuhn, with respect to the history of science, called a 'paradigm' " (Overgaard et al., 2013, p. 48).

Both scientists and philosophers, for instance, search for *explanations*.

---

[1] Timothy Williamson, for instance, clearly states that the aim of philosophy is to *know* (Williamson, 2007, p. 5).

Scientists, in general, are not satisfied with merely describing phenomena. They want to understand them. The lepidopterologist, for instance, wants to understand why there are so many kinds of butterflies and why they look like they do. Someone who is satisfied with only catching and categorizing butterflies is a butterfly collector, not a scientist. Similarly, philosophers want to *understand* and *explain* philosophical matters. Those who work in ethics, for instance, don't just want to tabulate and codify the moral laws (if there are any). They also want to know *why* these laws, and not others, are the true moral laws, *why* they hold, and *why* we should be motivated by them. For many, this attitude of demanding explanations is precisely what distinguished philosophy from other intellectual practices in the Ancient Greek world (see, for instance, (Adamson, 2014, p. 6)).

However, there is a serious question within the philosophy of philosophy whether philosophy is making any progress in achieving its epistemic aims.[2] The idea here is that philosophy might be *aiming* at giving explanations, but that it might be ill-suited for actually supplying any. In this sense, philosophy might be a pseudoscience, or, for those who think that it goes too far to place philosophy in the same category as astrology, homeopathy, and intelligent design just because it doesn't make much progress, philosophy

---

[2]The branch of philosophy that studies philosophy itself is also called 'metaphilosophy'. I will use the terms 'philosophy of philosophy' and 'metaphilosophy' interchangeably, as is common in the literature, despite Williamson's protestations that the term 'metaphilosophy' is misleading because the philosophy of philosophy is a proper part of philosophy and not some kind of higher-order discipline above or beyond it (Williamson, 2007, loc 96). Although I don't object to the term 'metaphilosophy', I do agree with Williamson that metaphilosophy is most fruitfully thought of as a proper part of philosophy itself.

might be, in the words of Kuhn's famous critic Imre Lakatos, a *degenerative research program* (Lakatos, 1970).

Influential philosophers who have been discussing this issue in the past decade include Peter van Inwagen (2006), Timothy Williamson (2007), Amie Thomasson (2015), David Chalmers (2015), and John Searle (2016). Furthermore, several collections of papers have recently been published on this issue, such as (D'Oro & Overgaard, 2017) and (Blackford & Broderick, 2017), and the literature is continuously growing.

One might perhaps think that one shouldn't think too much about these methodological matters, and, instead, just do philosophy. By doing good philosophy, the problem of progress will take care of itself. "[P]reoccupation with questions about methods tends to distract us from prosecuting the methods themselves. We run as a rule, worse, not better, if we think a lot about our feet", as Gilbert Ryle once put it provocatively (Ryle, 1953, p. 185). Ryle is right that overthinking hurts one's running performance. But that does not mean that one should not continuously hone one's technique. Bad running technique leads to injury and damages the body in the long run. To echo Williamson, "[p]hilosophizing is not like [running], best done without thinking about it — or rather: the best [runners] surely *do* think about what they are doing". (Williamson, 2007, p. 9).

With respect to philosophical progress, I distinguish three basic positions within the debate. The first is to maintain that philosophy does actually make progress by accumulating substantial philosophical knowledge.

Williamson, for instance, defends the view that philosophy makes such accumulative progress, albeit slowly and by trial and error, as long as philosophical theorizing is properly constrained by other disciplines such as semantics, logic, mathematics, and physics (Williamson, 2007, loc. 4947). Williamson, of course, although optimistic, does not deny that there is a serious question concerning the progress that philosophy is making, and holds that in many parts of philosophy such progress does not occur because of lax methodological standards (Williamson, 2007, loc. 4998). One can also find more pessimistic versions of Williamson's view that hold that philosophy might not yet make such progress, but will once proper methodological principles have finally been found. The idea here is that such a methodological breakthrough is right around the corner.[3]

The second view, which one can, for instance, find in (Chalmers, 2015), is that philosophy does not accumulate knowledge, but that philosophical theorizing might slowly be converging on the truth. That is, our current best philosophical theories might not be completely accurate, but they are closer to the truth then the philosophical theories of, say, Aristotle, Descartes, or Kant. To be fair, Chalmers is quite pessimistic about how little progress philosophy is actually making, calling his own view a glass-half empty view, although importantly he is not a skeptic because, according to his view, philosophy's epistemic position improves over time (Chalmers, 2015, p. 31).

---

[3]We often find philosophers express such optimism when philosophy is taking another *turn*.

The third view that I distinguish is the skeptical view. The skeptical view denies that there is philosophical progress in the sense that philosophical theorizing leads to the accumulation of knowledge or the convergence on truth. One can read (Thomasson, 2015) as a defender of such a view with regards to metaphysical truths. She claims that the best we can do in metaphysics is acquire knowledge about our own conceptual scheme, not about deep metaphysical truths. Philosophical skepticism, however, need not be committed to the idea that philosophical practice is epistemically valueless. Thomasson, for instance, believes that we can improve our conceptual schemes through philosophical practice, and (Searle, 2016) holds that, although philosophizing might not lead to substantial philosophical knowledge, it can lead to substantial scientific knowledge through a process known as philosophical midwifery.[4]

Although there might be some philosophers of science who defend a skeptical position with regards to scientific progress in general, it is usually thought that the sciences are clearly making epistemic progress. Vivid examples abound, such as the recent discovery of the Higgs-boson by AT-LAS and CMS ((ATLAS, 2012) & (CMS, 2012)), the development of the CRISPR/Cas9 technology for gene modification ((Gasiunas et al., 2012) and (Jinek et al, 2012)), and the proof of Fermat's Last Theorem (Wiles, 1995).

---

[4]I deal with the question what distinguishes philosophy from other epistemic practices in Section 3.1. I discuss philosophical midwifery in Section 4.5. Basically, I agree that philosophy sometimes births new scientific knowledge, but that these are limiting cases that happen on the periphery of philosophy, not at its core.

I could easily keep adding to this list.

Note, however, that such vivid examples are much harder to find in philosophy; the best positive example that I can think of is the development of modern logic by Gottlob Frege (Frege, 1879), and Alfred North Whitehead and Bertrand Russell (Whitehead & Russell, 1910), but this is a bit of a special case because it concerns the mathematization of a branch of philosophy, and as such, this breakthrough might not belong completely to philosophy proper. The mathematical development of logic is definitely an instance of progress, but the philosophical program connected to it, in this case *logicism*, has not led to many positive results.[5] More or less the same holds for the development of modal logic that instigated the rebirth of analytic metaphysics in the mid-twentieth century, and, if the general arguments of this dissertation are correct, the same will hold for the applications of Bayesian probability theory that currently revitalizes epistemology. The general problem is, of course, that each of these well-developed formal frameworks make certain controversial philosophical assumptions that can always be questioned and never be demonstrated. For instance, classical first-order logic assumes that the world must be consistent, modal logics make assumptions about the nature of possible worlds, and Bayesian epistemology makes certain assumptions about the nature of ideal rational agents. Each of these assumptions can, and has been, challenged in the philosophical literature, and are there-

---

[5]Again, I deal with the issue of philosophical midwifery, the idea that new sciences are born from philosophy during philosophical breakthroughs, in Section 4.5.

fore the subject of peer-disagreement.[6] This does not, of course, mean that such frameworks cannot be fruitfully employed in the investigation of philosophical issues.[7] The claim here is only that the development of such formal programs do not necessarily lead to substantial philosophical *knowledge*.

Even though positive results might be rare in philosophy, there seem to be, at least, some vivid examples of refutations of philosophical theories. Edmund Gettier's rejection of the justified-true-belief analysis of *knowledge* (Gettier, 1963) and Kripke's rejection of the theory of descriptions (Kripke, 1980) come to mind. Note, however, that up to this date, professional philosophers are still defending these "refuted" philosophical positions and they cannot, therefore, be considered to be definitively refuted. For instance, Laurence BonJour, among others, still defends a justified-true-belief account of knowledge—having strengthened the justification condition such that it entails truth (BonJour, 2010)—and Stephen Neale thinks that accepting the theory of descriptions is crucial for a proper understanding of anaphora such as 'every farmer who owns a donkey beats it' (Neale, 1990). Also note that all of these philosophical breakthroughs, both positive and negative, took place decades ago, which suggests that, even if there is philosophical progress, it is indeed slow. I can think of no philosophical breakthrough in the last decades that is comparable in scope with finding the Higgs-boson or the development of CRISPR/Cas9.

---

[6]I investigate the problem of *peer-disagreement* in detail in 2.2.

[7]See chapter 5 on my views on which epistemic goals are attainable in philosophy.

The reader might grant that philosophy does not make progress in the sense that physics, biology, or mathematics do, but nevertheless insist that there clearly is progress in philosophy. However, it often remains unclear how such philosophical progress should be conceived of and how it could be measured. That it is important to answer this question first is shown by the fact that, at least by some measures used by university administrators all over the world, philosophy is making excellent progress. There are now more professional philosophers than ever, and university administrators are happy to note that the number of publications in peer-reviewed journals is growing at an accelerating rate.[8] However, when we are told this positive news, our doubts concerning philosophical progress are not really addressed. This is not the sense of progress that we are interested in. But what sense of progress we are interested in is hardly ever made fully explicit.

In this dissertation, I investigate one possible explicitation of what this

---

[8]The annual number of people who received a Bachelor's Degree in philosophy in the USA has roughly doubled between 1987 and 2014, according to the APA (`http://c.ymcdn.com/sites/www.apaonline.org/resource/resmgr/data_on_profession/degree_completion_in_philoso.pdf`), and we see similar numbers for Master's and Doctoral Degrees. During the same time, APA membership has increased by 37%, from 58,222 to 79,796 members (`http://www.apa.org/about/apa/archives/membership.aspx`). There does seem to be a bump around 2010, which might be related to the economic crisis, although more research is needed here to be able to say anything definitive. Furthermore, according to SCIMAGOJR (see `http://www.scimagojr.com/countryrank.php`), 1251 philosophy articles were published in peer-reviewed journals in the USA in the year 1996. In 2016, that number has increased by more than 200% to 3815. A brief calculation over the years also shows that the annual growth rate of the number of peer-reviewed articles is increasing; from roughly 15% growth per year at the end of the 90's to roughly 25% now. As above, it is remarkable that there is a bump in growth-rate around 2010.

Although these numbers focus on the USA and might not completely generalize, it is plausible that similar trends hold throughout the Western world.

notion of philosophical progress could consist of, namely, *epistemic* progress, and defend a skeptical position regarding this kind of progress. That is, I claim that philosophical practice does not lead to epistemic progress. With epistemic progress, I mean progress with regards to either accumulating *knowledge* or improving *justification* of philosophical theories.[9]

In Chapter 2, I first discuss three arguments that aim to show that there might indeed be a problem with epistemic progress in philosophy. In Section 2.1, I present a thought experiment first developed by Dietrich (2011) that aims to show that philosophers from the distant past, such as Aristotle, do not have to revise their substantial philosophical beliefs to the same extent as they would have to revise substantial scientific beliefs. Substantive philosophical positions of the past are still, by and large, defensible. Then, in Section 2.2, I investigate the curious phenomenon that philosophers tend to disagree about almost every substantial philosophical proposition. Such widespread peer-disagreement on every substantial philosophical issue can be used to argue for philosophical skepticism. Lastly, in Section 2.3, I argue that the argument from pessimistic meta-induction is much stronger in the particular case of philosophy than for science in general. Due to the non-empirical nature of philosophical issues, the no-miracles argument that is normally used to counter-balance the pessimistic meta-induction is not

---

[9]Other forms of progress that fall outside the scope of this dissertation include *ethical progress* and *political progress*, such as the abolishment of slavery and the establishment of democracy in most of the Western world. It is sometimes claimed that philosophers such as Locke and Mill played a crucial role in these developments.

available in the case of philosophy.

The above arguments suggest that there is indeed something special about philosophy with regards to epistemic progress. However, the aim of investigating these problems concerning *philosophical* progress immediately raises the question of what demarcates philosophy from other epistemic practices. I deal with this demarcation problem in Chapter 3. It turns out that it is seemingly impossible to demarcate philosophical practices from other kinds of epistemic practices by stipulating necessary and sufficient conditions for a practice to be philosophical. As we will see in Section 3.1, we cannot demarcate philosophy by reference to either its subject-matter or its methodology. Philosophy is probably better conceived of as a family resemblance concept. That does not mean, however, that philosophical practices do not have any characteristics. In particular, I identify three important characteristics in Section 3.2 and explicate them: philosophical practices depend on arguments and reasons, philosophers are considered to be epistemically autonomous, and philosophical practices are performed in contexts of high epistemic standards. In Section 3.3, I argue that these characteristics explain why deductive arguments play a crucial role in philosophical methodology. Finally, in Section 3.4, I present an explication of philosophical practice that is inspired by Robert Brandom's *Game of Giving and Asking for Reasons* as a model that will facilitate our further discussion of philosophical practice and the possibility of epistemic progress.

After having introduced the problem of epistemic progress in Chapter

2, and having presented an explication of philosophical practice in Chapter 3, I argue in Chapter 4 that it is indeed impossible to achieve substantial epistemic progress in philosophy. In Section 4.1, I first look at two models of epistemic progress: epistemic progress through the construction of a foundational system and epistemic progress through hypothesis testing. I show that both of these models require a set of determinate inference rules and a set of immediate data. Then, in Section 4.2, I argue that the first of these requirements is undermined in philosophy due to the problem of *logical pluralism.* That is, there is peer-disagreement about what the correct logical principles for philosophical practices are (and even whether there exists such a unique correct logic). Then, in Section 4.3, I undermine the second requirement of immediate data. I argue that there is no set of objective philosophical data due to the fact that philosophical intuitions are unreliable and therefore cannot play the role of foundation or touchstone. Finally, in Section 4.4, I bring this all together to argue for *philosophical skepticism.* I end this chapter with a caveat concerning the possibility of philosophical midwifery (Section 4.5).

So if philosophical knowledge is indeed impossible in most cases, what about justified philosophical beliefs? In Section 5.1, I argue that, although demonstrations and refutations cannot be the epistemic product of a philosophical practice, a reflective equilibrium is a possible epistemic product of an ideal play of the Game of Giving and Asking for Reasons that I presented as a model of philosophical practice. At first sight, such equilibria seem to show that, although philosophical knowledge might be out of reach, justified

philosophical positions are at least possible. However, in 5.1.1, I rely on a distinction made by, e.g. (Littlejohn, 2009), between *doxastic* justification and *personal* justification, and claim that achieving a reflective equilibrium only personally justifies a philosophical position. That is, although a reflective equilibrium cannot establish a philosophical position as highly probable, it does indicate that the epistemic agent that accepts such an equilibrium cannot be faulted. We will also deal with problems concerning the underdetermination of such equilibria, and the role of epistemic values in resolving such problems, in Section 5.2 and Section 5.3 respectively. We will see that it is impossible within philosophy to solve these problems of underdetermination without reference to problematic contextual values. This does not, of course, mean that anything goes, as I argue in Section 5.4. What it does mean, however, is that a completely explicit philosophical position says as much about how the world could be as it does about what the one who accepts it values (Chapter 6).

Chapter 2

The Problem of Progress

In the previous chapter, I introduced the problem of philosophical progress: the question whether philosophy makes epistemic progress. I also indicated that, at least *prima facia*, there seems to be a case against the idea that philosophy makes epistemic progress. In this chapter, I flesh out this case in more detail.

In this chapter, I present three arguments against philosophical progress. First, in Section 2.1, I present a thought experiment aimed at showing that the philosophical positions of our ancestors need not necessarily be strongly revised. Second, in Section 2.2, I discuss the problem of disagreement in full detail. Third, in Section 2.3, I apply the pessimistic meta-induction in the case of philosophy and argue that it is much stronger in the case of philosophy then in the general case of science.

2.1   A Thought Experiment; Aristotle in our Time

To make plausible that something strange is indeed going on in philosophy, I present a thought experiment which was developed by Eric Dietrich in his paper "There is No Progress in Philosophy" (Dietrich, 2011). This thought experiment is meant to show that there is, at least *prima facie*, a real problem to think through here.

Dietrich's original argument is meant to show that, philosophically speaking, we seem to know little more than Aristotle. As I'll argue below, I think this conclusion is too strong. However, I do think that Dietrich's thought experiment illustrates an interesting difference between philosophy and sciences like physics and biology. Although I think that Aristotle might still be able to learn a lot about philosophy from us, the difference is that Aristotle would not need to substantially revise his philosophical viewpoints.

Let me first present the thought experiment. Dietrich asks us to imagine a scenario in which Aristotle suddenly finds himself transported to a modern university. Let's suppose that he also, magically, understands English. Dietrich does not specify the scenario to such detail, but we can easily imagine that Aristotle has been transported by The Doctor and that the translation circuit of the TARDIS takes care of any linguistic problems. I'm sure it would make a great episode of *Dr. Who*.

Adapting Dietrich's thought experiment slightly to illustrate my point

rather than his, Aristotle, whose intellectual curiosity is boundless, would, of course, start to follow classes. Being a physicist and a biologist himself, he starts by enrolling himself into a course on physics and a course on biology. He also plans to attend philosophy courses on metaphysics and ethics.

Suppose that Aristotle goes to his physics and biology lectures first. He would learn so much during these courses. And he would have nothing to say himself, he would only be able to listen to these lectures in awed silence. He would realize that in his own time, he had gotten everything wrong. There is no difference in principle between the heavens and the earth, the same laws govern both, nor is continuous force necessary to keep bodies moving, rather, force is needed to make them stop, and the natural world is not perfectly ordered for once and for all, rather, species eternally develop themselves in a struggle for survival. One can imagine that this would be a total alienating and, perhaps, humbling experience for him.

Feeling dumbfounded, he would go to his philosophy classes expecting to be bedazzled by the progress metaphysics and ethics has made too. But here, things would be different. It is not, pace Dietrich, that I think that Aristotle would not learn anything essentially new during these classes. For instance, Aristotle might pick up some knowledge of modal logic or the theory of rigid designation, and I'm sure Kantian ethics would be completely new to him.[1] But, as Dietrich points out, rather than having nothing to say during these

---

[1]And, as (Overgaard et al., 2013, p. 53) point out, a lot of what happens in philosophy of mind would also surprise Aristotle, who believed that the seat of intelligence, movement, and sensation was located in the heart (*Parts of Animals*, 656a).

classes, Aristotle would find that he is still considered to be an authority. His teachers and fellow students would make use of this valuable opportunity to ask him about the finer details of his views on *essences* and *virtues*. And even if they disbelieved his claim to be Aristotle, or the Doctor had forbidden him to reveal his identity, he would quickly stand out, making distinctions others, both students and teachers, had not properly made, etc. Aristotle would be good at neo-Aristotelian metaphysics and neo-Aristotelian ethics.

It is important to realize that Aristotle's presence in a modern university would not be just of historic interest. We would not just be interested in Aristotle's views on essences and virtues because they are Aristotle's, but because we might reasonably expect to learn something about essences and virtues from Aristotle. That is, not only the historians of philosophy would want to talk to him. So would metaphysicians and virtue ethicists.

This would be different if we transported, say, Newton or Darwin to a modern university. Although their views on gravity and evolution would be of historic interest, no physicist would expect Newton to tell her anything new about gravity, nor would Darwin be considered to be the perfect person to teach first year undergraduates about the theory of evolution.

Here we see a clear difference between philosophy and the (other) sciences. We have made so much progress in sciences like physics and biology, that the views of even the most famous of historical practitioners are completely outdated. In a sense, this holds also for social sciences, like economics. Aristotle's views on these sciences are completely irrelevant for modern day prac-

tice. A sudden revival of neo-Aristotelian physics or biology is unthinkable. But the recent growth of neo-Aristotelian metaphysics and neo-Aristotelian ethics shows that in philosophy, Aristotle still has something to say. And it is not just Aristotle. We could repeat our thought experiment with other great philosophers from all over history such as Plato, Aquinas, Descartes, Hume, Kant and Frege.

But what does the above thought experiment show? Does it indeed show that philosophy has not made much progress, at least not where traditional philosophical questions of interest to Aristotle about metaphysics and ethics are concerned? Whatever else this thought experiment shows, it does point us to an important difference between philosophy and the (other) sciences as far as progress is concerned. Philosophical viewpoints of past philosophers, no matter how distant or exotic, tend to remain live-options for a philosopher who is dedicated enough to defend them. In contrast, it is often the case that all scientists agree that a certain scientific viewpoint of the past is false. No resuscitation possible. There are no physicists that defend a neo-Aristotelian theory of gravity or biologists that defend a neo-Aristotelian theory of spontaneous generation.

This points to a second difference between philosophers and (other) scientists. Scientists tend to agree with one another about which substantial scientific theses are true and which are false. But there seems to be no substantial thesis to which every philosophers agrees. This leads me to the first argument aimed to show that philosophy does not make progress, the

argument from disagreement.

2.2   The Argument from Disagreement

One main difference between philosophy and the (other) sciences is that there is widespread disagreement in philosophy, seemingly about everything.[2] Or, as the authors of the recent *An Introduction to Metaphilosophy* put it: "Disagreement and debate seem endemic in philosophy to an extent that appears pathological from the point of view of the natural sciences" (Overgaard et al., 2013, p. 51). Such widespread disagreement is often a source of worry that philosophy might not be progressing as the (other) sciences are.

The idea that there is wide-spread disagreement in philosophy is, unfortunately, backed-up by empirical data. The recent survey by David Bourget and David Chalmers indeed shows that philosophers disagree about every substantial philosophical thesis the investigators could think of.

It is true that their data shows that some philosophical positions are more widely held than others. The popularity of naturalism, for one, comes to mind. Other popular views among professional philosophers are the existence of *a priori* knowledge and the analytic-synthetic distinction, the acceptance of non-skeptical and scientific realism, non-Humeanism about physical laws, compatibalism concerning free-will, atheism concerning the existence

---

[2]Weber (2011) has argued that it is perhaps the attitude of philosophers themselves that is to blame. While (other) scientists have a background attitude of trying to agree, philosophers tend to try to find disagreement.

of God, cognitivism about moral judgment, classicism concerning logic, and externalism about mental content (Bourget & Chalmers, 2014, p. 492). But even though some views are more popular than others, their data indicates that there is no substantial thesis in philosophy for which there is a (near) universal consensus.[3] The best Bourget and Chalmers could find is that 82% of the questioned philosophers agree that there is an external world (Bourget & Chalmers, 2014, p. 476). Compare this to the 97% of scientists who believe that living things have evolved over time (Pew Research Center, 2009, p. 37). Even something as politically controversial as global warming *due to human activity* is supported by 84% of all scientists (Pew Research Center, 2009, p. 39).[4] I would bet that, among philosophers, there is more consensus that evolution and climate change are real than there is about the existence of an external world. Someone should run this experiment.

Of course, pervasive disagreement in a field does not necessarily show that no progress is being made. As Peter van Inwagen reminds us, "the 'cutting edge' of elementary-particle physics looks a lot like philosophy in point of pervasive and fundamental disagreement among its respected practitioners" (Van Inwagen, 2004, p. 332). For instance, although the Standard Model

---

[3]Interestingly, their data also shows that philosophers hold unreliable views about their own field. For instance, their data shows that philosophers tend to be objectivists concerning art, contextualists concerning knowledge, consequentialists concerning ethics and nominalists concerning abstract objects (Bourget & Chalmers, 2014, p. 489–91). Would the reader have guessed this? The philosophers that Bourget and Chalmers surveyed did not.

[4]Note that these percentages include *all* scientists, not just specialists in biology or climate science.

of Particle Physics is consistent and hugely predictively successful, among other problems, it does not fully explain the best theory of gravity that we have, i.e. general relativity theory, and it might be inconsistent with the emerging Standard Model of Cosmology. Therefore, alternative theories, in particular the many controversial variants of string theory, are currently being developed. However, as the recent discovery of a particle strongly suspected to be the Higgs-Boson shows (see (ATLAS, 2012) & (CMS, 2012)), it would be ludicrous to claim that particle physics is not making any progress although physicists disagree about what a Theory of Everything should look like. What makes the difference in physics, Van Inwagen continues, is that "there is in physics a large body of settled, usable, uncontroversial theory and of measurements known to be accurate, within limits that have been specified. The cutting edge of philosophy, however, is pretty much the whole of it" (Van Inwagen, 2004, p. 332).

Van Inwagen's considerations point to a major difference between philosophy and the (other) sciences. There is no settled body of basic philosophical knowledge that every philosophy student has to learn, nor a paradigmatic set of philosophical problems that every philosophy student needs to know how to solve. In contrast, every student of physics knows classical mechanics and how to solve paradigmatic two dimensional motion problems, and every student of mathematics knows basic geometry and how to determine missing angles in basic geometry problems.

What makes this pervasive disagreement among philosophers about even

the most basic of philosophical thesis so worrisome? Well, this pervasive disagreement can be turned into an argument for philosophical skepticism, that is, into an argument that the widespread disagreement within philosophy shows that philosophy is not making any progress and that there is no store of philosophical knowledge.

The picture that is being painted here is that of philosophers engaging in an everlasting intellectual brawl about every substantial philosophical claim without ever getting anywhere. Perhaps Kant is most famous for painting this picture. He described philosophy, which he still called metaphysics to distinguish it from natural philosophy, as

> so far from reaching unanimity in the assertions of its adherents that it is rather a battlefield, and indeed one that appears to be especially determined for testing one's powers in mock combat; on this battlefield no combatant has ever gained the least bit of ground, nor has any been able to base any lasting possession on his victory. (Kant, 1787, B xv)

Let's flesh out this brief characterization of philosophy by Kant into a full blown argument for philosophical skepticism.

We can think of the philosophical community as forming a group of epistemic peers, i.e. a group of epistemic equals. Usually, this epistemic equality is taken to mean that each of the epistemic peers has equally good evidence concerning relevant theses and that they are equally good at processing that

evidence. A good example of an epistemic peer group is a group of native speakers of a language. All speakers of a language are roughly equally good at determining whether a given sentence is grammatically well-formed or not, despite the fact that there might be some fluctuations in actual performance.[5] That is, every native speaker of a language can more or less judge equally well whether a non-native speaker of that language is producing a grammatical sentence or not.

Although thinking of philosophers as comprising a group of peers is somewhat of an idealization (some philosophers are better than others), I believe it is indeed fair and fruitful to think that philosophers in general are roughly equally well-informed about the topics that they are working on and are roughly equally capable of thinking about these topics. If we have such an epistemic peer-group, there is an important issue concerning what the rational response is when one is faced with peer-disagreement.[6] As we have seen above, philosophy is rife with such disagreement.

For an instance of peer-disagreement, authoritative native English speakers disagree about whether the sentence 'I am the philosopher *that* writes about philosophical progress' is grammatical or not. For instance, we find Patricia O'Connor, a former editor at The New York Times Book Review, claim in her grammar guide *Woe is I* that in this case both 'who' and 'that' are correct (O'Connor, 2010). I'm assuming that the reader disagrees with

---

[5]This ability underlies the method of intuition in linguistics. I will come back to this point in Section 4.4.1.

[6]For an introduction into issues concerning peer-disagreement, see (Matheson, 2015).

her, but in any case, 'who' and 'that' cannot be used interchangeably in such sentences according to the SAT guidelines, a standardized test widely used for college admissions in the United States. Now, when we find that we disagree with O'Connor, what is the rational thing for us to do? Stick to our guns and argue that O'Connor must be wrong? Or suspend judgment and investigate further by, say, looking at what the Merriam-Webster dictionary has to say about it or argue about it with a fellow language enthusiast?[7]

These two attitudes that I described in the case above, sticking to your guns or suspending judgment, are also the two general views defended in the literature about peer-disagreement. Views that promote sticking to your guns are called *steadfast views*. Steadfast views are defended by, among others, Foley (2001), Kelly (2005), and Lackey (2010a & 2010b). Views in which a peer-disagreement leads to at least some revision of one's beliefs are called *conciliatory views*. Christensen (2007), Elga (2007), and Matheson (2015) are prominent defenders of a conciliatory view.

There are different kinds of conciliatory views, depending on how much weight one thinks the beliefs of one's peers should have. Some philosophers such as Christensen (2007) believe that the opinions of each peer should have *equal weight*, other philosophers such as Douven (2010) think one's own beliefs should weigh more, but that one should attach some weight to the opposing beliefs of one's peers and therefore should revise one's credences.

---

[7]Surprisingly, we find that, according to Merriam-Webster, 'that' can indeed refer to a person.

Although this is an important debate, these differences in views are unimportant for the argument for philosophical skepticism that I'm sketching here. For this argument, it is sufficient that a peer-disagreement shows that there is reasonable doubt concerning a proposition.

In the next chapter, I'll argue that philosophical knowledge should be beyond a reasonable doubt, but let us assume for now that this is indeed the case. If we couple this standard for philosophical knowledge with the claim that philosophers form an epistemic peer-group and the view that a peer-disagreement shows that there is reasonable doubt concerning a proposition, then this leads us to conclude that, since philosophers disagree about all substantial philosophical theses, no substantial philosophical thesis is known by any philosopher. That is, there is no philosophical knowledge.

One could, of course, try to resist this conclusion. Other than claiming that the standard for knowledge that the argument needs is too stringent, one could, for instance, defend a steadfast view concerning peer-disagreements instead. In the rest of this section, I want to look at this strategy in more detail.

According to steadfast views, finding oneself in a peer-disagreement does not demand any doxastic revision. Steadfast views appeal to the intuition that one's own perspective is somehow privileged.

Kelly (2005), for instance, argues that what you are justified in believing in a case of peer-disagreement is entirely a matter of the *first-order evidence*. First-order evidence is evidence that bears directly on the proposition and

should be distinguished from *higher-order evidence.* Although the exact nature of higher-order evidence is a topic of active research and not entirely clear at the moment,[8] the idea is that higher-order evidence is evidence about either the evidence itself, or evidence about subjects' capacities and dispositions for responding rationally to that evidence (Kelly, 2014).

Finding out that one disagrees with one's peers about the truth of a thesis is usually considered to be higher-order evidence about that thesis. If higher-order evidence is indeed irrelevant for the justificatory status of a thesis, and the justification of one's beliefs are entirely a matter of one's first-order evidence, then one is indeed not required to re-evaluate one's position when faced with a peer-disagreement.

A second strategy to defend steadfast views is to claim that, rather than revise one's own beliefs, the rational thing to do if one is confronted with a peer-disagreement is to demote those who disagree with you from peer-status. In the words of Foley, if I find myself in disagreement with you then[9]

> the prima facie reason I have to trust your opinion is defeated, and hence I have no reason to move my opinion in the direction of your opinion unless I have special reasons for thinking that you are in an especially good position to assess P. (Foley, 2001, p. 114)

But to me, neither of these strategies look very convincing. Surely, higher-

---

[8]See, among others, Matheson (2009), Fitelson (2012), and Feldman (2014).
[9]We find similar considerations expressed by King (2011) and Lackey (2010a & 2010b).

order evidence must have *some* evidential bearing on your beliefs. A near-universal consensus among experts that a certain proposition is true or false is at least *a* reason, no matter how weak and defeasible, to accept or reject it. Similarly, having a disagreement about a proposition with someone you consider to be a peer is at least a reason, no matter how weak and defeasible, to reconsider the proposition under discussion. The intuition here is that, if you find yourself in a peer-disagreement, you have gained a reason to believe that you might have made a mistake in judging the evidence.

The standard case to pump this intuition is called the *Restaurant Check Case*. It describes the case where we are both calculating the same thing, but arrive at different outcomes. It goes as follows:

> Suppose that five of us go out to dinner. It's time to pay the check, so the question we're interested in is how much we each owe. We can all see the bill total clearly, we all agree to give a 20 percent tip, and we further agree to split the whole cost evenly, not worrying over who asked for imported water, or skipped desert, or drank more of the wine. I do the math in my head and become highly confident that our shares are $ 43 each. Meanwhile, my friend does the math in her head and becomes highly confident that our shares are $ 45 each. (Christensen, 2007, p. 193)

The intuition that is being pumped by this is case is that, after you learn

that your equally capable friend has reached a different outcome, you cannot be sure that you are the one who got it right. Therefore, the rational thing to do here is to suspend one's judgment about what the result is and recalculate the bill. Similarly, when confronted with the fact that a former editor at The New York Times Book Review believes that 'I am the philosopher *that* writes about philosophical progress' is a grammatical sentence, the rational thing to do is to suspend judgment and investigate. It would be arrogant to suppose, without further evidence, that it was your friend or O'Connor who made the mistake, rather than you.

Worsnip (2014) takes these cases to conclusively refute any radical steadfast view on peer-disagreement, but this, of course, goes too fast. Some subtleties apply here. If we start to turn the knobs on this thought experiment and vary some of the variables, we will see that the steadfast response, i.e. sticking to your guns, sometimes seems to be the rational response.

This is the case, for instance, when one's friend has obviously made a *thinko* and has calculated that each owes $430 on a total $180 bill, something that happens to the best of us. In these kinds of cases, it seems that the higher-order evidence of a disagreement is not evidence for believing that I might have made a mistake, but evidence for believing that you are not equally likely as me to have reached the correct total, i.e. I have reasons to believe that, in this case, you are not my peer.

Nevertheless, being steadfast when one finds oneself in a philosophical disagreement is not the right general strategy.

First, although I feel the pull of the first person perspective, i.e. I tend to think I'm right when faced with disagreement on a philosophical topic by a colleague, I do consider this to be a bias and an intellectual vice. Surely every philosopher has learned valuable things when faced with disagreements and criticism from colleagues. And, unless faced with an obvious *thinko*, one should at least take comments from a peer-review report seriously as a sign to look into the matter again. Not every colleague that disagrees with you should be dismissed as just not getting it.

Second, we can show that, in general, ignoring higher-order evidence about propositions is irrational in the sense that it leads to losing betting behaviour. Elkin & Wheeler (forthcoming) have shown that ignoring that one is in a peer-disagreement over a proposition exposes one to risk.

Simplified, Elkin and Wheeler's argument goes as follows. Suppose that we are in the restaurant case. Now, suppose that I am 90% confident that each share is $43 and my friend is 90% confident that each share is $45. Furthermore, suppose that I do not take our disagreement as evidence that my calculation might be incorrect, but my friend does. In contrast, my friend believes that the right answer could be either $43 or $45. This means that while I am willing to bet 9:1 that the correct answer is not $45, my friend is unwilling to bet 9:1 that the correct answer is not $43. By the peer-disagreement hypothesis, my friend and I are equally likely to have calculated the total correctly. That is, it is equally likely that the correct answer is $43 as it is that the correct answer is $45. That means that I will lose my 9:1 bet in

50% of the cases. The expected gain on my bet is $(-\$9*.5)+(\$1*.5) = -\$4$. In other words, I'm expected to lose \$4 on this bet. My friend, on the other hand, is not exposed to this risk because she took our peer-disagreement as evidence that she might have miscalculated and was therefore not willing to take a corresponding 9:1 bet that the correct answer is not \$43.

Third, even if we assume that a steadfast view is correct when dealing with one's epistemic peers, it is obviously not the right response when one finds oneself in a disagreement with an epistemic superior, i.e. someone that you consider to have better evidence than you or more likely to have correctly processed evidence that you both have.[10] Surely, the right response when one is faced with a philosophical disagreement on a matter with an epistemic superior is some kind of conciliatory response. But barring solipsism, which has never been seriously defended in the history of philosophy as far as I'm aware,[11] we can likely find a historical philosopher that you consider to be an epistemic superior that disagrees with you regarding a substantial thesis that you hold. Even if the argument from peer-disagreement does not work, we can amend the argument using epistemic superiors instead. We could call this the argument from superior-disagreement.

The main problem with the argument from peer-disagreement is that its

---

[10]See Frances (2010) for this argument.

[11]One could argue that L.E.J. Brouwer, the founder of intuitionism, was a solipsist, as his biographer Dirk van Dalen does (van Dalen, 2001). Personally, I think it is better to think of Brouwer as a radical Kantian constructionist, although I grant that Brouwer's mystical views from, e.g.(Brouwer, 1905) are hard to interpret. I think that Wittgenstein's mystical remarks in (Wittgenstein, 1922, §5.62) should be taken in a similar manner.

conclusion is hard to accept. If the conciliatory view is correct and peer-disagreement is always a sign of reasonable doubt, then the constant squabbles among philosophers about every substantial philosophical thesis show that there is no philosophical knowledge. Such a wide ranging skepticism is hard to swallow, although some philosophers, such as Francis (2005), Goldberg (2009), and Kornblith (2013), see this skeptical attitude as entirely appropriate.

Whether this disagreement does indeed warrant philosophical skepticism or not, the fact that philosophers disagree about every substantial philosophical thesis does indicate that there is something peculiar about philosophy that differentiates it from the (other) sciences. And philosophers should come to terms with this fact. Here is this sentiment expressed by Christensen:

> There are great bodies of belief in mathematics, in the sciences, and in our everyday conception of the world that are not subject to significant peer-to-peer disagreement. On the other hand, there are areas of morality, religion, politics, and economics, and, unfortunately, philosophy which are rife with disagreement. Why is this? It seems clear that disagreement flourishes when evidence is meager or poorly distributed, or when, due to our emotional or intellectual limitations, we are just not very good at reacting correctly to the evidence. In other words, disagreement flourishes when epistemic conditions are bad. To focus in on my own field, I think that we all should acknowledge that epistemic conditions

are not so great in philosophy. (Christensen, 2007, p. 214)

Let us for now assume that the argument from disagreement by itself is not sufficient to show that philosophy makes little, if any, progress. One could insist that its conclusion is so hard to swallow that there must be something wrong with it. As we have seen, it has sufficient moving parts that there are many avenues to resist it.

However, there is a second powerful argument for philosophical skepticism; the Pessimistic Meta-Induction. This argument comes from the philosophy of science. I discuss this argument in the next section. As we will see, this argument is especially powerful in the case of philosophy.

2.3   The Pessimistic Meta-Induction

In the previous section, I problematized the notion of philosophical progress by presenting the Argument from Peer-Disagreement. The fact that we philosophers disagree about seemingly every substantial philosophical thesis can be used as a premise to argue that we philosophers ought to suspend judgment concerning every philosophical issue (assuming we do indeed form a group of epistemic peers, of course). In this section, I present a second argument for the claim that philosophy has made little, if any, progress. This argument comes from the philosophy of science and is called the Pessimistic Meta-Induction.

Let me first present the Pessimistic Meta-Induction. The argument as

we know it was first presented by Larry Laudan (1981), although we can find early precursors in e.g. (Poincaré, 1905, p. 178). Laudan developed it to attack the No Miracle Argument that was presented by Hilary Putnam in his *Mathematics, Matter and Method* (Putnam, 1975).

Putnam's No Miracle Argument starts from the premise that our current scientific theories, such as General Relativity Theory and Quantum Mechanics, are extremely successful. Success, in this case, means that the theories allow for many accurate predictions of observable phenomena and their manipulation. This fact is hard to deny. We regularly experience evidence of their success in daily life, for instance whenever we use a GPS-device or a mobile phone. But this fact also calls for an explanation. According to Putnam, the hypothesis that these theories are at least approximately true is the only explanation "that does not make the success of science a miracle" (Putnam, 1975, p. 73). As such, the No-Miracles Argument is an instance of an *inference to the best explanation.*

The Pessimistic Meta-Induction aims to show that, pace Putnam, we are not warranted in believing that our most successful theories must be at least approximately true. It does so by using counterexamples of successful theories of the past that have proven to be false. The success of these abandoned theories cannot be explained by their truth, or even their approximate truth, since they are simply false.[12] In his original paper, Laudan presented many

---

[12]It is generally accepted in this debate that the falsity of these past theories follows immediately from the fact that their central theoretic terms, such as *aether* and *phlogiston*, do not refer.

such examples, among which the theory of the electromagnetic aether and phlogiston theory, while noting that he could extend his list *ad nauseam* (Laudan, 1981, p. 33).

We can draw two morals from these historical facts that Laudan has pointed out.[13] First, if there are counterexamples to the principle that successful theories are at least approximately true, then we have no warrant to believe that our current successful theories are (approximately) true. This is the most difficult to answer form of the argument because one convincing counterexample is sufficient to make the inference problematic, and its cogency is still hotly debated today. Second, if so many of our successful theories of the past have turned out to be false, then by induction we might have good reasons to believe that our current successful theories are false as well. This is a more radical but less compelling version of the argument since, at least in the general context of science, there seem to be strategies to counter it as I will discuss below. I will nevertheless focus on the second argument here because I think that the strategies that are used to balance this argument in the general case of science are not applicable in the case of philosophy.

There are ways to resist the Pessimistic Meta-Induction. One could, for instance, claim that the inductive base is not large enough to warrant the induction. Perhaps the final history of science will show a couple of false starts

---

[13]See Saatsi (2005) for a discussion of why it is important to keep these two morals apart.

in the beginning and then the development of an endless stream of unrefuted successful theories.[14] A related strategy is to argue that the induction is illicit because e.g. it commits the base rate fallacy (Magnus & Callender, 2004) or some other statistical error (Lange, 2002). However, such strategies crucially depend on a specific way the Pessimistic Meta-Induction is formalized using probability theory (see Worrall (2009) for detailed criticisms). As such, they do not seem to be generally applicable to the argument, only against certain formulations of it. Furthermore, as Saatsi (2005) argues, the above strategies help at best against the inductive argument that our current theories are false, not against the attack on the connection between success and truth. This also shows that the argument is a bit unfortunately named. Its most powerful version is not an induction but a *reductio*.

A second strategy is to restrict the No Miracles Argument to mature theories only (see, e.g. Psillos (1999, p. 107)). Since the beginning of the twentieth century, we have seen unprecedented scientific progress in the sense of prediction and control of phenomena. The claim here is that no theory of the past is as successful as our current scientific theories are, and as such, the purported counterexamples that the Pessimistic Meta-Induction depends on fail to be genuine counterexamples. The idea here is that we cannot compare our current scientific theories to the theories of the past because our current theories differ in kind. One form of this strategy can be called the Argument from Scientific Progress and was recently independently developed by Ludwig

---

[14]See, for instance, Psillos (1999, p. 105).

Fahrbach (2009) and Seungbae Park (2011). It is an open question for now whether such a defense is successful. For a critique of this argument in the general case of science, see Müller (2015).

A third strategy is to deny that the successful theories of the past are false, and insist that our previous scientific theories are continuous with current science, and as such, their success is explained by the fact that they were partly true (Psillos, 1999). Current theories are just more successful because they are more approximately true.

I am uncertain at the moment whether any of these strategies are sufficient to save the No Miracle Argument and counter the Pessimistic Meta-Induction in the philosophy of science. But however these debates will play out for the general case of science, I claim that philosophy has no access to any of these defense strategies.

We can apply the Pessimistic Meta-Induction to philosophical theories as well. Given how I presented the Pessimistic Meta-Induction above, there are two complications in the case of philosophy. First, although there is an almost universal consensus in science that e.g. phlogiston theory is false, there are not many philosophical theories that are universally rejected. Perhaps solipsism is a good candidate, as are the skeptical hypotheses that we are misled by evil demons and evil neuroscientists into believing that the world is radically different from what it is, but such radical theories were never accepted by many philosophers in the first place. Second, philosophical theories do not seem successful in the sense that they allow for the prediction and control

of phenomena. Since philosophy typically deals with non-empirical topics, there is nothing for it to predict or control.[15] I, therefore, think that, instead of focusing on successful philosophical theories of the past, it is better to look at the rise and fall of philosophical paradigms in the history of philosophy to serve as our counterexamples for the Pessimistic Meta-Induction.

There are many examples of paradigm shifts in philosophy. For instance, the unification of Platonism and Aristotelianism in scholastic philosophy comes to mind, as does the rise of Cartesianism and Kant's Copernican revolution. More recently, we have experienced the Logical Positivist turn, the linguistic turn, the naturalistic turn, the cognitive turn, and the metaphysical turn. Each of these turns came with a rejection of previous paradigms and a promise that this new paradigm will finally allow us to solve certain philosophical problems for good. But none of these turns deliver. We philosophers still bicker about the nature of essences (and even whether there are any essences to begin with), haven't managed to lay the Cartesian Evil Demon to rest, struggle to explain how substantial a priori knowledge is possible (or even whether such knowledge exists), are unsure about whether ethical judgments are just an expression of how we feel, find it controversial whether how we speak is a reliable guide to how the world is, believe that natural-

---

[15]I reject the idea that philosophical theories are successful in so far as they predict and explain philosophical intuitions because I think that philosophical intuitions are not reliable signs of truth. Furthermore, and this is a point that is also made by Williamson (2007), philosophical theories are not about psychological phenomena (unless we're interested in some aspect of the philosophy of cognition, of course). As such, psychological states seem to be irrelevant. I present these arguments in more detail in chapter five.

ism does not allow for the proper investigation of the normative, have many more philosophical interests than just the mind, and fear that many of our metaphysical speculations are just castles in the sand.

The above seems to be exactly the structure we need to perform our Pessimistic Meta-Induction. Our current philosophical paradigm might seem successful now and we take this as a sign that we are making good progress. However, we have many instances of philosophical paradigms in the past that seemed successful when they were at their height but did not lead to the ultimate solution of philosophical problems. Also, past paradigms are often in conflict with current paradigms, and as such, should be rejected as false by the philosopher who thinks that the success of her current paradigm is an indication that her philosophical theories are (approximately) true.

Again, two morals emerge. First, we should be hesitant to believe that the perceived success of a philosophical paradigm is a sign that we have finally solved certain philosophical problems for good. Many of our philosophical ancestors have thought that they had finally solved a problem or had finally found philosophy's real method, but we now debate, criticize, and reject their solutions and tools. Second, the ultimate rejection of every philosophical paradigm in the past (although each still seems to have their modern day supporters) might present us with good reasons to think that our current paradigms will be rejected and replaced in the future as well. That is, we current day philosophers are making just as little philosophical progress as our philosophical ancestors.

Above, I discussed three strategies by which the Pessimistic Meta-Induction is resisted in the philosophy of science. I suspended judgment on how well these strategies do in the case of science in general. They might very well work in the case of empirical sciences such as physics. But none of these strategies work in the case of philosophy.

First, it seems unlikely that philosophy has finally come of the ground after a 2500 year series of false starts, although we do see such claims from philosophers who lead a paradigm shift. For instance, linguistic philosophers, like Ryle (1953), promised us to finally solve our philosophical problems fifty years ago, and some proponents of formal philosophy promise us so now. But boasts like that have come to nothing in the past—here, Kant and Hegel come to mind—and it just doesn't sound plausible in the face of all the disagreement that we see in philosophy even today.

Second, unlike in the natural sciences, it doesn't seem that our current philosophical theories are much more mature than our philosophical theories of the past. In ethics, for instance, Aristotelian Virtue Ethics, Kantian Deontology, and Millian Utilitarianism look very similar in kind. We would not say that Utilitarianism, although much younger, is a more scientifically mature moral theory than Virtue Ethics. These three theories are all more or less equally successful as moral theories, and each has its strengths and weaknesses.

Furthermore, the *Argument from Progress* seems to be biased towards the natural sciences. General Relativity and Quantum Mechanics are indeed

unprecedentedly successful, but theories seem far less mature in other fields, say, psychology or medicine. Theories are still regularly being discarded and replaced in these fields. Furthermore, the methodological foundations of these sciences seem far less secure than the methods of, e.g. physics and chemistry, as medicine and psychology shudder under the current replication crisis (Open Science Collaboration, 2015). Although it does seem fair to say that there is a difference in kind between, on the one hand, evidence-based medicine and experimental psychology, and, on the other, the seemingly blind groping about of early medicine and early psychology, it seems better to say that these sciences are currently in the throes of adolescence rather than fully matured as the natural sciences are. They have been successful in controlling physical and mental diseases, but, in hind-sight, many medical and psychological theories that were generally accepted in the 20th century have also done unimaginable harm. For instance, radical mastectomies deformed countless women in the early 20th century while having served little purpose (Mukherjee, 2011). Either a lesser operation would have done, or the cancer had spread and no amount of surgery would have helped. Similarly, homosexuality was still classified as a mental disease by the American Psychological Association until 1973 and by the World Health Organization until 1990. Conversion therapy, once thought necessary to "cure" people, is now generally considered to be ineffective, risky, and harmful by the scientific psychological community. The *Argument from Progress* might therefore not generalize beyond the natural sciences.

To be fair, medicine and psychology have also done a lot of good. If we take success as a yardstick, philosophy is much less mature than e.g. evidence-based medicine or experimental psychology. First, philosophical theories hardly ever reach the point where they are generally accepted to begin with. Second, philosophical methods seem much less secure than the statistical methods on which medicine and psychology rely. Last, it is hard to find success stories where application of a philosophical theory led to the fruitful prediction or control of phenomena.[16] It therefore seems that the *Argument from Progress* cannot be used in philosophy to counter the Pessimistic Meta-Induction. It is hard to maintain that current philosophical theories are so much more successful than the theories of the past that they differ in kind.

Third, even though we wouldn't say that philosophical theories differ in kind, it would also be hard to argue that our current philosophical paradigms are continuous with past ones and therefore form a progression. To be sure, many philosophical paradigms reach back into the past. For instance, there are modern versions of Aristotelian, Kantian, Fregean, Wittgensteinian, and Carnapian philosophy. But e.g. Kantian ethics is fundamentally incompatible with Utilitarian ethics. And for any philosophical paradigm of the present, we can find at least one successful philosophical paradigm of the past with which it is incompatible. And as we have seen above, one such example is sufficient.

---

[16]I will somewhat qualify this claim in the last chapter, where I deal with philosophical midwifery.

I conclude that philosophers currently have no answer to the Pessimistic Meta-Induction.

2.4 Conclusion

In this chapter, I presented the problem of philosophical progress. The problem is that it seems that philosophy, as a scientific discipline, is hardly making any progress.

I introduced this issue by presenting a thought experiment where Aristotle was transported to a modern-day university. This thought experiment aimed to make prima facie plausible that many philosophers of the distant past are still considered to be the main authorities on substantial philosophical issues even today. This suggests that there is indeed a problem regarding philosophical progress.

I then presented two arguments that aim to demonstrate that philosophical progress is indeed problematic. First, unlike any of the (other) sciences, empirical evidence suggests that academic philosophers do not universally agree on any substantial philosophical thesis. If we philosophers conceive of ourselves as a group of epistemic peers, then such wide-spread peer-disagreement among us can be used as a premise to argue that we should suspend judgment regarding every philosophical thesis. Second, the argument for a pessimistic meta-induction from the philosophy of science is also applicable to philosophy. Almost every philosophical paradigm has been

rejected at some point in time by a majority of philosophers and even now no paradigm is universally accepted. But, unlike in the (other) sciences, in philosophy the pessimistic meta-induction cannot be countered by strengthening the No Miracles Argument. This makes the Pessimistic Meta-Induction much more compelling in the case of philosophy.

Chapter 3

An Explication of Philosophical Practice

In the previous chapter, I presented a thought experiment and two arguments that aimed to show a *prima facie* problem regarding philosophical progress: prominent philosophical viewpoints of the past need not be substantially revised in light of more recent work, and the pervasive disagreement on every substantive issue in philosophy together with a history of aborted efforts do not seem to allow us to be optimistic for the future. But in order to properly evaluate the issues surrounding the problem of philosophical progress, a more thorough analysis of philosophy as an epistemic practice is needed. This analysis is the topic of the rest of this dissertation.

The first issue that needs to be dealt with in performing this analysis is the nature of philosophical practice itself. If philosophy, in contrast with other

epistemic practices such as physics and biology, does not make epistemic progress, then we need to know what distinguishes philosophy from these other practices. That is the main topic of this chapter. Secondly, if we want to understand philosophical practices, it is helpful to create an idealized model that we can study. Such an explication will guide us in the rest of our study, and makes precise what assumptions our analysis depends on exactly.

In Section 3.1, I first argue that it is impossible to define philosophy by stating necessary and sufficient conditions of a practice to be philosophical and that it is therefore better to think of 'philosophy' as a family-resemblance concept. Although this means that philosophy has no essence in the traditional sense, conceiving of 'philosophy' as a family resemblance concept does entail that philosophical practices have certain family characteristics. I discuss typical features of philosophical practices in section 3.2. In particular, I argue that philosophical practices typically depend on arguments and reasons, are typically executed in contexts of high epistemic standards, and that philosophers should, ideally speaking, be epistemically autonomous. These three characteristics explain, or so I argue in section 3.3, why deductive arguments play an essential role in philosophical practice and why philosophical knowledge needs to be established beyond a reasonable doubt. I will use these two facts in the next chapter to argue for philosophical skepticism, the thesis that philosophical knowledge is impossible. But before I argue for philosophical skepticism, I first offer an explication of philosophical practice in terms of Robert Brandom's *Game of Giving and Asking for Reasons* in

section 3.4, so as to render our discussion in chapter 4 and chapter 5 more precise.

3.1 The Impossibility of Defining Philosophy

In their recent *Introduction to Metaphilosophy*, Søren Overgaard, Paul Gilbert, and Stephen Burwood discuss how difficult it is to define philosophy (Overgaard et al., 2013, pp. 17–44). It proves hard to say what exactly distinguishes philosophy from other practices. Is it its subject-matter? Its methodology? What do all philosophical practices have in common? In this section, I argue that the answer to this question is 'nothing'. There is nothing that all philosophical practices have in common. However, that philosophical practices are diverse doesn't mean that we cannot theorize about philosophy. In this section, I argue that philosophy is best seen as a family resemblance concept, and that we can point to certain family characteristics that philosophical practices tend to have in common.[1] In so far as this is the case, we can theorize about paradigmatic philosophical practices, always keeping in mind that these general claims might not necessarily hold for the practices at the fringes.[2]

When I say that philosophy cannot be defined, let me first say that, of

---

[1]It is perhaps better to think of philosophy as a clan of families. This dissertation mainly focuses on the Anglo-Saxon Analytic family of that clan since that is the philosophical tradition that I myself am a part of.

[2]That philosophical practices at the fringes can be significantly different from the ideal type is why philosophy can birth new disciplines through what is called *philosophical midwifery*. I discuss philosophical midwifery in more detail in section 4.5.

course, I do not mean that it is impossible to provide a dictionary defini-
tion of philosophy. Many English dictionaries provide such a definition. The
OED, for instance, gives many, one of which defines 'philosophy' as "a partic-
ular system of ideas or beliefs relating to the general scheme of existence and
the universe". But it also gives another definition, now obsolete, according to
which 'philosophy' means "knowledge of the occult; magic; alchemy". Impor-
tantly, dictionary definitions are not aimed at capturing the essence of the
things defined. Rather, they aim to impart enough linguistic understanding
of a term so that the reader can grasp it and use it in daily life.

In a similar vein, someone could, of course, define what she means when
she uses the word 'philosophy'. Every writer is free to use words as they
please, especially when they make this use explicit. In a sense, that is what I
will be doing in this dissertation as well when I restrict my theoretical discus-
sion of philosophy to what I take to be the paradigmatic cases of philosophical
practice. But such definitions do not aim to capture the nature of the thing
defined either, and as such cannot be correct or incorrect. They can only be
useful or not.

When I say that philosophy cannot be defined, what I deny is that philos-
ophy has a nature that can be captured in a real definition. A *real definition*,
i.e. a definition that aims to capture the nature or essence of a thing, states
necessary and sufficient conditions for something to count as such a thing.
That is, everything that satisfies the sufficient conditions is such a thing,
and every such thing necessarily satisfies all the necessary conditions. Real

definitions play, and have played, an important role in the history of philosophy. Perhaps they are best known for their role in Plato's dialogues, where Socrates is depicted as scouring the streets of Athens in search of real definitions of e.g piety, the good, and knowledge.

Usually, real definitions take the logical form of *if and only if* statements. Here are two standard examples:

1. Someone is a bachelor if and only if he is an unmarried man.

2. A proposition is known by a person if and only if that proposition is true and that person believes it on the basis of sufficient justification.

It has proven hard to find real definitions in philosophy that are unproblematic. The second definition, long thought to be the correct analysis of *knowledge* and first expounded by Plato in the *Theaetetus* (201d–210a), is now generally believed to have been refuted by the kind of counter-examples made famous by Edmund Gettier (1963).

A simple Gettier-style counterexample to the definition of knowledge of (2) looks as follows:

> After arranging to meet with Mark for help with homework, Luke arrives at the appointed time and place. Walking into Mark's office Luke clearly sees Mark at his desk; Luke immediately forms the belief 'Mark is in the room. He can help me with my logic homework'. Luke is justified in his belief; he clearly sees Mark at his desk. In fact, it's not Mark that Luke saw; it was a marvelous

hologram, perfect in every respect, giving the appearance of Mark diligently grading papers at his desk. Nevertheless, Mark is in the room; he is crouched under his desk reading Frege. (Kennard, 2015, p. 38)

In the above case, Luke believes that Mark is in the room on the basis of sufficient justification—he sees him sitting there—and the proposition is also true. This proposition therefore satisfies the sufficient condition given by the second definition above. Nevertheless, many philosophers feel that Luke does not really know that Mark is in the room. He would, for instance, be surprised and learn something new when Mark suddenly crawled out from under the desk.

There have been many attempts in contemporary epistemology that try to deal with these Gettier-cases. For instance, some people have tried to amend the above definition of knowledge with a fourth condition to block these kinds of cases. To this day, no solution to the problems posed by Gettier-cases has been generally accepted, and discussing all the thorny details here goes beyond the scope of this dissertation.[3]

Having made these preliminary remarks, let me argue that it is impossible to give a real definition of philosophy. I discuss two general definitional approaches that intend to capture the nature of philosophy; the *descriptive approach* and the *prescriptive approach.* In both cases, the problem will

---

[3]The interested reader could look at Jenkins Ichikawa & Steup (2012) for an introduction to the analysis of knowledge.

be the same. Any definition will either be too inclusive, and will allow for practices that are obviously not philosophical to fall under the definition, or they will be too exclusive, and will refuse to count certain philosophical practices as philosophical. Let me start with the descriptive approach.

3.1.1 Descriptive Approaches

The descriptive approach tries to answer the question what all actual philosophical practices have in common. That is, the descriptive approach tries to capture the essence of what philosophers actually do. One problem with any descriptive approach reveals itself immediately because it is unclear who should count as a philosopher and who shouldn't, but let us, for now, simply count everyone who self-identifies as a philosopher.[4]

One way to try and capture the essence of what philosophers actually do is by focusing on philosophical methodology. One may, for instance, be tempted to argue that all philosophical practices are defined by their dependency on arm-chair methods. But, first, that would exclude *experimental philosophy* (and remember that we are looking for a descriptively adequate definition at the moment), and, second, 'arm-chair methods' is too inclusive a term.

---

[4]This characterization, of course, will not really do because e.g. Thales, who should intuitively count, probably did not self-identify as a philosopher, while people who espouse unintelligible metaphysical worldviews, and should intuitively speaking not count, often do. For instance, by this characterization, *Time Cube*, a famous rambling manifesto that was posted on the internet, would count as a work of philosophy because it's author, Gene Ray, claims to be "the Greatest Philosopher". The original *Time Cube* website unfortunately no longer exists, but brave readers can access it via the Wayback Machine of the Internet Archive at `https://web.archive.org/web/20160112193916/http://timecube.com/`.

Mathematics, for instance, also depends on the arm-chair method of *proof*.

However, as soon as you pick out a specific arm-chair method over another, say *conceptual analysis*, that would exclude too much. With a bit of violence, logicians might perhaps be thought of as aiming to analyze the concept of *logical consequence*, but it is harder to see what concept meta-ethicists who are interested in the nature of moral judgments are analyzing (aren't they studying judgments rather than concepts?), and feminist philosophers are definitely not just articulating concepts. Feminist philosophers are actively using their analyses of concepts, such as *sex* and *gender*, to further their political aims. Sometimes, philosophers even engage in purposefully redesigning concepts, not just understanding them. One could call such a conceptual redesign project an *ameliorative analysis*, or a conceptual engineering project, and such projects are perhaps best known through the work of Sally Haslanger (see e.g. (Haslanger, 2000) or (Haslanger, 2012)).[5] Furthermore, it would be false to say that scientists are not analyzing the concepts that are central to their sciences, even if one grants the questionable but somewhat popular thesis that these scientists need help from professional philosophers. Surely, physicists are trying to develop a fruitful conception of *matter*.[6]

Although I have discussed the particularity of only two attempts above, it seems to hold quite generally that there is no method, or collection of meth-

---

[5]Feminists are not the only philosophers who are interested in ameliorative analyses. The method of explication that I use in this dissertation is also a form of conceptual engineering, as is e.g. Hannes Leitgeb's recently articulated notion of *Rational Reconstruction*.

[6]Matter, for instance, can be conceptualized in terms of *elementary particles*, in terms of the *Pauli exclusion principle*, or in terms of the *energy-momentum tensor* of a system.

ods, that has been used by each and every philosopher through history (and only them). Actual philosophical practice is simply too diverse, although the reader is welcome to try and identify a philosophical methodology that unites e.g. Aristotle, Erasmus, Hegel, Marx, Deleuze, Gödel, and De Beauvoir.

If we cannot give an adequate descriptive definition that captures all actual philosophical practices by method, what about trying to find a definition based on subject-matter?

The problem with defining philosophy by subject-matter is that, just as it is hard to see which methods all philosophical practices have in common, it is hard to see what subject-matter they all have in common as well. Aren't philosophy of mind, philosophy of language, philosophy of biology, and philosophy of physics obviously about different things?

One attempt is to say that philosophy is about *concepts*, but, first, we have already seen that other sciences deal with concepts as well, and, second, philosophers think about other subject-matters too. Epistemology might perhaps be conceived to be about the concept of *knowledge*,[7], but metaphysics is about reality and metaphilosophy is about philosophy.

Perhaps we could try and say that philosophy is defined by dealing with the most *general* questions. To quote Sellars, perhaps philosophers aim "to understand how things in the broadest possible sense of the term hang together in the broadest possible sense" (Sellars, 1963, p. 35).[8] Another sug-

---

[7]Although see Williamson (2007) for the claim that epistemology is about knowledge itself, not its concept.

[8]Let us ignore Marx's dictum for now, first set out in his *Theses on Feuerbach* (Marx,

gestion might be that philosophers try to answer *normative* questions.

But, of course, philosophers are not just trying to answer the most general questions. 'What is the logical form of donkey anaphora?', however interesting, is not a very general question, and questions from (other) scientific disciplines can be quite general as well. For instance, it is hard to find a more general question than the question 'what is everything in the universe made of?' that some physicists are trying to answer. A similar concern holds in the case where we try to say that all philosophical questions are normative. "Are mental states identical to physical states?" is not a normative question, while "should the EU supply agricultural subsidies?" is an important normative question, albeit not a philosophical one.

As we have seen above, it proves hard to say what all philosophical practices have in common. As such, the problem with trying to capture what all philosophical practices actually have in common is that it invites a *deflationary* response. "It might be hard to see what, if anything, unites the efforts of the logician, the political philosopher, the metaethicist, the epistemologist and the feminist philosopher" (Overgaard et al., 2013, p. 20). The answer seems to be that philosophers have nothing in common but the fact that they tend to work in the same kind of departments at Western universities, and perhaps, given that e.g. Ayn Rand was never associated with any philosophy

---

1888) and famously engraved on his tombstone (and in the entrance hall to the *Humboldt-Universität zu Berlin*), that the point of philosophy is not to understand the world but to change it; a slogan that has been taken up by feminist philosophers such as Nancy Bauer (2015) and applied ethicists such as Peter Singer (2009).

department, not even that.[9] This leads one to side with Quine, who held
that:

> Philosophy is not a unified profession with a great core of shaped
> competence, as medicine is. 'Philosophy' is one of a number of
> blanket terms used by deans and librarians in their necessary
> task of grouping the myriad topics and problems of science and
> scholarship under a manageable number of headings. (Quine,
> 1975, p. 293)

This deflationary response defeats the purpose of trying to find a descriptively
adequate definition of the nature of philosophy that we started out with.

3.1.2   Prescriptive Approaches

What about a prescriptive approach? If we take a prescriptive approach, we
try to answer the question what all philosophers *ought* to be doing rather
than what they *are* doing. Perhaps all philosophers should be restricting

---

[9]The reader might object that Ayn Rand is not a "real" philosopher. Although I agree
that her philosophical work is not worth reading, I wish to point out that, before John
David Lewis lost his battle to cancer, *Objectivism* was studied and taught in the PPE
program at *Duke University*, one of the top universities in the world. There is even an
academic journal dedicated to the study of her work, and *Objectivism* in general, called
The Journal of Ayn Rand Studies. Thirty issues of this journal, which is currently being
published by *The Pennsylvania State University Press*, have appeared since 1999, and
the journal is indexed in *The Philosopher's Index.* Furthermore, an edited volume called
Ayn Rand's Normative Ethics: The Virtuous Egoist (Smith, 2007) was published with
Cambridge University Press, a prestigious academic publisher. I take all of this to mean
that at least some professional philosophers take her work seriously.

themselves to *conceptual analysis* as a method, even if they don't, or perhaps they should only be dealing with general or normative question.

When discussing descriptive approaches to defining philosophy, we saw that the danger was that a definition was either too inclusive or too exclusive of what actual philosophers have been doing. Being too inclusive is a problem for both approaches, but while being too exclusive is a fatal flaw for a descriptive approach, prescriptive approaches may simply bite the bullet and insist that, indeed, many people who claim to be doing philosophy are actually doing something else.

We can find many examples of this attitude in the history of philosophy. In particular, Early Analytic Philosophers like Wittgenstein come to mind, who claimed that all philosophy was critique of language (Wittgenstein, 1922, 4.0031), as do Logical Positivists like Carnap, who claimed that philosophical knowledge should have empirical content and metaphysics, like art and poetry, is nothing but a way to express one's attitude towards life (Carnap, 1931, p. 238).

The problem with such a prescriptive approach is, of course, that if such a definition is to reject some purported philosophical research projects as non-philosophical, then this rejection needs to be justified. Why should philosophers only study certain subjects or only use certain methods? This should not be a matter of taste. Just because you and your friends don't like a certain subject or method doesn't mean that it isn't worth studying or using.

The reason that is often given is that only by doing philosophy in a cer-

tain way is philosophical progress possible. E.g., philosophical problems can only be solved by better understanding the language we use. The idea here is analogous to why certain subjects and methods are deemed unscientific in the (other) scientific disciplines. For instance, in the medical sciences we reject research programs that do not use randomized trials as bad, because we know that they do not lead to reliable results, and physicists reject research programs that investigate questions like "what happened before the Big Bang?" because such questions don't make physical sense.

Although this looks similar, the difference with philosophy is that we thoroughly understand why non-randomized trials lead to unreliable results and why our physical theories cannot make sense of questions that ask what happened before the Big Bang, and there is broad consensus about these issues. In contrast, there is no consensus in the philosophical community over which methods are reliable and which are not. For instance, we hardly understand what constitutes the method of analysis to begin with, let alone have insight into its reliability. This is somewhat changing with the advent of metaphilosophy and experimental philosophy, but, as I show in the next chapters, what little we do know about the reliability of philosophical methods, most notably about the method of intuitions, seems to suggest that philosophical methods are not particularly reliable in general. Historically speaking, philosophers have also been bad at predicting which philosophical subjects can fruitfully be studied and which cannot. The critiques leveled by philosophical behaviorists, for instance, did not foresee the power of cognitive

modeling that gave rise to the Philosophy of Cognitive Science, and metaphysics flourished again against all expectations with the advent of modal logic as a strong contributing factor.

This then is the main problem with any prescriptive definition that is suggested. Unless such a definition comes with good theoretical reasons that show why certain practices that self-identify as philosophical need to be excluded, reasons that need to be acceptable to the whole of the philosophical community in time, no good prescriptive definition of philosophy is going to be forthcoming.[10]

We have seen above that both descriptive and prescriptive attempts to define philosophy fail. Actual philosophical practices are too heterogeneous to be captured by a non-deflationary definition, and prescriptive definitions require a justification that does not seem to be forthcoming.

Is there nothing then that we can say about philosophical practices as a whole? Of course there is. In the next section I argue that, while we cannot give necessary and sufficient conditions that fully determine the nature of philosophy, we can characterize philosophical practices based on characteristics that these practices tend to have in common.

---

[10]If a potential prescriptive definition aims to capture all practices that self-identify as philosophical we get the same problem we had with providing a descriptive definition. These practices are so heterogeneous that it looks impossible to capture any interesting feature that they all have in common.

3.2 Philosophy: Some Family Characteristics

We need not be able to define something in order to say something characteristic about it.

The *Habsburg Jaw*, for instance, famously characterizes members of the House of Habsburg, although not all members of the Habsburg family suffered from *mandibular prognathism*, the medical term for their protruding lower jaw. A notable exception, for instance, was Maria Theresa, the only woman to have ruled the Habsburg dominions. Also, not every living thing that suffered from mandibular prognathism is a member of the House of Habsburg. Many Shih Tzus, a dog breed thought to have originated in Tibet, suffer from this genetic defect while clearly not being members of the Habsburg Family. And this same unfortunate fact also holds for many people unrelated to the Habsburg family.[11]

Although the Habsburg Jaw does not define members of the House of Habsburg, we can still use this characteristic jaw to theorize about the House as a whole. We can, for instance, infer from the protruding jaw of many of its family members that the family as a whole suffered from genetic defects. And this general fact that we have inferred, for instance, can serve to explain particular facts, such as why Theresa's children did sprout the characteristic

---

[11]Although all humans evolved from a common ancestor, I don't expect the reader to seriously insist that, technically, all people are relatives of the Habsburg family. Although there are problems concerning the identity-criteria of families, and as such it can be vague whether one belongs to a family or not (wars have been fought over these issues), there is a clear-cut sense in which I am not a member of the Dutch royal house even though we do share a distant common ancestor.

family jaw even though she did not. Furthermore, the protruding jaw can serve as a phenomenon we wish to explain. For instance, the characteristic appearance of the Habsburg family can be explained by their tendency to intermarry, a practice that led to significant inbreeding. Lastly, we can also draw morals from their protruded chin. For example, Charles II, the last Habsburg ruler of Spain, suffered so badly from the characteristic chin that he was barely able to chew, frequently drooled, and had such a large tongue that he could barely be understood when he spoke. The moral being here, of course, to be careful with royal intermarriage should you care about the royal appearance.[12]

Just as we can theorize about the characteristics of the House of Habsburg in fruitful ways, such as explaining particular features and drawing morals, we can also theorize about the family of practices that we call philosophy. There might not be an essence that all philosophical practices have in common, but there certainly seem to be characteristics that philosophical practices have in common. That is, it makes sense to think of philosophy as being, what Wittgenstein called, a family resemblance concept (Wittgenstein, 1953, §67).

Although I disagree with James Tartaglia (2011) that tracing back the ancestry of the myriad practices we now consider to be philosophical shows that philosophy has a core subject-matter —gaining an understanding of the world and our place in it— I do think that he presents an illuminating image

---

[12]Charles II's mental defects and sterility might serve as other lessons to be careful with intermarriage, should you care about ruling competence or preserving the lineage.

of philosophy as a family that is held together by a common history.[13] The members of this philosophical family might not share a common essence, but they do tend to exhibit certain characteristics that they have in common. Some practices, say mainstream analytic philosophy, are prototypical members of the family, exhibiting many of the characteristics prominently, other practices, say experimental philosophy, are more distant cousins who some of us hardly recognize as part of the family and who might exhibit only a few of the family characteristic. Still others, such as physics and cognitive science, have perhaps grown so distant that it hardly makes sense to claim that they are still part of the family at all, although there is a common ancestry.

Just as we could use typical family characteristics to theorize about the House of Habsburg as a whole, we can use typical characteristics of philosophical practices to theorize about philosophy as a whole. In so far as it is correct to think of philosophy as a family resemblance concept, we can theorize about philosophical practice as an ideal type and make general, though defeasible, claims about it.

Even though, as I argued above, there is no subject-matter or method-

---

[13]Tartaglia argues that, even though modern philosophers do not tend to primarily study questions that aim to enhance our understanding of the world and our place in it, that is what motivated our philosophical ancestors and therefore is the core of philosophy. I reject this kind of argument as flawed because I also do not believe that measuring land is the subject-matter of geometry, even though the study of geometry was probably originally motivated by the desire to make precise measurements to reestablish boundaries after the yearly flooding of the Nile. The core subjects that a science studies can evolve over time as the science progresses. As such, geometry evolved from the study of land to the study of space, and can now perhaps be seen as the study of all possible spaces, though a topologist would disagree.

ology that defines philosophy, it is still the case that philosophical practices tend to study certain kinds of subject-matters and use certain kinds of methods. For example, philosophical problems tend to be non-empirical, i.e. they resist being fully solved by empirical methods, and philosophical practices tend to share certain methodological characteristics; philosophers tend to give reasons and arguments for their positions, as they tend to want their views to be accepted based on those reasons, not on their authority, and they tend to want to establish their philosophical views beyond a reasonable doubt. In the next chapter we will see that the problem of philosophical progress is related to these characteristics. But first, let me discuss these characteristics in detail.

3.2.1   Dependency on Arguments and Reasons

The first characteristic of philosophical practices that I want to discuss is philosophy's dependency on arguments and reasons. Philosophers typically do not simply state what they have come to accept. They argue for their beliefs. Similarly, philosophers typically do not simply reject a view. They give reasons why they think a view is untenable. Although there is little consensus in philosophy, a recurring theme of this dissertation, professional philosophers do tend to agree that merely expressing personal opinions on controversial topics is not the same as doing philosophy.

One could even argue that this insistence on reasons and arguments is

what originally distinguished philosophy from other intellectual practices in Ancient Greece. According to Peter Adamson, author of the magnificent podcast and spin-off book *History of Philosophy without any gaps* (Adamson, 2014), this is the crucial difference between Presocratics like Thales, who is generally believed to be the first philosopher, and epic poets like Homer and Hesiod, who we generally do not count as philosophers. According to Adamson,

> [The Pre-Socratics'] views were, at least implicitly, grounded in *arguments*. This, to me, is the difference between early Greek philosophy and other early Greek cultural productions. (Adamson, 2014, p. 6)

Here, of course, we use the term 'argument' in its technical philosophical sense. 'Having an argument' can also mean having a contention or a debate. Conflating these two meanings could lead to something like we see in the famous Monty Python's *Argument Clinic* sketch.[14] In this sketch, a man (M) enters an argument clinic looking to have a good argument. After having paid five pounds, he finds that the professional arguer (O), instead of arguing, just contradicts everything the man says. Exasperated, the man exclaims

> M: I came here for a good argument!
>
> O: AH, no you didn't, you came here for an argument!

---

[14]The Argument Clinic sketch can be viewed on Monty Python's official YouTube channel: `https://www.youtube.com/watch?v=kQFKtI6gn9Y`.

M: An argument isn't just contradiction.

O: Well! it CAN be!

M: No it can't! An argument is a connected series of statements intended to establish a proposition.

O: No it isn't![15]

The definition that the man gives in this sketch—an argument is a connected series of statements intended to establish a proposition—is sufficiently good for our purposes. However, I wish to make the connection between arguments and reasons explicit and will therefore follow Hodges (2001, p. 36) in maintaining that an argument is what someone produces when they make a statement, which I call a *thesis*, and when they give reasons for believing that statement. These reasons are the premises of the argument.

Reasons are statements that are either used to justify an action or to prove an assertion. In the first case, we are dealing with practical reasons, in the second case, theoretical ones. Practical reasons (aim to) give rise to intentional actions, while theoretical reasons (aim to) give rise to a modification of a belief.[16] In the rest of this dissertation, I focus only on philosophers giving theoretical reasons and, therefore, only on philosophers who aim to modify beliefs, not on those who aim to modify actions.

Although I think Hodges's definition is a good one, there is one notion

---

[15]The reader can find this script on Monty Python's official website: `http://www.montypython.net/scripts/argument.php`.

[16]See Harman (1986) for a more detailed discussion of this distinction.

used in the definition given by the man in the Monty Python sketch that I do wish to emphasize and which cannot explicitly be found in Hodges's own definition; the notion of proposition.[17] If we think about the nature of arguments, it is important to note that arguments consist of expressed propositions, not sentences or utterances.

It is hard to say what propositions exactly are, and I won't be able to solve all the issues regarding the nature of proposition here. But there are some things I do want to explicitly say.

A proposition is often thought to be the meaning of a declarative sentence, or that which could be expressed by a declarative sentence (and the context in which this sentence is uttered, and perhaps, evaluated).[18] I say "could be expressed by a declarative sentence" because there probably are many true propositions which have never actually been expressed. Uninteresting true arithmetical propositions come to mind, but I'm also confident that there are a lot of things that could truthfully be said about my dog that have never actually been expressed since many of her bad behaviors are simply never observed by anyone.

Propositions are often associated with the set of possible worlds in which they are true, and I will follow this convention since it will make it easier for me to relate a Bayesian Epistemology framework to a logical one later in

---

[17]I do, however, take Hodges to follow Strawson, and use the term 'statement' to mean what I mean with the term 'proposition'.

[18]See the intriguing and controversial recent book by John MacFarlane for a detailed discussion on the notion of context of evaluation (MacFarlane, 2014).

this dissertation.

Although propositions are abstract objects, one need not necessarily be a die-hard Platonist about them. That is, one need not think that they have existence independent of any mind in a substantial metaphysical sense. As Quine (1960) has argued, such a view is problematic in itself because propositions do not seem to have clear identity conditions. Furthermore, there are alternatives. Propositions could, for instance, be sentence types in the language of thought that need not be instantiated.[19] Or one could defend a deflationary account of propositions in which propositions are nothing but a convenient fiction needed to explain what we say and think.[20] Propositions would then be just as metaphysically spooky as centers of gravity are.

I do not know which of these views on propositions, if any, is correct and wish to stay neutral on this issue. But propositions, whatever they are, are assigned two important properties. First, they are considered to be shareable objects. For instance, two people can think about or express the same proposition. Second, they are considered to be the primary bearers of truth and falsity. That is, sentences are only true or false because they express true or false propositions.

It is hard to see how sentences themselves could be the primary bearers of truth, because sentences can be ambiguous. A sentence is a linguistic object that consists of words, and the same sentence could be true in one sense, and

---

[19]A view like this is explored by (Soames, 2010).

[20]For instance, a deflationary account of propositions is defended in (Schiffer, 2003).

false in another. For example, my friend could truthfully write to me about a mutual friend that "He fed her dog food", but I need not conclude from this that I should refrain from eating his famous stew when I visit. Furthermore, the same sentence can be used to state a truth in one circumstance but a falsehood in another. For instance, my friend could write me truthfully that "He fed her dog" on Wednesday morning, but falsely on Wednesday evening. It is hard to explain why the same sentence is true and false at the same time, or at different times, if one cannot refer to the different meanings that the sentence expresses in different circumstances.

Furthermore, it is important that the constituents of an argument, that is, the thesis that is argued for, and the reasons that are given for it, are shareable. One wants to persuade someone to believe the very same thing that one is arguing for, and one wants her to be persuaded because she has grasped the very same reasons that one has given. Of course, one could argue that this only means that arguments need to consist of sentence types, not tokens, because sentence types are also shareable. However, one can give the same argument by using different sentences. For instance, "Every man is mortal, therefore Socrates is mortal" is the same argument as "All men are mortal, therefore Socrates is mortal".[21] Furthermore, I can give the same argument again in a different language. For instance, "Alle Menschen sind sterblich, also ist Sokrates sterblich" expresses the same argument, but

---

[21]I'm aware that it is tricky to say what the identity conditions of arguments exactly are. Nevertheless, I take the identity of the above arguments to be intuitively evident because the two arguments have the same meaning.

now in German. In order to explain this, one needs to appeal to what these sentences mean. That is, one needs to be able to say that the premise and the conclusion pair in all three arguments express the same thing.

I have already said that I will only focus on arguments that invoke theoretical reasons, that is, reasons that aim to modify beliefs. But whose beliefs are philosophers trying to modify?

It is often assumed that the aim of an argument is to persuade another to accept a thesis that the arguer herself already supports. Such persuasive arguments definitely exist in philosophy. However, focusing solely on persuasive arguments leaves out an important class of arguments that can be called *exploratory arguments*. These exploratory arguments are aimed at investigating what, if anything, it is rational to believe about a certain thesis. That is, exploratory arguments are aimed at modifying one's own beliefs by putting forward possible reasons for certain beliefs and test those reasons.[22]

Of course, philosophers are not interested in persuading others, or themselves, to accept or reject a thesis by any argument that does the trick. In ancient times, this is exactly what philosophers like Plato and Aristotle thought distinguished their argumentative practices from the rhetorical practices of sophists like Gorgias and Callicles. Philosophers are not primarily interested in persuasion. They are primarily interested in persuading someone to believe the truth. This is why philosophers only want to modify

---

[22]See Meiland (1989) for a detailed discussion of the distinction between persuasive arguments and exploratory arguments, and why it is important to acknowledge both kinds. A more recent discussion of these two kinds can be found in (Johnson & Blair, 2006).

beliefs, of others or themselves, by producing good reasons to accept a certain thesis (ideally speaking, of course). That means that, in a philosophical argument, the thesis and the reasons given for accepting it need to stand in a certain relation to one another. In particular, the reasons given need to sufficiently support the conclusion.

There are different ways in which reasons can support a conclusion.[23] In the next two sections, I argue that, for paradigmatic philosophical arguments, the reasons given for the thesis argued for must support the thesis in the strongest possible sense, i.e. the premises of an argument need to rationally necessitate the acceptance of its conclusion. This is because of two other characteristics of philosophical practices that I wish to discuss: in the ideal case, philosophical theses should only be accepted in virtue of conclusive reasons (and not in virtue of any authority), and philosophical theses need to be established in a context of high epistemic standards, i.e. beyond a reasonable doubt.

3.2.2 Epistemic Autonomy

In the previous subsection, I argued that one family characteristic of philosophy is that claims need to be supported by arguments and reasons. In this subsection, I discuss a second important family characteristic of philosophical practices, namely that philosophical theses should only be accepted on

---

[23]Charles Sanders Pierce's threefold distinction between deductive, inductive, and abductive support comes to mind here. In this dissertation, however, I only distinguish between deductive support and inductive support. See section 3.3.1.

the basis of these arguments and reasons. Philosophical claims should not be accepted because of the authority of the speaker or the emotions that the argument invokes. That is, in philosophy there is the ideal of *epistemic autonomy.* In particular, it means that one's knowledge does not essentially depend on trusting the testimony of another.

The point is not, of course, that philosophers never accept, or even should accept, a thesis on the basis of testimony. Such a requirement would be too demanding. We are finite beings with limited resources. Even if we had the mental capacity, none of us have the time to learn everything that there is to know. Nobody can learn all of mathematics, all of physics, or even all of philosophy. This is why specialization is necessary. Much of what we know depends on trusting experts. Rather, the point is that, if one has an ideal philosophical argument, a philosopher doesn't need to (Easwaran, 2015, n7).

A case in point is Gödel's *Incompleteness Theorem* (Gödel, 1931). Many philosophers have simply accepted Gödel's claim that no consistent finitely axiomatized formal system containing basic arithmetic can prove all true arithmetic statements, without having completely understood its proof.[24] That is, of course, not problematic. Many competent people have checked Gödel's proof, and others have devised alternative proofs. There is nothing wrong with trusting the experts here.[25]

---

[24]Could the reader produce any proof of the Incompleteness Theorem immediately, i.e. without consulting a reference book, if asked?

[25]One could perhaps argue that accepting arguments on the basis of testimony or authority is, in a sense, a philosophical vice. It is a bad habit, even though it might sometimes be necessary because of our finite and fallible natures. One could think of this as playing

However, trusting the expert is not sufficient. Most mathematicians believe that Pierre de Fermat was an exceptionally gifted mathematician whose mathematical judgments can, in general, be trusted. Nevertheless, the note that Fermat left in the margin of his copy of Diophantus's *Arithmetica*,[26] was not sufficient to establish what is now known as *Fermat's Last Theorem.* This theorem was only a conjecture until Andrew Wiles presented a proof in (Wiles, 1995). The difference between Gödel's Theorem and Fermat's Last Theorem (before it was proved) was that one didn't need to rely on Gödel's authority (or anyone else's) if one didn't want to. Gödel's proof is transferable, while Fermat's note is not.

The notion of transferability that I want to use here was first developed by Kenny Easwaran in (**?**) and (Easwaran, 2015). Although Easwaran mostly talks about transferability as a necessary condition for mathematical *proofs*, his focus on proofs is not essential. Since proofs are a kind of argument, we can easily apply his reasoning to *arguments* in general. An argument is transferable, according to Easwaran, if it presents a sequence of reasons that a suitable audience will find compelling on mere consideration of those

---

an opening sequence in chess from memory, without really understanding the purpose of any of the moves. Playing in this way might improve one's game in the short run, and perhaps this is all one cares about, but will ultimately hamper one's ability to understand the game in the long run. Similarly, presenting arguments from memory rather than understanding might impress one's audience in the short run, but one would have become a better philosopher if one had thought through the argument for oneself.

[26]"[I]t is impossible to separate a cube into two cubes, or a biquadrate into two biquadrates, or generally *any power except a square into two powers with the same exponent.* I have discovered a truly marvellous proof of this, which however the margin is not large enough to contain". (Heath, 1910, p. 144-5)

reasons with no essential dependence on the authority or trustworthiness of the author. In his own words (while writing 'argument' for 'proof')

> the basic idea is that a[n argument] must be such that a relevant expert will become convinced of the truth of the conclusion of the [argument] just by consideration of each of the steps in the [argument]. (**?**, p. 343)

It is precisely because we can check for ourselves whether Gödel has really established the Incompleteness Theorem (given, of course, that we are willing to invest in acquiring sufficient expertise to be able to judge this) that we may accept these results on trust. In contrast, we may not accept Fermat's Last Theorem because we trust in Fermat's expertise, even though we all do believe him to be an expert, because his note is not a transferable argument. If we accept Fermat's Last Theorem at all, it is because we trust Wiles (and the peer-review process), not because we trust Fermat. And that is because Wiles's proof is transferable, just as Gödel's is.[27]

We have seen above that it is not sufficient to trust someone's testimony that they are indeed an expert. In her paper "Second-Hand Knowledge", Elisabeth Fricker tried to formulate the conditions under which it is indeed allowed to accept a proposition P based on trust in an expert:[28]

---

[27]One reason for this is that, given the complexity of Wiles's proof and no reason to think a much simpler proof is forthcoming, it is highly unlikely that Fermat had actually found a proof of Fermat's Last Theorem. However, this is inessential. We would not accept Fermat's Last Theorem because Fermat claimed it to be true even if we suppose that Fermat did manage to find a proof for it.

[28]The reader can also consult Douglas Walton's *Appeal to Expert Opinion* (1997) for a

> One properly accepts P on the basis of trust in another's tes-
> timony that P—her word that P—just if she speaks sincerely,
> and she is epistemically well enough placed with respect to P so
> that were she to have, or make a judgment to form, a conscious
> belief regarding whether P, her belief would almost certainly be
> knowledge; and she is better epistemically placed with respect to
> P than oneself; and one recognizes all these things to be so and
> one is not aware of significant contrary testimony regarding P.
> (Fricker, 2006, p. 232)

The above characterization is sufficient to rule out accepting Fermat's
Last Theorem on the basis of Fermat's note because, even though we can
assume Fermat to have written down the note sincerely, given the centuries
long search for a proof by competent mathematicians and the complexity
of the proof Wiles eventually found, we have good reasons to believe that
Fermat was not sufficiently well enough placed epistemically for us to trust his
testimony. Note, however, that Fricker does not demand that the testimony
of the authority is grounded in a transferable argument. This is because
we need to make a distinction between being *epistemically autonomous* and
acting *epistemically responsibly.*

Acting epistemically responsibly means that one trusts some testimony
but only because it comes from a reliable source. Properly speaking, Fricker
tries to formulate conditions for acting epistemically responsibly, not for be-

---

detailed discussion of when reliance on expert opinion is acceptable in argumentation.

ing epistemically autonomous. In the example above, a philosopher who trusts the mathematical community that the first Incompleteness Theorem and Fermat's Last Theorem are indeed true are not epistemically autonomous but they are acting epistemically responsibly. However, as Fricker notes, "this knowledge from trust in testimony is knowledge at second hand (or third, or fourth . . .), and as such my epistemic position vis-à-vis what I know is in at least one respect inferior to when I know at first hand" (Fricker, 2006, p. 240).

Even though epistemically inferior, depending on reliable testimony cannot always be problematic in science. That is because not all scientific knowledge is transferable. Many scientific results depend on experiments that cannot be easily repeated. This could be because of practical reasons. I cannot, for instance, verify the results that indicate the existence of the Higgs-Boson reported in (ATLAS, 2012) and (CMS, 2012), not even if I really wanted to, because I do not have access to the Large Hadron Collider (LHC) and it is unlikely that I'll be able to build a similar particle accelerator. Projects like the LHC and the International Space Station are so expensive and complex that they are necessarily international collaborative projects. Of course, the reader could insist that these experiments can still be replicated by anyone who feels the need to do so *in principle*. Building and operating a sufficiently powerful particle accelerator is perhaps a more substantial investment of resources than learning advanced logic, but it does not differ in kind. However, some scientific experiments cannot be replicated even in principle. There are

scientific experiments that can only be performed once (or a limited number of times) because they destroy (part of) the object under study. This is, for instance, prevalent in the field of archaeology. Once a sliver of a religious icon has been removed and dissolved in order to determine its chemical constitution, the object under study is damaged forever. Such experiments can only be repeated so many times before there is nothing left to study. Also, speaking of history more broadly, historical events are unique and can only be witnessed once. Ultimately, as a historian, one needs to trust the testimony of one's sources. One cannot go back in time and go look for oneself.[29]

However, it is important to note that, although not every person can know everything independent from testimony, this does not mean that there are known propositions that cannot be known but by testimony. When a proposition is known through testimony there must be a chain of testimony, no matter how long, that ends with someone knowing the proposition immediately. This is why Fricker formulates what she calls criterion **T**:

> If H knows that P through being told that P and trusting the teller, there is or was someone who knows that P in some other way—*not* in virtue of having been told that P and trusting the teller. (Fricker, 2006, p. 240)

From this it follows immediately as a corollary that

---

[29]Many sciences have such a temporal component where ultimately recorded observation must be trusted. This is also the case, for instance, for the observation of astronomical events. These events cannot (currently) be reproduced, only observed as they happen.

> For any proposition *P* that can be known, there must be some
> way other than trust in testimony through which *P* can or once
> could be known. (Fricker, 2006, p. 241)

We can now use this corollary to argue that the regulative ideal of complete philosophical autonomy is achievable in principle, just as it is achievable in mathematics.

Easwaran himself does not claim that all ideal philosophical arguments are transferable, although he grants that some are. As an example, he mentions Gettier's famous argument against the definition of knowledge as justified true belief that I discussed earlier (Gettier, 1963). According to Easwaran,[30]

> [a]s long as [the philosopher] shares the intuitions about the specific cases at hand, she will come away believing that knowledge is not justified true belief. Thus this argument is just as transferable as any mathematical argument. (**?**, p. 360)

From Fricker's corollary it follows that all philosophical theses that can be known, can be known independent of testimony. This means that either

---

[30]The reader might fear that there is an important difference here between the philosopher and the mathematician because the philosophical argument is only transferable conditionally. According to Easwaran, Gettier's argument is only transferable if the philosopher shares the same philosophical intuitions. But I don't think that this difference is essential for the possibility of transferability of philosophical arguments. The mathematician similarly only accepts a proof on the condition that certain axioms are true (and through them, previous results). I do, however, agree that there is an important epistemic difference in kind between accepting philosophical intuitions and mathematical axioms. I will return to this issue in the next chapter.

they can be immediately known, or they can be established by a transferable argument. Since, as I argue in the next chapter, immediate philosophical knowledge such as intuitions should, ideally speaking, be *a priori*, there is no need in philosophy to ever trust the testimony of another.

3.2.3 Highest Epistemic Standards

The third and final family characteristic of philosophical practices that I want to discuss is that philosophical debates are typically performed in contexts of the highest epistemic standards. That is, they must be established beyond any reasonable doubt.

Compare the following three knowledge claims:

Mathematician:

I know that there are infinitely many prime numbers.

Doctor:

I know that smoking causes lung cancer.

Philosopher:

I know that I have two hands, and that, therefore an external world exists.

Let us grant that the mathematician and the doctor can indeed claim knowledge. This is a quite standard assumption in epistemological debates, so I take it that it is an unproblematical assumption.

It is clear that, although both the mathematician and the doctor claim to know something, they do not claim to know their respective theses in the same sense. If we take the meaning of a word to be its use (Wittgenstein, 1953), we can appreciate that the mathematician will use the word "know" differently from the doctor. For example, when we ask the mathematician and the doctor how they know what they claim, they will give different answers:

Mathematician:

I can prove it.

Doctor:

I can't prove it, but I have practically irrefutable scientific evidence.

That is, the mathematician and the doctor have different standards that they must meet before they are entitled to claim knowledge. Similarly, they have different methods by which they acquire knowledge and methods by which they critique claims of knowledge.[31] These standards are, in general, well understood, and practitioners of the formal and the empirical sciences are rigorously trained in what the standards of their disciplines are.

One important difference, for instance, is the role of counter-examples, or even the possibility of finding counter-examples. In mathematics, ruling out

---

[31]The above view is a form of *Epistemic Contextualism*. See (Rysiew, 2016) for an introduction to *Epistemic Contextualism*. Prominent defenders include Fred Dretske (1981) and David Lewis (1996). For a strong critique of Epistemic Contextualism, see Williamson (2010).

every possible counter-example is paramount. For instance, it is not sufficient that a property holds for all numbers that you have observed, even if that has been many.

A famous example is the claim that for every natural number n, $n^{17} + 9$ and $(n + 1)^{17} + 9$ are relatively prime. If you would program a computer to check this for every n, assuming that you could perform the couple of hundred trillion operations per second that is the current technological best we can do for such projects,[32] you could check for a trillion trillion years and not find a single counter-example. The first counter-example is 8424432925592889329288197322308900672459420460792433.[33] While such a check is running, mathematicians are not entitled to claim that they know such a thesis holds, even if their computer program has been running for, say, ten thousand years with no counter-example observed. Just the bare fact of not having ruled out every possible counter-example is sufficient to block a mathematician from claiming knowledge, no matter how unlikely it is that a counter-example can be found.

On the other hand, finding a hundred year old life-long smoker who did not develop lung cancer does not refute the doctors claim that smoking causes cancer. It is maybe true that not all smokers develop lung cancer, but we can show that it is much more likely for a smoker to develop lung cancer than for a non-smoker. In the context of medicine, that is good enough.

---

[32]See https://www.mersenne.org/
[33]See https://oeis.org/A010034.

Philosophical claims are much more similar to mathematics in this regard, than to medicine. Like the mathematician, philosophical claims have to be immune to all kinds of counter-examples, no matter how far-fetched, to exclude such possibilities of being misled by evil demons, being prodded in the brain by evil scientists, and being in a country full of fake barns. Insisting that you have no reason to believe that you are not a brain in a vat, and that you find it extremely unlikely that you are, is not considered to be a sufficient response in the face of philosophical skepticism.[34]

To see that the standards of being able to claim knowledge are raised in a philosophical context even above the already strict scientific standards, consider the following. Philosophers can problematize unproblematic scientific claims of knowledge, like the above, by applying their own philosophical standards to them.

We philosophers, for instance, can accept that the mathematicians have proven that there are an infinite number of primes to their own satisfaction while still subjecting their claims to further critique. Strict Finitists, like Jean Paul van Bendegem, for instance, deny the existence of infinite mathematical objects (Bendegem, 2012). According to them, mathematical claims are only meaningful when they are about potentially concrete things and concrete operations. That is, a construction is only accepted if it could actually be performed. Since the universe is believed to be finite, constructions

---

[34]Mooreans might disagree with this claim, but Mooreanism itself is highly controversial within philosophy.

that takes an infinite amount of time and resources are not accepted as legitimate. Therefore, mathematical claims that depend on such constructions are considered to be meaningless.

Of course, Strict Finitisim is a fringe position, and is not broadly defended within the philosophy of mathematics. It is, however, considered to be a legitimate philosophical position that is up for philosophical debate. This means that, until Strict Finitism is conclusively refuted, philosophers cannot claim to know that there really are an infinite number of primes. Similarly, philosophers but not the medical community need to take seriously the possibility that a powerful evil demon could be tampering with all the empirical evidence, rendering all our medical knowledge into doubt. That shows that, typically, philosopher need to take relative alternatives into account that the mathematician and the doctor can simply ignore.

The reader might believe these philosophical standards to be too strict, and prefer a fallibilist account of philosophical knowledge, although they are widely accepted within the philosophical community as we can see from philosophical practice. But, if nothing else, this section has then hopefully convinced these reader that philosophers should be more exact in explicating when they can successfully claim philosophical knowledge and when they cannot.

3.3   The Deductive Nature of Philosophical Arguments

In the last section, I discussed three family characteristics of philosophical practices; 1) they depend on arguments and reasons, 2) there is an ideal of epistemic autonomy, and 3) philosophical knowledge is evaluated in a context of high epistemic standards.

In this section, I aim to show that, due to the last two characteristics discussed above, only deductive arguments can be used to establish philosophical knowledge. Inductive and abductive arguments do not suffice because they cannot be evaluated by solely considering the argument themselves.

In order to make these claims clear, it would be helpful to have a formal framework in which these claims can be clearly understood. In the next subsection, I therefore first present a discussion of deductive and inductive arguments within a semi-formal possible world framework.

3.3.1   A Bayesian Framework

As indicated above, an argument consists of a proposition, the thesis (also known as the conclusion of the argument), and a set of propositions which are the reasons given for accepting the thesis. This set of propositions is often called the set of premises.

This set of reasons can offer various degrees of support to the thesis under discussion. In this dissertation, I distinguish between two kinds of support.

First, the set of reasons can guarantee the thesis. When this is the case, I'll say that the set of reasons deductively supports the thesis and I'll call the argument a deductive argument. Second, the set of reasons can make the thesis more likely. In this case, I'll say that the set of reasons inductively supports the thesis and I'll call the argument an inductive argument. These two kinds of arguments can be understood from a single formal framework that makes use of the idea of *possible worlds.*

For the moment, I wish to remain neutral on what possible worlds exactly are.[35] For now, it is sufficient to think of each possible world as a maximally consistent set of propositions.[36] This set of propositions indicates which propositions are true in that world. All other propositions are false in that world. We'll say that possible worlds are the same when they agree on the truth-value of every proposition.[37]

We usually distinguish between two kinds of possibility: alethic possibility

---

[35]The three best developed accounts in the literature are: (a) real existing maximally connected objects (Lewis, 1986), (b) total sets of possible states of affairs (Plantinga, 1976), or (c) a recombination of the simple atomic facts of the actual world (Armstrong, 1997). Importantly, one need not be ontologically committed to the existence of possible worlds, either as really existing or as *possibilia*, if one wants to make use of a possible world semantic framework.

[36]A set of propositions is maximally consistent if a) that set is consistent and b) every proposition that is consistent with that set is an element of that set. Intuitively, this means that no proposition that is not in the set can be added without the set becoming inconsistent.

[37]As we will see when we discuss *logical pluralism*, note that conceiving the nature of possible worlds differently has serious consequences for which propositions follow from which. From the brief characterization above, it is already easy to see that, every proposition is assumed to have a determinate truth-value in every possible world, must be either true or false in that world, and there are no worlds in which contradictory propositions are true. It turns out that we can fruitfully change any of these assumptions and get a different logic.

and epistemic possibility. Alethic possibility is about which worlds are really possible (and which are not). As such, the study of alethic possible worlds falls under metaphysics. In contrast, a world is epistemically possible if, *for all we know*, the possible world under consideration could describe the actual world. It is important to see that these are two different notions of possibility and that they can come apart. For instance, *for all we know*, Goldbach's conjecture that every even integer greater than 2 can be expressed as the sum of two primes can be either true or false, although many think that it is highly likely that it is true. Metaphysically speaking, however, Goldbach's conjecture is usually thought to be either necessarily true or necessarily false because most metaphysicians think that mathematical truths are necessarily true (i.e. true in all alethic possible worlds). In this dissertation, I only deal with epistemic possibility. As is customary, I'll designate the set of all epistemically possible worlds with W.

We can clarify both the nature of deductive arguments and the nature of inductive arguments by using this notion of epistemic possible worlds. I start with discussing arguments where the reasons given for the thesis, if true, are thought to guarantee the truth of the thesis, i.e. deductive arguments. Here is a simple example of a deductive argument taken from Lewis Carol's "What the Tortoise Said to Achilles" (Carroll, 1895):[38]

Things that are equal to the same are equal to each other ($P1$).

---

[38]When I discuss *logical pluralism*, the reader will find out that this example is taken on purpose.

The two sides of this triangle are things that are equal to the same (*P2*). Therefore, the two sides of this triangle are equal to each other (*C*).

In the above example, it is usually taken to be impossible for the conclusion (*C*) to be false, while the premises (*P1* & *P2*) are true.

If it is indeed true of an argument that it is impossible for the conclusion (*C*) to be false while the set of premises are all true (Π), we say that the argument is deductively valid and that *C* is a *logical consequence* of Π. We denote this by writing

$$\Pi \models C$$

I trust that the reader immediately grants that when two sides of a triangle are equal to something else, say a ruler, it is indeed a logical consequence that the two sides of the triangle are equal to each other (given that two things that are both equal to a third thing are equal to another). Nevertheless, let us suppose that we are confronted with an interlocutor who does not immediately see that the conclusion necessarily follows from the premises. What could we say to convince this interlocutor? Here is a quick reminder in terms of possible worlds.

As indicated above, I associate every proposition with the set of possible worlds in which the proposition is true. I don't necessarily mean to say that a proposition is this set. Just that every proposition is associated with such

a set. For now, as also indicated above, I wish to remain neutral on the exact nature of propositions.[39]

So how do we convince our interlocutor who is still in doubt about whether the above argument is a valid argument? We will first ask him to conceive of the set of all possible worlds W (although we will allow him to make use of any representational tool to augment his powers of imagination; truth-tables and logical diagrams in particular). We then tell him to only consider those possible worlds where $P1$ and $P2$ are true, that is, only those worlds in which it is the case that things that are equal to a third thing must be equal to another and where the sides of a certain triangle are both equal to the same thing. If it is true in all these possible worlds that C is true, that is, that there is no possible world where the sides of the triangle are of unequal length although they are both equal to a third thing, then the truth of $P1$ and $P2$ necessitates the truth of C. Since our interlocutor would most likely agree that it is inconceivable that there is such a world in which the two sides of a triangle are not equal given the truth of the premises he would therefore agree that the truth of the conclusion is indeed necessitated by the truth of the premises.[40]

---

[39]Personally, I believe that there is no single semantic framework that captures the essence of meanings. Rather, like any framework, a semantic framework may or may not illuminate certain aspects of propositions in fruitful ways. In any case, affirming that a proposition is nothing but the set of possible worlds in which it is true comes with a high cost, such as the commitment that every true mathematical statement says the same thing because they are all true in every possible world. The idea that coming to know that, e.g. Gödel's Incompleteness Theorem holds is nothing but coming to learn that $1 + 1 = 2$ in a different way, is hard to swallow.

[40]There is actually an issue here with what it means to conceive of a possible world, and

Although the technique works in this case, in general, inspecting semantic relationships immediately (either in imagination or on paper) is often insufficient to demonstrate the validity of an argument because, in order to see whether the conclusion indeed holds in every possible world where the premises are true, we sometimes need to be able to represent and inspect an infinite number of possible worlds. Our finite minds unfortunately do not allow for this. Rather, this technique is more often used to show that an argument is invalid by conceiving of a possible world in which the premises are true but the conclusion false, i.e. by conceiving of a counter-example. But, instead of asking our interlocutor to inspect the semantic relationships between the different sets of possible worlds in his imagination immediately, we could have given an argument instead. We would ask our doubting interlocutor to assume that the conclusion is false given the premises and show that this is impossible by forcing her to admit a contradiction.

Such an argument would go as follows: we ask our doubting interlocutor to assume that there exists a possible world where the sides of the triangle are unequal to each other (negation of (C)) but are equal to the same thing

---

whether the ability to conceive or not conceive of a possible world is a sufficient ground for knowing that such a world can or cannot exist. This is because whether something is conceivable or not is a question about the human mind, while whether something is possible or not is a question about the world. Why would certain features of our mind reliably mirror certain features of the world? We often see appeals to evolution here, that is, if our minds hadn't evolved to be able to reliably conceive of what is possible or not, we would not have survived, but a) why wouldn't the mind use heuristics here? and b) why should we believe that we have this ability for all domains? (Perhaps we only evolved to reliably conceive of practical possibilities). For a sophisticated discussion of whether conceivability entails possibility, see (Chalmers, 2002).

($P2$), i.e. that there exists a possible world where things that are equal to the same are unequal to each other given ($P1$). But, given ($P1$), all things that are equal to the same are equal to each other in every possible world that we may consider. Contradiction! Such a world can therefore not exist in the set of possible worlds under consideration. Therefore, given ($P1$) and ($P2$), (C) must be true in every possible world under consideration.[41]

From the above discussion it has now hopefully become sufficiently clear what deductive arguments are, i.e. arguments where the truth of the premises guarantees the truth of the conclusion, and that we can think of deductive arguments as establishing that the conclusion of the argument must be true in every possible world in which the premises are true. We also discussed two methods of how to convince someone that a conclusion does indeed follow if he does not grant so immediately; first, if an argument is valid, than it is impossible to conceive of a possible world in which the conclusion is false given the reasons that are supplied. Second, if one denies the conclusion but

---

[41]This argument can be formalized using a deductive system known as *semantic tableaux*. A semantic tableaux is basically a systematic way of trying to build a counter-example. When every such attempt fails, the conclusion must be a logical consequence of the set of premises. When a conclusion (C) can be derived from a set of premises (Π) by use of a deductive system, we say that that the conclusion is *derivable* from the premises, and denote that by $\Pi \vdash C$. If we set up things in the right way, we can show that every valid argument is derivable and every derivable argument is valid ($\Pi \models C \Leftrightarrow \Pi \vdash C$). For an easy to understand proof of this, see (Smith, 2003). There are other deductive systems, such as *natural deduction* and the *sequent calculus*, that are also sound and complete with respect to the semantics I presented but I have chosen semantic tableaux as the background deductive system here because the relationship between semantics and deductive system is easy to see in the case of semantic tableaux. I assume that this is the reason why semantic tableaux have become so popular in elementary logic textbooks. Furthermore, semantic tableaux have a nice relationship with the *Game of Giving and Asking for Reasons* that I describe below.

grants the premises, then one could be forced to grant a contradiction.

However, not all arguments that are used in life and philosophy are deductive arguments. There are also arguments where the premises do support the conclusion, but do not guarantee it; these arguments are called inductive arguments.[42] Let us see how we can use the possible world framework I set out in this section to think about these kinds of arguments as well.

In inductive arguments the premises, if true, do not guarantee the truth of the conclusion but they *do* support the conclusion to some degree. The idea here is that a set of premises supports a conclusion when the truth of the premises makes it *more likely* that the conclusion is true.

A statistical inference is a good example of an inductive argument; that 45% of a random sample of 3,669 voters from all over the country say that they have voted for the Democratic candidate while only 41% claim to have voted for the Republican candidate does support the thesis that the Democratic candidate will win. Indeed, if a large number of such polls show similar results, this might support the thesis that the democratic candidate has won to such a high degree that it becomes increasingly unreasonable to deny it. Nevertheless, at the end of the day, the truth of the conclusion is not guaranteed by the reasons supplied for it; it is still possible that the Democratic candidate has lost no matter how unlikely that has seemed.

---

[42]Arguments are sometimes categorized as either deductive, inductive, or abductive. In this part, I take inductive and abductive arguments under a single header because they are both ampliative, i.e. their conclusion contains more information than is contained in their joined premises. Inference to the best explanation, which is closely related to abductive reasoning, will play an important part in the next chapter, and will be treated there.

Polling is, of course, a complicated process, and many things can go wrong in setting up and interpreting a poll. But even a theoretically perfectly performed poll can support the wrong prediction even if that becomes increasingly unlikely as the sample size grows larger.[43]

In the possible world framework that I introduced above, the support that an inductive argument provides can be measured by the ratio of the worlds in which the thesis is conceived to be true and the total number of possible worlds in the set of possible worlds that is left after we consider only those possible worlds of W in which the premises are true.[44] We will call this ratio the probability that the conclusion (C) holds given that the premises (R) hold and denote this with $P(C|R)$.[45] Simple probabilities are shorthand

---

[43]That is, of course, until the sample size grows so large that the findings start to guarantee the truth of the thesis under discussion. If we have, for instance, polled 95% of the voters and the majority of them have truthfully claimed to have voted for the Republican candidate, then this guarantees the thesis that the Republican candidate has won. At some point, as the sample size grows, it will become impossible that he has lost. Since this is a case where the thesis is guaranteed by the premises rather than just supported by it, I would call such an argument a deductive argument rather than an inductive one.

[44]There are problems with this characterization of support regarding infinity, since $\frac{\infty}{\infty}$ is undefined and our conceptualization of probability does not seem to be able to deal with irrational probabilities. In case of countable infinities, the spirit of this interpretation can be kept by making use of the notion of *maximum entropy* (see (Jaynes, 1968) for details). In the uncountable case, things get complicated, since there is the risk of running into Bertrand's Paradox. That is, different conceptualizations of the possibility space can lead to different answers concerning what the ratio is. However, accepting Bertrand's Paradox is quite compatible with the argument that I am presenting here, since it implies that two agents could rationally disagree about the amount of support that an inductive argument provides, depending on what background knowledge they decide to bring.

[45]This is only one of two ways in which we can conceive of probabilities. In this thesis, I prefer this interpretation over the *subjectivist* interpretation where probabilities are taken to indicate the betting behavior of ideal agents because it provides a nice uniform framework in which I can consider deductive arguments and inductive arguments together.

for the ratio of all the worlds in which the proposition under consideration holds and all possible worlds. So the probability that C, denoted by $P(C)$, is shorthand for $P(C|W)$. We will say that a reason (R) supports a thesis (T) if $P(T|R) > P(T)$.[46]

The idea here is that, although there are counter-examples, i.e. possible worlds in which the thesis is false while the premises are true, they are comparatively rarer when they are compared to the number of possible worlds in which the thesis holds given that the premises hold. The truth of the premises can therefore be taken as a *sign* that the conclusion is true as well.

It is easy to see that when $R \models C$, i.e. when the conclusion is a logical consequence of the set of reasons, it should be the case that $P(C|R) = 1$, i.e. the probability of the conclusion given the set of reasons should be 1. This is because, since the truth of R necessitates the truth of C, every possible world in which R is true must be a world in which C is true. Normally, this relationship is axiomatically guaranteed when the above is made formally precise in an inductive logic. The development of such a logic is, however, the subject of active research and beyond the scope of this dissertation. In contrast to deductive logics, which are quite well understood, it is still under dispute how an adequate inductive logic should exactly be formulated.[47]

---

[46]Technically, $P(C|R)$ is only defined if R is a proposition. Usually, however, we have a set of premises rather than a single one. When we have a set of premises Π, the proposition R is simply the conjunction of all the premises in Π.

[47]For an introductory presentation of inductive logics, see e.g. (Howson and Urbach, 2006) or (Hawthorne, 2012). The exact details of how an inductive logic should be formulated are not crucial to this dissertation. I will just use and explain relevant details as I need them. This dissertation should be self-contained, but should any details of the frame-

For my argument, only two differences between deductive arguments and inductive arguments will turn out to be crucial. First, in deductive arguments as I have defined them, adding premises to a valid argument cannot undermine the guarantee that the conclusion is true.[48] If an argument establishes a thesis, then it does so definitively. Another way of putting this, is to say that deductive arguments are *indefeasible.* This is not true for inductive arguments, where new information can shrink the set of possible worlds to a subset where the probability of the conclusion could be anywhere from 0 to 1. That is, adding premises to an inductive argument can undermine the support for a conclusion.

Second, in deductive arguments, whether a set of premises guarantees a conclusion is only dependent on the logical structure of the premises and the conclusion. That is, no matter what the value of the thesis $P(T)$ is taken to be, a valid deductive argument establishes that $P(T|R) = 1$. All possible worlds in which the set of reasons R hold are worlds in which the thesis T is true. This is not true for inductive arguments in general.

We can see that, in inductive arguments, the level of support that a set of

---

work I'm working with be unclear, I follow the inductive logic presented in (Hawthorne, 2012).

   [48]This is also true for relevance logics although they are non-monotonic; that is, adding irrelevant premises to a relevantly valid argument can make the argument invalid. However, importantly, this is not because there is a change in the ratio of counter-examples versus possible worlds under consideration. Rather, adding irrelevant premises breaks an additional constraints on the logical consequence relation. Relevant consequence relations do have the property of Necessary Truth Preservation. Like (Beall and Restall, 2006), I will only call an argument deductive if the consequence relation that relates the set of premises and the conclusion necessitates the truth of the conclusion. That does mean that, in my terminology, arguments depending on e.g. default logic are not deductive.

reasons lends to its conclusion does not solely depend on its logical structure by looking at an important theorem of probability theory. Bayes' Theorem, a theorem derivable from the axioms of probability first proved in the 18th-century by Reverend Thomas Bayes, states that

$$P(T|R) = \frac{P(R|T)}{P(R)} \cdot P(T)$$

Importantly, this means that the probability of the thesis given the reasons depends on how probable the reasons are given the thesis $P(R|T)$, how probable the reasons are $P(R)$ and how probable the thesis $P(T)$ is thought to be. The latter is often called the *prior probability*. Here I will only focus on the dependency of how probable the thesis is thought to be, also called the *initial plausibility* of the thesis, but it is good to bear in mind that the support that an inductive argument supplies also depends on how plausible the premises are taken to be. In our possible world framework, this is expressed by the ratio of worlds where T is conceived to hold, the ratio of worlds where R is conceived to hold, and the conceived ratio of R worlds in the subset of T worlds.

According to our current best knowledge, there is no unique way to determine the initial plausibility of a thesis *a priori* (Hawthorne, 2012). This is because the laws of probability do not determine a unique probability distribution but rather a whole family of distributions. This is known as *the problem of the priors*. Some philosophers (called *Subjective Bayesians*) are not

bothered by this, and claim that an epistemic agent may indeed choose any coherent prior probability distribution. Others (called *Objective Bayesians*) attempt to put additional *a priori* constraints on what prior probabilities are acceptable. One such example is demanding a *principle of indifference*, which means that an agent should assign equal probability to all independent possibilities. However, to date, every such proposal has turned out to be problematic and no proposal has reached universal agreement. This means that, until the problem of the priors is solved, the support that a set of reasons gives to a thesis is dependent not only on the content of the propositions, but also on how the agent has chosen to assign her prior probabilities. This means that different agents, depending on how they have assigned their prior probabilities, can disagree about the degree of support that a set of reasons gives to a thesis and even about the fact whether a thesis is supported by a set of reasons at all. In short, there is a lot of room for differences between rational agents in how to conceptualize epistemic possibilities.

3.3.2   Inductive Arguments are not Fully Transferable

In the previous subsection, I introduced a semi-formal framework that makes use of *possible worlds* in which I accounted for both deductive arguments and inductive arguments in order to be able to compare them. We saw that these two kinds of arguments have different formal features. Remember that, while whether a deductive arguments supports a conclusion or not only depends on the content of its propositions, whether an inductive argument supports

a conclusion or not also depends on how likely the thesis under consideration is thought to be. In this subsection, I show that this feature of inductive arguments means that inductive arguments are not fully transferable. Using inductive arguments to ground philosophical knowledge is therefore in tension with the ideal of epistemic autonomy, and, as such, we cannot truly know any philosophical truths on the basis of inductive arguments (although we may come to believe them on the basis of such arguments, of course).

For now, let us just say that there is some threshold of support that an argument needs to lend to a philosophical thesis in order to sufficiently justify it for a relevant expert to become convinced. In this discussion, we will only demand that this threshold be greater than 0.5 but, in reality, it seems more reasonable to demand that this threshold has to be close to 1 given our assumption that philosophical arguments are evaluated in contexts of high epistemic standards. This seems reasonable, because it seems unacceptable to say that someone knows a proposition if she thinks its negation is equally likely or perhaps even more likely. Furthermore, in order for a relevant expert to be able to become convinced, her degree of belief in the thesis should be smaller than this threshold.

Given these preliminaries we can now see that something other than the mere consideration of the steps of the argument is needed for a relevant expert to become convinced of a philosophical thesis by an inductive argument. Namely, this relevant expert must also find the thesis sufficiently likely to begin with.

Remember that, in an inductive argument, the degree of support that a set of reasons lends to a thesis is not only dependent on the content of the propositions (as it is in the case of deductive arguments). It also depends on the *prior plausibility* of the thesis. However, as discussed above, the prior plausibility of the thesis is not, in general, thought to be uniquely determined. So, in general, there is no reason to suppose that the one who is presenting the argument and the one who is considering the argument agree on how likely the thesis under discussion was before they considered the argument. This means that the one who presents the inductive argument and the one who considers it can disagree about the degree of support that the reasons provide to the thesis under consideration. In order for a relevant expert to become convinced of the conclusion of an inductive argument, it is not sufficient for her to consider only the steps of the argument, she must also (roughly) agree with the originator of the argument about how likely the conclusion was to begin with, or defer to his judgment of the initial plausibility, in order to become convinced. In other words, inductive arguments are not fully transferable.

She must either trust that her judgment of the initial plausibility was wrong and act like she was convinced although she was not, or she'll remain unconvinced. Even if an inductive argument happens to convince a particular expert on consideration, that does not mean that she can now use the argument to convince any expert that is not yet convinced of the thesis. This leads me to the second aspect of epistemic autonomy; if philosophical argu-

ments should aim to convince a universal audience, then inductive arguments are not good enough, because they cannot fulfill this aim given the different prior distributions within this audience.

3.3.3 Inductive Arguments cannot Universally Convince

In the previous subsection, I argued that inductive arguments are not fully transferable. Whether a relevant expert who is considering the argument believes that its conclusion is sufficiently supported by its premises will depend on objective features of the argument itself, but also on the subjective feature of how plausible the expert thought the conclusion to be in the first place. This is closely related to a second feature that inductive arguments lack: they cannot be decisive.

Being epistemically autonomous means that one is allowed to deny the conclusion of an argument if it does not convince you even if it convinces everyone else. However, it is obvious that the particular details of who the relevant expert is should not matter. We don't want the success of a philosophical argument to depend on the particulars of any specific expert, not even when it is the completely epistemically virtuous Socrates. If we tailor our arguments to only those that we speak to, we might be all too prone to create philosophical *echo chambers*. Similarly, Van Inwagen claims that a philosophical argument is only successful if it would always win the assent

from the members of a neutral audience who have listened to an ideal presentation of the argument (van Inwagen, 2006, Ch. III). The qualification of *neutral* here only means that the members of the audience are still open to argument. They haven't definitively decided yet that the thesis is true or false. That is, they can still be persuaded either way.[49]

In order to screen off any of the particulars of the audience, we should therefore conceive of the audience of relevant experts as exemplifying what Perelman and Olbrechts-Tyteca have called the *universal audience* (Perelman & Olbrechts-Tyteca, 1969). This universal audience consists of all possible reasonable (human) beings.[50]

Let's interpret Van Inwagen's criterion and the idea that philosophical arguments are aimed at a universal audience in terms of the framework that I have set out above. In this framework, the criterion that an ideal philosophical argument convinces all neutral rational beings can now be translated into the demand that the argument brings every possible member of the audience over the threshold for belief. That is, it should convince any agent no matter what their prior probabilities for the thesis under discussion are.

I take it that what it means to be a neutral person is that one can still be persuaded of either the truth or the falsity of the thesis. This means that

---

[49]This doesn't mean that the members of the audience do not have strong reasons to think the thesis true or false. What is important here is that they do not dogmatically hold either the thesis or its contradictory, but can still be rationally persuaded of either alternative.

[50]For a recent textbook discussion of different kinds of audiences, see (van Eemeren et al., 2014, §5.4).

her prior degree of belief of the thesis under dispute cannot be 0 or 1. Using our possible world framework this is easy to see; someone who thinks that $P(T) = 1$ believes that T is true in every possible world; there is no subset to which the premises can shrink the space of possible worlds where T is false. Similarly, someone who holds that $P(T) = 0$ believes that T is false in every possible world; there is no subset to which the premises can shrink the space of possible worlds in which T is true.[51] That is, no matter what evidence is presented, those who have made up their minds are immune to evidence. However, apart from the extreme values of 1 or 0, our neutral rational agent my have any prior degree of belief with respect to the thesis under dispute.

Here I argue that only arguments that establish their conclusion as certain, i.e. arguments such that for any acceptable prior distribution $P(T|R) = 1$, can meet this criterion. This is because if $P(T|R) < 1$ there is always an acceptable prior distribution $0 < P(T) < 1$ such that the posterior probability after updating falls short of the knowledge threshold (whatever we take this threshold to be as long as it is bigger than 0.5).

We can prove this fact mathematically rigorously, but using our possible world framework it is easy to see it intuitively. Suppose that we have a rational agent who thinks that the thesis is so unlikely that there is only

---

[51]There is a small technical caveat here. If the premises shrink the space of possible worlds to the empty set, T is necessarily true given the premises since, technically, T is true in all possible worlds in which the premises are true. This is why in classical first-order logic everything follows from an inconsistent set of premises (*ex falso sequitur quodlibet*). However, no inconsistent argument will convince our universal audience because it inhabits the actual world and as such knows that the space of possible worlds is not empty (I think, so at least the actual world exists; a modern version of the Cartesian *cogito ergo sum*).

one possible world in which it is true. Remember that we have excluded all arguments such that $P(T|R) = 1$ as being instances of deductive arguments. The best any inductive argument can therefore do in this case is to shrink the space of all possible worlds to two worlds; one where T is true and one where T is false. In that best-case scenario $P(T|R) = 0.5$, which is insufficient to bring our agent over the threshold for being convinced that T.

3.4   Philosophy as a Game of Giving and Asking for Reasons

Now that I have argued that philosophical practices tend to share certain family characteristics, most importantly that philosophy proceeds by giving arguments and reasons, and that these arguments must be deductive in nature and need to establish their thesis as beyond a reasonable doubt, I am ready to present my model of philosophical practice.

The model that I present in this section is inspired by Robert Brandom's *Game of Giving and Asking for Reasons.*[52] The model is based on the idea that philosophy can fruitfully be conceived of as a language game that is played by a philosophical community of epistemic peers where these peers

---

[52]The reader can find Brandom discuss the game of giving and asking for reasons, which he believes underlies all meaning, in (Brandom, 1994) and (Brandom, 2000). For those who feel that textual discussions of this kind of game remain too abstract, John MacFarlane wrote an interesting computer program in *Ruby* which implements a simplified version of this game that the reader can actually play. The reader can find it at `http://johnmacfarlane.net/gogar.html`.

are trying to persuade each other of philosophical claims.[53] That scientific practices in general can fruitfully be conceived of in terms of such games of giving and asking for reasons is discussed in (Zamora Bonilla, 2006).

Two normative notions will turn out to play a major role in the game of giving and asking for reasons; *commitments* and *entitlements.* A player is *acknowledging* a commitment to a proposition (or is *attributed* a commitment to a proposition), when that player acknowledges that she can be held accountable for holding it (or is taken to be accountable for holding it). When a player is acknowledging a commitment (or being attributed it) she can be asked a reason why she thinks that she is entitled to this belief. A player is entitled to a proposition (or claims to be entitled) when her belief in that proposition is grounded in *sufficient reasons* (or thought to be grounded in such reasons).

Now that these preliminaries are out of the way, let us look in more detail at this game. Let me first set out the different players engaged in this game, and what they aim to achieve.

---

[53]The notion of a language game, or a *Sprachspiel*, is due to Wittgenstein (Wittgenstein, 1953). It is not completely clear what Wittgenstein intended with this notion. What is clear is that, according to Wittgenstein, a language game consist of "language and the actions into which it is woven" (Wittgenstein, 1953, §7). Wittgenstein used the term 'language *game*' to "emphasize the fact that the *speaking* of language is part of an activity, or of a form of life" (Wittgenstein, 1953, §23) (emphasis in original).

3.4.1   The Players and the Aim of the Game

The game is played by a **proponent** of a substantial philosophical thesis who aims to persuade *a group of epistemic peers* of the truth of this proposition. In the previous sections, I argued that such a thesis needs to be established *beyond a reasonable doubt*, so we can think of the aim of the proponent as trying to convince all possible epistemic peers. This idealized epistemic peer group is similar to what we called the universal audience.

The above suggests that the game of giving and asking for reasons is played by a single player against a group. In actual philosophical practice, this is indeed the case. A paper in a journal or a talk at a conference is, indeed, an instance of a one versus all situation. However, since we have idealized this group as the universal audience, we can simplify our game by making a further idealization. We will idealize the group as comprising of a single opponent, which we will call the **skeptic**. This skeptic, since it is no opponent in particular, has a neutral attitude towards every proposition at the start (meaning that he neither thinks that the proposition is certainly true nor certainly false) unless rationally compelled to change his mind during the game.[54] This happens when the proponent has successfully demonstrated to the skeptic that she is entitled to that proposition.

In particular, the skeptic will only accept a proposition when the proponent establishes it *beyond a reasonable doubt*, i.e. either immediately or

---

[54]To facilitate reference, I conceive of the proponent as a "she" and the skeptic as a "he". This gendered conceptualization is random and has no relation to the roles that the players are supposed to play.

by deductive argument from premises that the proponent has already established herself as being entitled to. That is, the skeptic is not convinced of the truth of a proposition until he can no longer reasonably challenge whether the proponent is indeed entitled to it.[55]

From the above, it is clear that the main aim of the proponent is to persuade the skeptic to accept a substantial philosophical thesis. The main aim of the skeptic is to keep voicing reasonable doubts. An advantage of the way that I have set up the game of giving and asking for reasons is that, in contrast to the way (van Inwagen, 2006) conceives of philosophical practices, it has no need for a third party consisting of one or more judges.

The qualification that the philosophical thesis needs to express a *substantial* philosophical thesis is important. This criterion is meant to exclude e.g. conditional statements of the form "if you grant $\phi, \ldots, \psi$ and background logic $\mathcal{L}$, then you must grant that χ and $\phi, \ldots, \chi$ are mutually inconsistent (or lead to some other unacceptable conclusion)". I grant that hypothetical propositions of this form can sometimes be established. Although hypothetical propositions can indeed be established, they do not deliver any substantial knowledge in the sense that it does not determine which philosophical position is true. As such, we can explain how philosophical practice makes

---

[55]In the next chapter, I argue that because there is no reliable method by which a philosophical proposition can be established beyond a reasonable doubt, either immediately or through inference, no philosophical proponent is ever fully entitled to a philosophical proposition. The best she can do is show that she is permitted to believe what she believes. She cannot force a universal audience (embodied by the skeptic) to commit to believing her position as well.

epistemic progress by mapping the landscape of possible defensible positions, even though philosophers disagree about which substantial philosophical positions are true.[56]

Now that I have set out the aim of the game of giving and asking for reasons, and the two kinds of players engaged in it, let me discuss the moves that the players can make in the game in more detail.

3.4.2 The Moves of the Players

The game of giving and asking for reasons starts when the proponent puts forward a substantial philosophical thesis and claims to be entitled to it.[57] After that, the players take turns.

After the proponent has put forward a thesis, the skeptic can challenge her entitlement to this thesis by asking her why she thinks that she is entitled to this proposition. Alternatively, he can directly argue against the thesis by showing that by claiming entitlement to the thesis the proponent is also committed to a consequence that the skeptic thinks she would consider unacceptable. Such an unacceptable consequence can be a *contradiction*, or a proposition that the skeptic thinks the proponent would reject. As a response, the proponent may "bite the bullet" and decide to accept the consequence nevertheless, may offer a counter-example that challenges the inference, or she may revise the logical principles that she has accepted in

---

[56]See chapter 5 for a detailed discussion.

[57]A proponent argues against a philosophical position by putting forward a philosophical thesis for acceptance that is incompatible with the position that she wants to reject.

order to avoid the unwanted consequence.

If the skeptic asks for the reasons *why* the proponent believes to be entitled to a proposition, then the proponent must either claim that this is an immediate philosophical truth, or she must provide a *transferable* argument aimed at establishing the thesis *beyond a reasonable doubt.* If the proponent claims that the proposition holds immediately, then the skeptic can demand to know in virtue of what. A possible answer that we will investigate in this dissertation is that the proposition is claimed to be an *immediate truth.* The proponent can also provide a transferable argument on the basis of (new) reasons. In that case, she puts forward a set of propositions from which she claims that the proposition that was being challenged by the skeptic can be derived. This adds the premises of the argument to the set of propositions that the proponent has put forward for acceptance.

The skeptic can now challenge these new propositions that the proponent claims to be entitled to in a similar way as he could challenge the thesis. But he can do more. He can also challenge the argument by voicing doubts whether, even when granted the truth of the premises, the argument does indeed show that the thesis is true beyond a reasonable doubt by offering a counter-example to the rule of inference that was used. If he does that, then the proponent must either disarm the counter-example or she must offer another defense.

As the game progresses, a score-board is kept that contains three pieces

of information:[58]

- The propositions that the proponent claims to be entitled, and therefore committed, to.

- The inference rules to which the proponent commits herself.

- The consequences that the skeptic thinks those inference rules commit proponent to.

That is, while playing the game of giving and asking for reasons, the proponent is not only making explicit all of the propositions that she commits to, but also which inference rules she is using. I assume that it is clear by now why the score-board records the (purported) entitlements and (perceived) commitments of the proponent with regard to propositions. But I do want to say a bit more about the need for the proponent to commit to inference rules.

That there is an essential difference between propositions and inference rules was vividly shown by Lewis Carroll in his "What the Tortoise Said to Achilles" (Carroll, 1895), the paper from which I took the example of deductive reasoning in the previous section. The main point of that paper, as I see it, is that, if one wants to persuade an interlocutor of a conclusion that logically follows from a set of premises, one must not only get him to accept those premises, but also the inference rules needed to derive that

---

[58]Read (Lewis, 1979) for a detailed discussion of score-keeping in the game of giving and asking for reasons.

conclusion.[59]

In Carroll's paper, as discussed above, Achilles is trying to persuade the Tortoise to accept that the two sides of a triangle are equal to each other, given that the two sides of that triangle are equal to the same and that things that are equal to the same are equal to each other. The tortoise does accept the premises, but refuses the conclusion. He simply claims that he does not see how the conclusion is related to the premises. From a formal perspective, we could say that the Tortoise accepts $p \rightarrow q$ and p, but claims that he does not see how q follows from that. Achilles now asks whether the Tortoise also agrees that if the two sides of a triangle are equal to the same and that things that are equal to the same are equal to each other then the two sides of that triangle are equal to each other, i.e. $(p \rightarrow q \,\&\, p) \rightarrow q$. Again, the Tortoise agrees, but still claims that he does not see how this leads to q. Frustrated, Achilles keeps on adding conditional propositions that the Tortoise grants, the next one being $(((p \rightarrow q \,\&\, p) \rightarrow q) \,\&\, (p \rightarrow q) \,\&\, p) \rightarrow q$, and the Tortoise keeps insisting that he does not see how q follows from those propositions that he granted. The problem is, of course, that the Tortoise hasn't granted *modus ponens*, the inference rule that says that if you grant $\phi \rightarrow \psi$ and ϕ you must also grant ψ. Until the Tortoise grants this rule, he is not committed to its consequences, and Achilles can keep getting the tortoise to accept conditional propositions until the end of time without ever

---

[59]My interpretation of Carroll's paper is based on the lesson that Bertrand Russell took from it. See (Russell, 1903, p. 35).

being able to commit the tortoise to grant any of the consequences.

Similarly, until the proponent commits to an inference rule, she is not entitled to use it when she presents the reasons for her theses. And on the other side, when she does commit to a rule, she is also committed to its consequences.

It is not important that the proponent had these inference rules in mind, either explicitly or implicitly, when she claimed to have supplied a reason for a proposition that she deemed herself entitled to. As such, there is no fact of the matter which inference rule the proponent has followed, a problem that was made famous by Kripke (1982). However, when asked to make her inferential practice explicit, she must choose a rule and commit to it.[60]

This focus on making one's rules of inference explicit is a form of what Brandom (1994) calls *logical expressisvism*, the idea here being that the purpose of logic is not to *justify* our inferential practices but to make our inferential practices explicit. This ideal is grounded in Frege's dictum that

> Whatever is needed for a correct inference is fully expressed; what
>
> is not needed is . . . not. (Frege, 1879, §3)

However, unlike Brandom, I do not believe that we are guaranteed that every game of giving and asking for reasons will result in making explicit the

---

[60]This is also what I believe provides the normative force of logic. We are not committed to the rules of logic because we find ourselves in accord to their dictates when we investigate our discursive practices. Rather, we are autonomous thinkers, and as such, we put ourselves under laws of logic. I believe this idea of rational autonomy is the neo-Kantian core of Carnap's *principle of tolerance*; that one is free to adopt any form of logic that serves one's purpose (Carnap, 1934).

one and only unique set of inferential practices shared by every reasonable being (i.e. the One True Logic). In this sense, I follow Hartry Field's relativist expressivist view of logic (Field, 2009). I further develop this point in Section 4.2, where I discuss the problem of *logical pluralism.*

3.4.3 Common Ground and Bedrock

One way for the above game of giving and asking for reasons to stop is when what Stalnaker calls *common ground* has been found (Stalnaker, 2002). This happens when all the reasons that the proponent has given for her thesis and all the inference rules that she has used are considered to have been established beyond a reasonable doubt by the skeptic, and all the consequences that this commits the proponent to are deemed to be acceptable by her.

However, in order to make the audience universal, we have assumed that the proponent cannot count on any proposition or inference rule to have been established as common ground at the start of the game. This is an important difference with (other) scientific practices, where a set of basic propositions and inference rules seems to have been agreed, perhaps only tacitly, upon from the start. That is, in most contexts a mathematician can count on the fact that the axioms of ZFC and the principles of First-Order Logic are considered to be common ground, and physicists can presuppose the acceptability of the principle of the uniformity of nature. Since philosophers are thought to critically question everything, no common ground can be presupposed. This in one important difference between philosophy and other

epistemic practices that accounts for why there seems to be so little progress in philosophy.

According to Brandom, we must postulate the existence of common ground, on pain of skepticism (Brandom, 1994, p. 177). In the next chapter, I argue that we must instead accept this impalatable consequence because it is impossible for the proponent to establish common ground with a skeptic on paradigmatic philosophical theses. This is because, or so I argue, the kinds of propositions that could function as common ground, i.e. immediate truths, cannot be established beyond a reasonable doubt. Furthermore, or so I argue, there is no One True Logic that determines all the inference rules every epistemic agent must commit herself to independently of her purpose, and even if there were, at the moment, we would have no method of finding out what it is.

I do agree with Brandom, however, that

> nothing recognizable as a game of giving and asking for reasons results if justifications are not permitted to come to an end. (Brandom, 1994, p. 177)

Some propositions cannot require defense in the light of the skeptic's mere requests for justification. Otherwise, philosophical games of giving and asking for reasons would never come to an end and that would defeat the purpose of playing them.

In these cases, where the proponent has no further reasons to offer but

the skeptic also has no substantial reasons to doubt a proposition or a rule, the proponent can claim that this is, in the words of David Chalmers, bedrock (Chalmers, 2011). However, note that, according to Chalmers, "[t]he 'bedrock' card can be played only rarely. Most disputes are not plausibly bedrock disputes." (Chalmers, 2011, p. 545), and I agree.

This means that the game of giving and asking for reasons comes to a stop either when common ground between the proponent and skeptic has been established on every substantial issues under discussion, in which case the skeptic is persuaded of the thesis under discussion, or the proponent has successfully claimed bedrock on every substantial issue. In that case, she has not persuaded the skeptic, but she has ended up with a defensible philosophical position that is completely explicit in terms of its grounds, consequences, and its inferential structure. I deal with such structures in philosophy, also known as *reflective equilibria*, in detail in section 5.1.

In the next chapter, I argue first that the first of these cases will never occur for paradigmatic philosophical theses; an idealized skeptic can never be persuaded. As such, philosophical progress cannot be conceived of in terms of establishing or refuting a philosophical thesis in the face of every rational being, i.e. a universal audience. The best that a proponent of a substantial philosophical thesis can do is to develop a fully explicit defensible philosophical position. In Chapter 5, I intend to show what the epistemic aim of philosophical practice could be, if it cannot be philosophical knowledge.

Chapter 4

A Defense of Philosophical Skepticism

In previous chapters, I introduced the problem of philosophical progress and provided a general discussion on the nature of philosophical practices. In this chapter, I will make use of these preliminary discussions to present a defense of *philosophical skepticism*, the claim that substantial philosophical knowledge is impossible.

In the first section of this chapter, I discuss two systematic methods of acquiring knowledge; building a theory from first principles, and hypothesis testing. For both of these ways, two components turn out to be crucial; a set of established inference rules and a set of immediately established propositions, which we will call the *philosophical data*.

In section two, I discuss the availability in philosophy of a set of estab-

lished inference rules first by discussing the *problem of logical pluralism.* I discuss two distinct problems. Both problems start from the empirical fact that, at the moment, there are a variety of logical systems that are claimed to be the *correct* logic by their proponents. In particular, one can think here of classical logic, relevance logic (including paraconsistent logic), and intuitionistic logic. The first problem that I discuss is an *epistemological* problem. We do not *know* which, if any, of these logics is the One True Logic, and I argue that, at the moment, there is no principled way of deciding between these logics. Second, there is the *metaphysical* issue of *logical pluralism.* Some people currently claim that there might be more than one correct logic. Until these problems are solved, the Skeptic in our Game of Giving and Asking for Reasons can always raise reasonable doubt about the universal validity of the inference rules that the proponent needs to establish her case.

In section three, I move from inference rules to the nature of the philosophical data, and argue that Skeptic can always voice reasonable doubt regarding any potential immediate proposition. We can call this problem the *plausible deniability of the data.*

In the final section of this chapter, I show how the problem of logical pluralism and the plausible deniability of the data entail philosophical skepticism by defending two claims. First, it is impossible to build a secure system of philosophical truths from a foundation of first principles because no first principle is beyond a reasonable doubt. Second, it is impossible to provide a knockdown argument against a philosophical position because any

consequence can be accepted (by biting the bullet) or avoided (by changing the logic).

4.1 Methods of Scientific Knowledge Acquisition

In this section of this chapter, I discuss two systematic methods of knowledge acquisition that we see in the (other) sciences. I discuss the building of a system of knowledge from first principles, a method of knowledge acquisition that is customary in mathematics, in Section 4.1.1, and the rejection of rival theories in the hope of converging to the truth, as is the customary method of knowledge acquisition in the empirical sciences, in Section 4.1.2.

4.1.1 Building a System of Truths

In the formal sciences, practitioners are often thought to make scientific progress by (deductively) deriving new truths from propositions that were previously established as true or were postulated as axioms. As such, progress in these disciplines is conceptualized as a growing body of truths.

According to this conception of science, which Willem de Jong and Arianna Betti have described as the *Classical Model of Science* (de Jong & Betti, 2010), a science is constituted by a set of *first principles* and all of their *logical consequences.* Scientific progress, according to such a conception, consists in

finding the first principles by *regressive analysis* and deriving new knowledge from these principles by *synthesis*.[1].

 The classical model consist of seven main claims about a science S, which is conceived of as a systematic set of propositions about a specific domain which is its *subject-matter* (de Jong & Betti, 2010, p. 186):[2]

1. All propositions and all concepts (or terms) of S concern a specific set of objects or are about a certain domain of being(s).

2. There are in S a number of so-called fundamental concepts (or terms). All other concepts (or terms) occurring in S are composed of (or are definable from) these fundamental concepts (or terms).

3. There are in S a number of so-called fundamental propositions. All other propositions of S follow from or are grounded in (or are provable or demonstrable from) these fundamental propositions.

4. All propositions of S are true.

5. All propositions of S are universal and necessary in some sense or another.

---

[1]See (Beaney, 2014) for a discussion of regressive analysis.

[2]For earlier explications of a similar model of science, see e.g. (Scholz, 1930) and (Beth, 1965). The explication in (de Jong & Betti, 2010) is essentially different from these earlier accounts by making the important distinction between the true order of a justification and the order in which the links of justification are discovered (or presented), and by emphasizing that the non-fundamental propositions need to be *grounded* in the immediate propositions, a requirement that is generally taken to be stronger than merely being a logical consequence of the immediate propositions.

6. All propositions of S are known to be true. A non-fundamental proposition is known to be true through its proof in S.

7. All concepts or terms of S are adequately known. A non-fundamental concept is adequately known through its composition (or definition).

When we think of progress in mathematics, we often think of progress in this sense. For instance, the standard story of mathematical progress goes something like this: after Thales proved the first theorem, and thereby became the founder of the science of mathematics, the Greeks proved many mathematical truths, mostly truths from geometry such as the *Pythagorean Theorem*. During the modern times, more branches of mathematics were discovered, and truths from number theory, algebra, calculus, etc. were added to the store of mathematical knowledge. Today, mathematicians continue to add to this store of knowledge by finding new branches of mathematics, such as topology and category theory, and by adding new results to the established and new branches of this science.[3]

Historically speaking, philosophy was also conceived by many philosophers, at least ideally speaking, to constitute a science in precisely this sense, and we can find many remains of this conception in discussions surrounding philosophical methodology (de Jong & Betti, 2010, p. 186). This is, for instance, the conception of science that we find in Aristotle's *Analytica Posteriora.* Other examples include the *more geometrico* that we find in Descartes'

---

[3]See e.g. (Merzbach & Boyer, 2011) for a more detailed and more scholarly presentation of the history of mathematics.

Meditations, the Cartesian *Logic of Port-Royal*, and Spinoza's *Ethics*. It is also especially clear in the 18th- and 19th-century German metaphysical tradition, where metaphysicians like Christian Wolff presented their philosophy as a deductive system from *first principles*. Another good example is the conceptualization of science by Bernard Bolzano that we find in his *Wissenschaftslehre*.

One core aspect of this conception is that there are two distinct ways of knowing a proposition; some propositions can be immediately known, others only through logical inference. It is obvious that the truth of the first principles cannot be known by inference. Usually, these immediate truths are said to be *self-evident* or *intuitively known* (de Jong & Betti, 2010, p. 193) With regard to the inferred truths, some philosophers have added the demand that the deductions used to prove new results should, at least ideally speaking, not only establish the result but should also be *explanatory* (de Jong & Betti, 2010, p. 190). I will side-step this issue for now, but I will note that this is one motivation for the investigation of whether there exist explanatory proofs in mathematics.[4] It does, however, seem clear in the philosophy of mathematics that non-explanatory proofs, although perhaps deficient in some sense, are sufficient to constitute scientific progress. The computer assisted proof by cases of the four-color theorem, for instance, might not explain why every planar map can be colored using at most four colors, but it was an absolute landmark result (Appel & Haken, 1989).[5]

---

[4]See (Mancosu, 2011) for an introduction to the notion of mathematical explanation.

[5]Let us, for now, also side-step the interesting issues surrounding the epistemic status

Note for now that a practice cannot make scientific progress in this sense unless the set of immediate propositions is *known* to be true and it is determinate when a proposition is grounded in others. Later in this chapter, I argue that neither of these assumptions is satisfied in the case of philosophy. A purported intuition can always be plausibly denied, and logical pluralism makes it indeterminate what the logical consequences of the foundations are.

Of course, this conception of what constitutes a science is not the only one on the market. It looks like a good notion of science for formal sciences like mathematics and logic, but not all scientific progress can be thought of as either finding first principles or as deriving their consequences. The main reason for this is that such a conception cannot account for progress in the sense of *scientific revolution.* To account for this kind of progress, we need to turn to a different notion of scientific progress.

4.1.2 Theory Building and Testing

In the previous subsection, I discussed scientific progress according to the *Classical Model of Science.* We concluded that discussion with the observation that the Classical Model fits well with progress in the formal sciences, like mathematics, but could not account for progress in the sense of *scientific revolution* which is prevalent in the empirical sciences. In this subsection, we'll look at another conception of scientific practice that does take account

---

of computer assisted proofs. Although considered to be controversial a few decades ago because these proofs could not practically be verified by a human, computer assisted mathematics seems to have become more acceptable nowadays.

of the possibility of scientific revolution; progress through testing. The hard empirical sciences, like physics, chemistry and biology are often thought of making scientific progress in this sense.[6]

Unlike the previous conception, where knowledge of the system is mediated by the knowledge of the immediate propositions and the inference rules, here the idea is that, as theories are tested and get confirmed or rejected by evidence, the modified theories we develop to take account of the results (slowly) converge to the truth. During periods of what Kuhn calls *normal science* (Kuhn, 1962), such progress is slow and incremental, but when one theory replaces another during a scientific revolution, such progress occurs in leaps.

Of course, in this conception, the truth that we aim at is only a *regulative ideal.* We need not ever reach such a true theory for a scientific practice to be progressive in this sense, nor do we need to believe that we could even reach such a true theory in principle. It is, however, crucial for this conception of scientific progress to believe that our scientific practices bring us closer to the ideal step by step.

Let us, for concreteness, look at the development of physics as a paradig-

---

[6]In this section, I could have also discuss Nagelian reduction, i.e. scientific progress in the sense that one theory is reduced to a more encompassing one. The standard example here is the reduction of thermodynamics into statistical mechanics. However, I won't go into this conception of progress in detail here, because such a conception of progress is only applicable when the to be reduced theory and the reducing theory can be formalized as deductive systems in the sense discussed in the previous section (see (Losee, 2003, loc. 670) for a detailed discussion). As such, when progress in philosophy is conceived of in this way, the problems of the *plausible deniability of the data* and the *indeterminacy of consequence* hold as well.

matic case of scientific progress in this sense. The standard story of scientific progress in physics goes something like this: after being established during the classical period, Aristotelian physics was dominant until it was challenged in early modernity by scientists who started to make observations that were in conflict with the Aristotelian world view. A famous example of this is Galileo's observation of mountains on the moon after designing his own telescope, an observation that was in conflict with the Aristotelian idea that the perfect heavenly bodies where qualitatively different from the imperfect Earth. At roughly the same time, the extremely accurate and comprehensive observations by the Danish astronomer Tycho Brahe led to the heliocentric theorizing by Copernicus and Kepler, a model that now slowly replaced the Ptolemaic system of concentric circles which considered the earth to be the center of the universe. The crowning achievement of this period was, of course, Isaac Newton's development of a theory of mechanics that accounted for bodies both in the heavens and on the earth, a development which definitively destroyed the Ancient idea that the heavens and the earth followed different rules. Newtonian mechanics was a giant leap forwards, but it turned out to be not quite correct. New phenomena were discovered that were not in accord with Newtonian mechanics, in particular, phenomena within the theory of electromagnetism. This gave rise to a second revolution, led by Albert Einstein, who developed the Special and General Theory of Relativity, and who played a fundamental part in the development of Quantum Mechanics.

At the moment, both Relativity Theory and Quantum Mechanics are the

best confirmed theories we have and they have withstood the most stringent of experimental tests. Especially the confirmation of the General Theory of Relativity has been in the news lately due to its prediction of the gravitational waves that scientists of the LIGO Scientific Collaboration (LSC) announced to have directly observed on February 11, 2016.[7] However, no matter how successful Relativity Theory and Quantum Mechanics seem to be, physicists do expect another scientific revolution when a theory is developed that unifies the force of gravity with the weak, strong, and electromagnetic force into a single framework. It has proven difficult to find such a theory, and it is expected that such a theory, once found, will lead to a fundamental change in our conception of the nature of the universe.[8]

In the above story, we clearly see the notion of progress under discussion at work; Aristotelian physics was replaced by Newtonian physics, which in turn was replaced by the even better theories of Relativity and Quantum Mechanics. But although these two theories are much better than any of the previous theories we had, we strongly suspect that they cannot be completely right, and we are searching to replace them with an even better theory.

Some contemporary philosophers also suppose that philosophy *ought* to make progress in this sense (see for instance (Sider, 2009, p. 385), (Love, 2013), and (Chalmers, 2015, p. 4)). According to such a view, a philosophical thesis is a hypothesis which should be tested against some kind of relevant

---

[7]See https://www.lsc-group.phys.uwm.edu/ppcomm/Papers.html for a large collection of papers from LSC members.

[8]For a more detailed introduction to the history of physics, see e.g. (Heilbron, 2015).

philosophical data, usually thought to be philosophical intuition. The ideal here would be to find a philosophical theory that confirms as much of the philosophical data as possible, while not being definitively refuted by any of it. In the words of John Pollock and Joseph Cruz:

> [O]ur basic data concerns what inferences we would or would not be permitted to make under various circumstances, real or imaginary. This data concerns individual cases and our task as epistemologists is to construct a general theory that accommodates it. (Pollock & Cruz, 1999, p. 5)

An important method of inference for this account is called inference to the best explanation (often known as IBE).[9] That is, a theory is developed to explain the data at hand, and is subsequently tested, i.e. confirmed or refuted, by new data.

Lately, some philosophers have been claiming that IBE is an essential philosophical tool. That this is also what philosophers should be doing was, for instance, recently defended by Andow (forthcoming) and Williamson (2016). Other examples, are Hawley (2006), who argues that IBE can justify beliefs in metaphysics no less than in science, Rayo (2013), who argues that IBE can justify beliefs about identity ("just is" statements), Biggs (2011), who argues that IBE can justify beliefs about necessity and possibility, and

---

[9]In this dissertation, I will use the terms "inference to the best explanation" and "abduction" interchangeably. For my purposes a distinction is not important, but note that there are good reasons to distinguish *Peircean* abduction from inference to the best explanation. See, e.g., (Mackonis, 2013) for details.

Sider (2009), who is explicit that metaphysical investigations generally appeal to abductive principles.

Let us, for concreteness, look at a quote from a paper published in 2017 by Stephen Biggs and Jessica Wilson, who express great enthusiasm for the use of IBE in philosophy:

> [A]bduction has clear advantages over other modes of a priori inference, involving intuition and conceiving, that have been traditionally appealed to as justifying philosophical claims that, one way or another, transcend the empirical. In particular, abduction is a distinctively rich mode of inference, with ampliative, fine-grained resources enabling it to bridge conceptual gaps and substantively legislate between competing philosophical views. Abduction is not just useful for scientific or everyday inference; on the contrary, it is the ultimate philosophical tool—and one that, we have here aimed to establish, can be largely employed from the armchair. (Biggs & Wilson, 2017, p. 756)

In the next section, we'll take a closer look at whether we are justified in thinking that philosophy progresses because subsequent philosophical theories converge to the truth. Let us just note a problem for now which we discuss in greater detail in the next section.

One problem with the current conception of scientific progress that philosophers of science have had to grapple with is that, since we do not have access

to the truth, we cannot actually measure whether we're indeed approaching the 'true' system, even supposing that we have decided on a single metric.[10] Philosophers of science therefore use indirect pragmatic measures to argue that the empirical sciences, like physics, have made progress in this sense. In particular, they often refer to the near *universal consensus* of scientists about scientific claims, and the *predictive power* of a theory. That predictive power is a sign of convergence to the truth is related to what is known as the *No Miracles Argument.*

I have already argued in Chapter 2, that we do not see anything resembling (near) universal consensus in philosophy, nor is there any evidence that the opinions of philosophers are converging towards such a consensus. Similarly, I have pointed out that the *No Miracles Argument* is not available to philosophers because philosophical theories generally do not allow for the prediction and control of empirical phenomena.

However, in the next sections, I aim to show that there is another problem. This problem is related to the crucial role that *inference rules* and *data* play in this conception of progress through testing. In particular, I show why the data that is used in philosophy cannot play the crucial role needed to be able to confirm or refute philosophical theories.

---

[10]People differ in view of how this notion of closeness to the truth is to be explicated even in the ideal case, i.e. when we are allowed to postulate a determinate true theory from which we can measure. I will not go into a detailed discussion here, but it is closely related to the pragmatic considerations concerning theory choice that I discuss in section 5.3.

4.2   Logical Pluralism: the Indeterminacy of Logical Consequences

In the previous section, I discussed two notions of scientific progress: system building on a solid foundation and convergence on truth through hypothesis testing. We saw that both notions of progress require there to be a set of immediate truths that function as either the foundation of the system or the testing bed of the theory (although these immediate truths may, of course, be theory-laden), and a determinate theory about what the logical consequences of a set of propositions are. In the first case, we need to know what other propositions follow from the foundational base, in the other, we need to know which propositions confirm and refute the hypotheses under investigation.

In the next section, I argue that there are no immediately known propositions in philosophy. In this section, I discuss how the empirical fact of *logical pluralism* problematizes our notion of philosophical progress.

4.2.1   Logical Pluralism

Logical Pluralism is the philosophical claim that there is more than one correct logic.[11]

It turns out to be difficult to say what logic is. According to some, logic is about the most general truths, i.e. the structure that all possible worls

---

[11]See (Russell, 2013) for an introduction to Logical Pluralism. See (Beall and Restall, 2006) for a book length treatment and defence of the claim.

share. We can call this the metaphysical conception of logic. According to others, logic is about the laws of reason, i.e. about the rules of how we should think. Perhaps we could call this the psychological conception of logic.[12] A third school thinks that logic is about specific relationships within (a) language. One could call this the linguistic conception of logic. But whatever conception of logic you happen to like, all accounts agree that the notion of *logical consequence* is one of logic's most central concepts. One way to argue for logical pluralism is therefore to claim that there is more than one way to explicate the notion of logical consequence (Beall and Restall, 2006).[13] This is also how we will interpret the claim in this dissertation. We will say that we have two significantly different notions of logical consequence when these notions do not agree on whether some ϕ is a logical consequence of some (possibly empty) set of propositions Γ.

In the past century, we have seen that the notion of logical consequence

---

[12]This is a normative view of logic. Ever since Frege's attack on psychologism, it is customary to strictly distinguish the descriptive laws of how people do reason and the normative rules of how people should be reasoning (Kusch, 1995). The first kind of laws belong to the psychology of reasoning, only the second kind to logic. Furthermore, it is important to distinguish the psychological conception of logic from the psychologistic conception of logic. The psychological conception of logic only says that the subject-matter of logic is about (ideal) reasoning. It says nothing about how the laws of logic are justified. The psychologistic conception of logic, by contrast, does claim that the laws of logic are justified by how empirical or ideal agents do actually reason.

[13]As with anything in philosophy, there are other possible ways to explicate logical pluralism. One possibility is to claim that there is more than one way to explicate the notion of *logical constant*. This is, for instance, defended by (Varzi, 2002). Even Williamson, who is a logical monist defending classical logic, believes that "the choice of logical constants is pragmatic. Varying the extension of 'logical constant' amounts to varying what one is investigating the general structural features of." (Williamson, 2015, p. 3) So although Williamson believes that classical logic is the correct logic in all domains, he does believe that the choice of what the logical constants are can vary across domains.

can indeed be explicated in various ways, although not all differences between the explications can serve as grounds for an interesting form of logical pluralism. One way in which the consequence relation can be differently explicated is by raising the question what kind of things the relation relates. Usually, we take logical consequence to relate a (possibly empty) set of propositions, which we usually call the set of premises, and a single proposition, which we usually call the conclusion. However, some believe that it is more fruitful to think of logical consequence as a relationships between a set of propositions, the set of premises, and another set of propositions, the set of conclusions (Restall, 2005). A third option is that logical consequence relates a set of sentences to a sentence (or a set of sentences). Although the observation that we can have different relations of logical consequence depending on what kind of things we take the relata to be is usually not believed to lead to a very exciting form of logical pluralism, and I agree with this assessment, it might turn out after full explication that these different consequence relations do actually differ from each other in interesting ways, as was argued by Gillian Russell (2008). One example of an interesting difference that Russell gives is that, supposing that logical consequence is reflexive under every explication, the proposition that $a = b$ is a consequence of the proposition that $a = a$ (since, for instance, it is impossible for Venus not to be Venus), although the sentence '$a = b$' might not be a logical consequence of the sentence '$a = a$' (since, for instance, the words 'Hesperus' and 'Phosphorus' cannot be interchanged in every context). Should this be the case, we would get a logical

pluralism, since $a = b$ would be a logical consequence of $a = a$ in one sense, but not in another. Although these are all fascinating questions to explore, in this dissertation, I will simply follow the usual convention that the consequence relation holds between a set of propositions, which I will generally designate by Γ, and a single proposition, which I will generally designate by ϕ.

As briefly discussed earlier in this dissertation, there are actually already two viewpoints under which the notion of logical consequence can be understood in any classical theory of logic: in terms of *validity* and in terms of *derivability.*

In terms of *validity*, we can say that ϕ is a logical consequence of Γ if it is impossible for ϕ to be false when every proposition in Γ are true. In symbols, this relationship is usually written as

$$\Gamma \models \phi$$

Although one need not necessarily interpret logical validity in a metaphysical manner, for instance by insisting that the semantics of a logic is about mathematical models and not about possible worlds, this account fits well with a metaphysical view of logic due to its use of the modal notions of *possibility, impossibility,* and *necessity.*

In terms of the notion of *derivability*, we can say that a proposition ϕ is a logical consequence of a set of propositions Γ if we can derive ϕ from Γ using

a set of logical derivation rules. In symbols, this is usually written as

$$\Gamma \vdash \phi$$

This conception of consequence fits well with the psychological conception of logic, since it is natural to interpret the derivation rules as steps in a reasoning process.

Usually, these two different viewpoints do not lead to a logical pluralism, because for first-order classical logic (as well as for a number of other prominent logical systems) it is the case that

$$\Gamma \models \phi \Leftrightarrow \Gamma \vdash \phi$$

This is called the *Completeness Theorem*, and was first proved by Kurt Gödel (Gödel, 1930). In essence, this theorem states that, if it is impossible for ϕ to be false while every proposition of Γ is true, then we can derive ϕ from Γ using the derivation rules of the system, and vice versa. This means that the possibility of explicating the notion of logical consequence in terms of *validity* and in terms of *derivability* does not lead to logical pluralism in our sense, since these different explications never differ in their judgement about whether a proposition ϕ is a logical consequence of a (possibly empty) set of premisses Γ.

There is a small caveat here. The *Completeness Theorem* does not hold for any logical theory that is expressive enough to express arithmetical truths.

This is Gödel's famous *Incompleteness Theorem.* Gödel proved that for these more expressive logics, either the derivation rules allow for the derivation of two inconsistent propositions, or there are valid consequences that cannot be derived by the derivation rules. Some take this result to mean that First-Order Classical Logic, the most expressive logical system for which the completeness theorem still holds, is the real subject-matter of logic, but I will stay neutral on this matter here.

There is, however, the question of primacy. In terms of justification, which of these two relations that I have discussed above come first? Usually, it is believed that the notion of *validity* is primary, because the derivation rules are usually justified by pointing out that they never lead to an invalid inference. We will follow this judgment in this dissertation, and will therefore only focus on *validity* in the rest of our discussion of logical pluralism. It is important to note, however, that there are people known as *inferentialists*, who think that it is derivation which is the primary concept (or at least an independent notion), and that the deductive systems can be justified in terms of, among others, the notion of *harmony*.[14] Although I find this is a highly interesting research project that I follow with care, up to date (as far as I know), no one has been able to fully justify a logic through this method yet.

---

[14]For an introductory discussion of this issue from the two sides, see (Shapiro, 2005) and (Prawitz, 2005).

4.2.2   The Generalised Tarski Thesis

After these preliminary remarks, let us now turn to logical pluralism proper. Remember that one form of logical pluralism states that there is more than one way to explicate the notion of logical consequence. Logical pluralism in this sense has recently been defended by J.C. Beall and Greg Restall in (Beall and Restall, 2006). Another defense of logical pluralism was recently provided by Stewart Shapiro (Shapiro, 2014). Famous opponents of logical pluralism are Stephen Read, Timothy Williamson and Graham Priest.

Given, for the purposes of this dissertation, that *logical consequence* should be explicated in terms of *validity*, the question is immediately raised how to determine what it means for a proposition ϕ to be necessarily true when all the propositions in Γ are true.

A natural explication of this modal understanding of *logical consequence* is to follow Beall and Restall, who cash out their logical pluralism using a model theoretic explication schema, which they call the Generalised Tarski Thesis. That is,

$$\Gamma \models \phi \Leftrightarrow \phi \quad \text{holds in every case where} \quad \Gamma \quad \text{holds}$$

Logical pluralism arises from giving the notion of 'case' different interpretations. Not any explication of what a case is counts as a correct explication. Beall and Restall do restrict what a consequence relation could look like. According to them, any relation must hold necessarily, be normative, and be

formal.

It is important to distinguish two kinds of logical pluralism. On the one hand, there is what we could call descriptive logical pluralism; the claim that there exists more than one logic. In a sense, this is not a very interesting claim, since it is easy to show. There are many different formal logical systems that disagree on what the logical consequences of a set of propositions are. For instance, there is classical logic, relevance logic, and intuitionist logic.

However, the more interesting question is which of these logic is the *correct* logic. It turns out to be hard to say what it means for a logic to be correct. In the literature, we often see that it is unclear whether a correct logic has to reflect the relations between propositions in a metaphysical sense, or whether it should capture the logical intuitions of language users (and if so, whether it should capture it for all (possible) language users). Every logic has its defenders. Classical logic, for instance, is defended by Williamson, and Stephen Read is a staunch defender of relevance logic. Intuitionistic logic is not so popular anymore, although it used to be defended by e.g. Michael Dummett.

Importantly, I follow (Hjortland, 2017), (Williamson, 2015), and (Priest, 2006) in the view that logic is not exceptional, and that it cannot be justified by immediate self-evidence. Rather, whatever logic we end up choosing, we need to give reasons for it. In philosophy, and elsewhere, we cannot simply assume that logic is given.

According to these authors, the best way to conceptualize this is to see

that the choice of logic depends on inference to the best explanation. One way to see that logic can't simply be grounded in our grasping the One True Logic *a priori*, is to recognize that logical intuitions are just as problematic as any other kind of intuition (see (Counihan, 2008) for a detailed study of logical intuitions in different cultures). That is, we choose logic because it allows us to derive the propositions that we wish to accept and have no (or few) consequences that we want to reject. For those who want to see this method at work, a good example of a philosopher investigating the consequences of different logics with respect to certain philosophical topics is Williamson's discussion of why classical logic is the right logical framework to think about *vagueness* (and anything else, for that matter) and not, e.g. any of the many-valued logics or any kind of supervaluationist theory.[15]

Of course, like any philosophical discussion, Williamson's defense of classical logic has not settled the debate once and for all. Priest (Priest, 2006), for instance, thinks that it is a paraconsistent logic that explains best how we ought to reason. (Hjortland, 2017), in turn, defends some kind of logical pluralism.

It is interesting to see that IBE actually played an important role in how classical logic got accepted in the first place. Russell's acceptance of the Theory of Descriptions, for instance, can already be seen as using the abductive method. In his classical 1905 paper, Russell weighs his new theory

---

[15]Williamson discusses these issues in many different places, among which (Williamson, 1994) and (Williamson, 2007).

against the theory he himself held earlier (but which he attributes to Frege) and the theory of Meinong by seeing how they fare against certain puzzles. He does something similar when he argues in favor of his "classical" conception of logic versus Boole's algebraic logic. And the logic used in Principia Mathematica is explicitly "inductively justified", to use Whitehead & Russell's terms, because they derive the right truths of mathematics which are immediately known, such as $1 + 1 = 2$.

Like Whitehead and Russell, Williamson appeals to the fact that classical logic is able to derive the classical mathematical truths. Alternative logics, like intuitionistic logic, cannot. Williamson takes this to be in great favor of classical logic.

> Once we assess logics abductively, it is obvious that classical logic has a head start on its rivals, none of which can match its combination of simplicity and strength. [...] The case may indeed be strengthened by reference to the track record of classical logic: it has been tested far more severely than any other logic in the history of science, most notably in the history of mathematics, and has withstood the test remarkably well. Nevertheless, the initial case for classical logic would be quite powerful, even if we had only stumbled across that logic a few weeks ago. (Williamson, 2015, p. 19)

However, as (Hjortland, 2017) argues, classical logic is not the only possi-

ble background logic for mathematics. First of all, mathematicians are hardly ever fully explicit about the logical structure of their proofs. Mathematical proofs could be compatible with a whole range of different logics. Different frameworks are being investigated by Priest, and also Shapiro. Second, even if classical logic was the best logic in the domain of mathematics, that does not mean that it fares equally well in any other domain. For instance, it might fare badly in the domain of quantum mechanics.

Penelope Maddy defends something similar in her *The Logical Must* (Maddy, 2014). Many phenomena are classical, which gives us the idea that classical logic is universally applicable. However, some phenomena, like Quantum Mechanics (and perhaps others like Truth and Vagueness) might not be. We need to investigate and choose a suitable logic for each and every domain.

One idea to save the notion of a One True Logic, is to insist that it contains those principles that hold across all domains. However, that will lead to the worry that this One True Logic will be extremely weak;

> That logic will likely turn out to be exceedingly weak, maybe even empty. We can call it the One True Logic if we want, but it will have a significant shortcoming. Unlike the pluralist theory it cannot explain the success of classical logic in mathematics. (Hjortland, 2017, p. 25)

Logical pluralism, as described above, has recently given rise to the investigation of *logical nihilism*. The best worked out investigation is by Gillian

Russell, although not in print, but we find Shapiro investigate it as well in his *Varieties of Logic* (Shapiro, 2014). The main thesis of *logical nihilism* is that there are no *logical truths*. That is, there are no truths in virtue of logic.

Note that we can distinguish between two problems here: an epistemological problem and a metaphysical problem. The epistemological problem is that, at the moment, we do not know what the One True Logic is. The metaphysical problem is that there might not be a single correct logic. If this is true, then logic is currently undergoing the same development as geometry did in the 19th century. It used to be believed that there was only one correct geometry, but now we know that there are infinitely many geometries that can (or cannot) be fruitfully used to describe the space of a system that we are interested in. Often, there is no one unique geometry that *must* be chosen, although some are more convenient to use than others.

4.3 Intuitions: Philosophical Evidence

In the previous section, I discussed the problem of logical pluralism. In this section, I discuss the philosophical data that could potentially function as evidence. Both accounts of scientific progress that I discussed at the start of this chapter needed some kind of immediate data, i.e. non-inferential propositions, whether we want to derive truths from a certain foundation, or wish to confirm or refute a theory on the basis of some kind of given data.

Philosophers have, of course, discussed the necessity of the immediate

in the acquisition of knowledge. Here is a particularly clear example from Locke:

> This part of knowledge is irresistible, and, like bright sunshine, forces itself immediately to be perceived, as soon as ever the mind turns its view that way; and leaves no room for hesitation, doubt, or examination, but the mind is presently filled with the clear light of it. 'Tis on this intuition, that depends all the certainty and evidence of all our knowledge. (Locke 1690: 264)

In the formal sciences, it is probably best to think of the immediate propositions as being either postulated (e.g.(Potter, 2004)) or assumed (this is usually called *if-thenism*), although some philosophers think that the mathematical axioms and derivation rules are justified by IBE from the assumption that there are self-evident mathematical truths, such as the elementary truths of arithmetic.[16]

In the empirical sciences it is clear that some kind of empirical observation functions as the immediate knowledge base. It's true that there are philosophical issues with the theory-ladenness of observation. It's also true that we can be mistaken about what we observe. We are, for instance, all susceptible to perceptual illusions like the Müller-Lyer illusion (Müller-Lyer, 1889). However, since I simply assume that the (other) sciences make

---

[16](Whitehead & Russell, 1910), for instance, defended this view. Empirical research has indeed shown that humans are capable of immediately grasping some mathematical propositions and relations through what is known as the *number sense* (Dehaene, 2011). In the terminology that I explain below, the number sense consists of a reliable maturational Type 1 cognitive process.

progress, discussing this goes beyond the scope of this dissertation. I do want to remark, however, that these problems can probably be dealt with since it is not at all clear that the observational data that is being used must be theory-laden with the theory that is being tested.[17]. Furthermore, we know many things about the reliability of scientific measurements and have robust accounts on standards of observational error. As I discuss below, we have no such accounts for the reliability of immediate philosophical evidence.

So what could play the role of the immediate in philosophy? In my general discussion of philosophy I have already pointed out the non-empirical nature of philosophy, so, normally speaking, empirical observation won't do.[18] So what kind of *a priori* propositions are available to a philosopher? Usually, it is believed that philosophers make use of philosophical *intuitions* as their data.

In the next section, I argue that philosophical intuitions cannot play the role of the immediate that is needed in the accounts of progress discussed earlier. The problem is, as was observed by David Chalmers (2015), the possibility of *premise denial.* In each of these cases, no matter what is proposed by a philosopher as an immediate truth, another philosopher can rationally

---

[17]See (Kuipers, 2001, §2.3) for a detailed discussion of this issue.

[18]Of course, sometimes the immediate propositions of some parts of philosophy could be empirically informed. What we see happening in the philosophy of cognitive science seems to be an especially good example of this. In the next chapter, I argue that when philosophers do indeed manage to connect their theories up to empirical observations or mathematical theorems, those theories could become part of that branch of science, or could perhaps start to develop as a science of its own. This is called philosophical midwifery.

deny it.

So, what are philosophical intuitions? A first attempt to define what philosophical intuitions are, is to say that they are propositions that *seem* to be true in an immediate sense. A common example of such a philosophical intuition is that *torturing a sentient being for fun is wrong.* We seem to need no argument to be willing to assent to such a proposition, and even if we cannot give reasons *why* it is wrong to torture sentient beings for fun, many of us are willing to claim that we *know* that it is wrong.

The above characterization is, of course, too imprecise for philosophical theorizing about the nature of intuitions. What does it exactly mean for a proposition to *seem* true? Furthermore, it is not sufficiently restrictive. It seems true to me that the house that I grew up in had a dark wooden floor, because that is how I remember it. However, for normal theoretical purposes, intuitions should be distinguished from memories. Similarly, philosophical intuitions should be distinguished from perceptions.

It has proven hard to define a generally acceptable definition of the term 'intuition'. Well-known definitions state that intuitions are propositional attitudes with a distinctive phenomenology (Bealer, 1998), dispositions to assent (Sosa, 1998), or theory-laden judgments that occur in the absence of conscious reflection (Devitt, 2006). Each of these definitions have come under attack in the literature. Bealer, for instance, insists against dispositional analyses that intuitions must be occurent (Bealer, 1998). Even the *prima facie* obvious idea that intuitions must, at the least, be beliefs has been

undercut. For instance, it still seems immediately true to me that every characterization of a set defines a set (if only the empty set) but I do not believe this because I know that it leads to Russell's paradox.

The question what intuitions could be has only become more unclear now that the nature of philosophical intuitions has come under extensive philosophical scrutiny. The above difficulties have led Williamson to claim that "philosophers might be better off not using the word 'intuition' and its cognates. Their main current function is not to answer questions about the nature of the evidence but to fudge them, by appearing to provide answers without really doing so" (Williamson, 2007, p. 220). Williamson prefers to talk of certain propositions being 'counterintuitive', i.e. *prima facie* unacceptable.

Although I agree with Williamson that the use of the term 'intuition' is fraught with difficulties, I do think that the notion of *intuition* can be used fruitfully. I think it is perhaps best to think of the concept of *intuition* as a *family resemblance* concept, following C.S.I. Jenkins who, in a recent paper, identified four core family traits of the notion of intuition; *immediacy, a prioricity, methodological starting point* and *common sensicality* (Jenkins, 2014). For my purposes here, it suffices to focus on only two of these family traits: immediacy and a prioricity. Whatever intuitions exactly are, if they are going to play the evidential role in philosophy that is needed, they must be immediate and a priori.

Before I deal with the evidential status of intuitions in philosophy and

the argument from Experimental Philosophy that aims to undermine their reliability, I first want to discuss some philosophers who deny that intuitions (should) have any evidential status in philosophy at all. Jennifer Nado calls these philosophers *intuition deniers* (Nado, 2016). Those are, most notably, (Williamson, 2007), (Cappelen, 2012), and (Deutsch, 2015). These intuition deniers point to several problems with the idea that intuitions (should) provide evidence in philosophical theorizing. I have put the word 'should' between parentheses, because intuition deniers differ slightly with respect to what they exactly claim. While (Williamson, 2007) argues that philosophers often think that they use psychological data as evidence, but that they should not think of themselves as doing that, the main claim of (Cappelen, 2012) is that philosophers do not even conceive of themselves as using intuitions in an evidential way.

(Cappelen, 2012) has argued that, although intuitions are discovered through an immediate Type 1 psychological process (see below), they are always justified by argument afterwards, i.e. by a Type 2 process. In essence, this means that Cappelen rejects the view that philosophical intuitions play the role of immediate propositions that I identified when discussing notions of progress. I can, of course, grant Cappelen that intuitions could be used as a method of discovery in philosophy. My main point in this section is that intuitions cannot play the role they are usually believed to play in immediate justification, a role which Cappelen claims they factually do not play since he believes that the justification of the intuitions themselves are always

mediated through argument.

A more challenging problem is posed by Williamson's form of intuition denial. Many philosophers take themselves as appealing to a *psychological* fact, an intuition, when they appeal to something immediate. (Williamson, 2007) has argued that this is the wrong way of looking at things. As Jonathan Ichikawa has pointed out, the term "the intuition that p" is ambiguous. It can either mean being in a certain psychological state, or it can refer to the content of that state (Ichikawa, 2014, p. 234). Now, given this ambiguity, Williamson points out that we should not psychologize our data. According to Williamson, philosophers do not appeal to psychological entities when they use philosophical intuitions, but to the immediate facts that are their content. For instance, we do not appeal to our *judgment* that the subjects have no knowledge in Gettier-cases, according to Williamson, but to the fact that they do not have any such knowledge.[19] We can appreciate this point when we realize that, when we give reasons in the game of giving and asking for reasons, we appeal to a proposition, and not to a seeming or some other such psychological entity.

Let us grant Williamson this point. However, as (Nado, 2016) has pointed out, *intuition deniers* like Williamson and Deutsch still owe us a story about why people are willing to assent to such immediate appeals to facts. How do they know that it is indeed a fact? And what can they tell the Skeptic

---

[19]Ichikawa himself follows Williamson, and also holds this view. A similar view is defended by (Deutsch, 2009) and (Deutsch, 2010).

when he demands a reason why he should accept these propositions as facts? Here the story seems to be that the willingness to immediately assent to such propositions are the product of some kind of reliable psychological process. So, although the immediate propositions themselves that the philosopher needs to appeal to are not best conceived of as a psychological entities, the willingness to immediately assent to them most certainly is a psychological phenomenon. Let us therefore turn to psychology and see if we can find help in our attempt to give an explication of the notion of philosophical intuition.

In order to pin down a fruitful notion of philosophical intuition, it is helpful to look at its use in psychology. Although philosophers do not always use the term 'intuition' in the same way as psychologists do, there seems to be a large semantic overlap between the two uses (see Hodgkinson et al. (2008)). In both disciplines, intuitions are regarded as assessments that come about without explicit reasoning, i.e. judgments that are immediate, and they seem to have some kind of *prima facie* credibility to those who hold them.

In this dissertation, I follow De Cruz (2015) in taking philosophical intuitions to be the result of Type 1 cognitive processes. Psychologists often think of intuitions in terms of *dual process theory*. According to dual process theory, first proposed by (Evans, 1984), and based on the traditional philosophical distinction between associative thinking and true reasoning, human cognition consists of two kinds of processes called Type 1 and Type 2 processes. Type 1 processes are typically fast, automatic, fluent, and effortless. Most importantly, the products of Type 1 processes are immediate in the

sense that they emerge without explicit reasoning. Type 2 processes are typically slower and require mental effort. These processes are often identified with reasoning. A deduction, such as a mathematical proof or a philosophical argument, is a typical product of a Type 2 process.

Here, a word of caution is in place. Some psychologists believe that these processes are products of different systems or modules, often called System 1 and System 2, but we need not make this assumption. There is no need to reify these two kinds of processes as consisting of two isolated systems (see also (Kahneman, 2011, p. 28-29)). Also, we need not think that there is only a single cognitive faculty that is responsible for all intuitions. Intuitions could be the product of any number of cognitive processes.

Furthermore, psychologists such as (McCauley, 2011) distinguish between two kinds of Type 1 processes: maturational Type 1 processes and practiced Type 1 processes. All typically developing children acquire maturational Type 1 processes early in their development, usually during infancy, and these processes are mastered more or less spontaneously, without extensive cultural support. Examples include learning to speak a language and developing a theory of mind. In contrast, practiced Type 1 processes are not acquired spontaneously, but develop only under certain (cultural) conditions, typically some kind of training or schooling. Examples include the ability to read and write, the ability to drive a car, and the ability to play a musical instrument.

The important difference between these two types of Type 1 processes is that, because one type develops spontaneously during typical child devel-

opment while the other does not, one can expect a low variability in the proficiency of maturational processes within any population, while practiced processes will show a larger range of variability because not everyone has practiced this capacity to the same extent.

The results from Experimental Philosophy that I discuss below show that there is a large variability concerning philosophical intuitions. As (Feltz et al., 2016) argue, there is simply no unique folk intuition about any philosophical question. But, most importantly, it seems that philosophers have significantly different intuitions than the general population concerning philosophical matters (Weinberg et al., 2001). This variability has led several philosophers like (Pritchard, 2014), to argue that there is indeed an essential difference between naïve philosophical intuition and expert philosophical intuition. This variability is exactly what one would expect if philosophical intuitions are the product of practiced Type 1 psychological processes. Note that, on this account, philosophical intuitions are not fully *a priori*, since they are the product of experience.

4.4   The Case for Philosophical Skepticism

In the previous sections of this chapter, I discussed two models of scientific progress—the building of an epistemic system on a certain foundation and the constant improvement of an epistemic system by hypothesis testing—and have shown that both these conceptions depend on the postulation of reliable

immediate propositions and a determinate relation of logical consequence. I then introduced the problem of *Logical Pluralism* and discussed the nature of philosophical intuitions. In this final section of this chapter, I now show how what was discussed leads to *Philosophical Skepticism*; the thesis that there is no substantial philosophical knowledge.[20]

4.4.1 There are no Reliable Immediate Propositions

In philosophy, the role of immediate propositions is played by philosophical intuitions. But if philosophical intuitions are to play this role, then they need to be *reliable.* In this section, I argue that they are not.

A first attempt to undermine the idea that intuitions can play an evidential role in philosophy would be to point out that there is widespread consensus in the literature that intuitions are fallible. But discounting intuitions on this basis is, of course, too quick. Immediate judgments in the (other) sciences need not be infallible. For instance, participants regularly make grammatical performance errors in linguistic production and evalua-

---

[20]We should distinguish between two kinds of *Philosophical Skepticism.* First, there is the claim that we, as a philosophical community, do not currently have any substantial philosophical knowledge. Second, there is the stronger claim that substantial philosophical knowledge is impossible. In this dissertation, I only defend the first, weaker, version of philosophical skepticism. I do not know whether philosophical knowledge is essentially impossible. I do, however, confidently defend the thesis that substantial philosophical knowledge is not possible by employing philosophical methodology as discussed in this dissertation. The reader could still reasonably hope, as many philosophers have done in centuries before, that a revolution in philosophical methodology will finally lead to substantial philosophical knowledge. Given the history of philosophy, I am pessimistic, but, as I show in the next chapter, that does not mean that current philosophical methodology does not lead to anything of epistemic value.

tion tasks (Chomsky, 1965).[21] Furthermore, nobody is really worried that observation cannot play an evidential role in the empirical sciences because of the possibility of, e.g., perceptual illusions like the Müller-Lyer illusion. This is because judgments of grammaticality and our powers of observation are, generally speaking, considered to be reliable even if they are sometimes mistaken. The question is therefore not whether philosophical intuitions are infallible, but whether they are reliable.

It is important to note that, given the essentially non-empirical nature of paradigmatic philosophical issues, philosophical intuitions often cannot be independently confirmed or denied. For example, it might seem that torturing sentient beings for fun is wrong, but there is no way for us to independently verify this.[22] In this sense, philosophical intuitions differ from, say, mathematical intuitions and physical intuitions, which can be checked. It might, for instance, also seem that time must pass at the same speed for every object, although this intuition is refuted by the fact that one cannot fruitfully build a Global Positioning System under that assumption (Ashby, 2002). In philosophy, unfortunately, we have no such tests for our intuitions.

---

[21]We say that a participant has made a performance error in judgment if, after reflection, this participant herself robustly judges her previous judgment to be a mistake.

[22]One might retort here that we can sometimes check intuitions against philosophical principles that we hold. Remember, however, that in the conceptions of science that we are working with, these principles themselves must either be more fundamental (and the question that is being raised then is how to independently verify these principles) or these principles are currently being tested by these intuitions. I will deal with the interplay between modifying intuitions and modifying principles in the next chapter when I discuss the method of reflective equilibrium. However, in that chapter, I argue that equilibrium does not entail sufficient justification for knowledge.

We need to depend on indirect evidence when we want to judge the reliability of philosophical intuition.

Earlier in this chapter, I argued that intuitions, whatever they are, are best conceived of as being the product of Practiced Type 1 psychological processes. Intuitions can therefore be investigated by using empirical methods from psychology. This is precisely what is being done in the branch of philosophy called *Experimental Philosophy.*

Many results of this research program can be used to undermine belief in the reliability of philosophical intuitions. In particular, research in what is called the *negative school* of Experimental Philosophy aims to show that intuitions are indeed unreliable, or, at least, display a lot of variability. One major problem that is being raised by these philosophers, is that their experiments show that philosophical intuitions vary along dimensions that should be epistemically irrelevant.

The classical paper within this tradition is the study by Weinberg, Nichols, and Stich (Weinberg et al., 2001) which found that subjects in different cultural and socioeconomic groups have significantly different epistemic intuitions with respect to a wide range of classical philosophical problems. Furthermore, this paper also reports a study that showed an effect on intuitive judgments based on the number of philosophy courses subjects had taken; in essence, subjects who had taken more philosophy courses were more susceptible to skeptical arguments than those who had taken fewer. Similarly, Machery, Mallon, Nichols, and Stich (Machery et al., 2004) found that sub-

jects in different cultural groups have significantly different intuitions about reference. Other studies within this tradition show that there is variation among other dimensions too, among which are language (Vaesen et al., 2013), gender (Buckwalter & Stich, 2014) and age (Colaco et al., 2014). These results concerning philosophical intuitions lead Alexander & Weinberg (2007) to conclude that:

> such instability impugns the status of intuitions as evidentiary, making it unlikely that there would be any fixed set of intuitions about a particular thought experiment that can stand in any sort of evidential relationship with a philosophical claim. (Alexander & Weinberg, 2007, p. 66)

That is, the fact that people seem to differ so much in what they intuitively judge to be the case, coupled with the fact that these judgments are strongly correlated with epistemically irrelevant factors, indicates that intuitive judgments are not reliable.

Not only does this instability show that intuitions are not reliable, it also gives us a grasp on what the difference is between linguistic intuitions, the paradigmatic model on which the evidential role of philosophical intuitions seem to be based, and philosophical intuitions themselves.

It is important to note that linguistic intuitions can only play their evidential role, because linguistics is a *descriptive* science. That is, the scientific models that are developed in linguistics attempt to capture the rules that

generate all and only the sentences that people actually think are grammatical, or, at least, try to explain these judgments.[23] In language, the final authority lies with the speakers of that language. If a linguistic community judges that a sentence is grammatical, then this grounds the fact that that sentence is grammatical. In philosophy, however, the purpose is not to explain the intuitions, but to use these intuitions to help understand the subject-matter the intuitions are about, and the truth of philosophical intuitions is not grounded in the authority of the philosopher. Even if it seems to everyone that it is wrong to torture sentient beings for fun, it is still an open question whether it is indeed wrong or not. Another way to make this difference clear is to see that, if every speaker of English judges that "She used to secretly admire him" is a grammatical sentence, then that is the end of the matter. No appeal to the principle that one cannot split infinitives can show that English speakers are mistaken. If every philosopher in the world judges that it is permissible to torture sentient beings for fun, that fact alone doesn't make it so.

The above seems to be a strong argument against the reliability of intuitions. From the above we have seen that we have good evidence to suppose that philosophical intuitions are unreliable. Furthermore, as (Cummins, 1998) has argued, although we do have independent reasons from psychology

---

[23]Of course, this method is not free from problems itself. For instance, there are borderline cases where people's judgments concerning the grammaticality of a sentence differ and people sometimes make performance errors. Furthermore, there are problems related to how to conceive of *dialects*, i.e. regional languages that are closely related but subtly different, and *idiolects*, i.e. the distinctive and unique use of languages of individuals.

to believe that philosophical intuitions are unreliable, we seem to have no independent reasons to believe that intuitions are reliable. A related problem with philosophical intuitions is that we hardly seem to have a naturalist account why they are reliable. In my eyes, pointing out that this does not necessarily mean that such an explanation is impossible, as (Pust, 2004) does, doesn't mitigate the problem that we currently don't seem to have any convincing explanation of how intuitions could possibly be reliable. But can nothing be said in their favor?

It has been argued that one possible defense of intuitions is to insist that it is intuitively true that intuitions are reliable (Pust, 2017). This defense might be epistemically circular, but one might try to argue that this circularity cannot always be problematic on pain of Philosophical Skepticism. But this kind of circular defense looks unappealing to me. First, I do not share the intuition that philosophical intuitions are reliable. Many forms of scientific intuition have proven to be unreliable. I see no intuitive reason why philosophical intuitions should be different. And second, since I defend a form of Philosophical Skepticism in this chapter, I do not feel that the consequence of rejecting epistemic circularity is unacceptable.

What other defense can be offered? There is the argument from (Williamson, 2007) that, in general, Type 1 processes must be reliable since we essentially depend on them in our daily practices. Perception, generally speaking, is also a Type 1 process on which we heavily rely in daily practice, even if we are sometimes misled under the artificial conditions of a laboratory. Why

isn't the philosopher's dependence on intuitions similar?

Let us grant Williamson that we constantly rely on Type 1 generated judgments in daily life. One can reply to Williamson that philosophical practice is not daily practice. Given how I defined philosophical knowledge above, I'm committed to believe that one needs to meet the higher standards in order to have philosophical knowledge compared to daily life.[24] I agree that meeting this standard is hard, but I have given independent reasons to believe philosophy has a true Skeptical problem.

There is also the worry that the studies from Experimental Philosophy on which I have relied might not indicate what they seem to indicate. First, there might be methodological problems with how philosophers set up their experiments (see, for instance, (Sosa, 2007) and (Deutsch, 2009) for this critique). For instance, experimenters might not make sufficiently certain whether subjects imagine the details of the cases in the same way. Similarly, experimenters might not make sufficiently certain whether subjects interpret crucial theoretical terms in a similar manner. Variation among subjects might therefore sometimes be explained by differences in interpretation rather than differences in intuition. Second, there has been some difficulty replicating some of these studies (See (Nagel et al., 2013), (Adleberg et al., 2014), and (Kim & Yuan, 2015)). Time will tell whether these critiques hold up as Experimental Philosophy matures. Personally, I am confident that future work

---

[24]Basically, here I agree with Nado (Nado, 2016, p. 787) that one can reply to Williamson by defending a contextualist account of knowledge.

in Experimental Philosophy will robustly show that philosophical intuitions do vary actually widely among epistemically irrelevant dimensions. However, even if intersubjective consensus about intuitions is larger than currently believed, we need not fully rely on the empirical results from experimental philosophy to argue that philosophical intuitions are not reliable.

First, Koriat's (Koriat, 2008) empirical evidence cautions us that the strength of an intuition correlates with the consensus that people have about it, not with its correctness. This means that intuitions cannot play the role of self-evidence needed. At most, it shows that there is more consensus in philosophy than is often thought.

Second, there is the conceptual argument based on an argument given by Stich (Stich, 1990). According to Stich, it is conceptually possible that we could have had different intuitions about philosophical matters than we actually have. If that is indeed so, then philosophical theories developed by people who have different intuitions would be just as justified as our current philosophical theories, even though their philosophical theories would disagree with our current theories in content. This possibility undermines the reliability of intuitions to serve as evidence for our philosophical theories, especially if we think philosophy is in the business of finding necessary truths.

It is true that Stich's argument itself depends on the modal intuition that creatures with different intuitions than our own are possible, and the argument therefore undermines itself if intuitions are unreliable (Pust, 2004). However, Stich's argument combines well with the work done in Experimental

Philosophy. If empirical evidence does show that there are people who have different intuitions from our own, even if this variability is not as big as the proponents of the negative school would claim it to be, then it is still clearly possible that there are people who have different intuitions. And given that this possibility is empirically rather than intuitively justified, Stich's argument goes through.

I conclude that intuitions are never self-evident; they always allow for further scrutiny as to why they are reliable. In the words of Ichikawa and Jarvis:

> What [the evidence of intuitions' unreliability] indubitably show is that many, perhaps all, judgments that are putatively a priori knowledge are susceptible to a posteriori defeaters. One can always acquire empirical reason to believe that one's a priori faculties are functioning poorly and ought not to be trusted. (Jenkins Ichikawa & Jarvis, 2013, p. 188)

This means that any philosophical intuition that is appealed to can reasonably be denied by a critic. Furthermore, the unreliability of intuitions leaves open the possibility of the radical revision of any putative philosophical system when intuitions that seemed reliable are discarded.

Given their unreliability, philosophical intuitions cannot be used to ground a philosophical view, nor test it. But can they, at the least, be used to abductively infer a theory? Let us look at a proposal by Timothy Williamson,

who argued in the recent paper "Abductive Philosophy" that philosophical theorizing has the form of an *inference to the best explanation* (Williamson, 2016).

Recall that abductive reasoning proceeds from a set of data to principles that best explain that data. We frequently see this pattern of inference in the (other) sciences. For instance, any statistical regression can be seen as an instantiation of this pattern. Williamson claims two things. First, he claims that philosophers already use abductive inferences in their theorizing. This is entirely correct. We see many instances of abductive reasoning in philosophical practice. The second claim is that philosophers *should* use more abductive arguments.

I do not necessarily disagree with Williamson about the method. We can see in the game of giving and asking for reasons that I used as a model of philosophy that when Skeptic asks "why", she asks for an explanation — preferable the best the proponent has to offer, of course. I do, however, deny that providing these arguments leads us to, or closer to, finding true theories.

First of all, as Williamson himself also notes, "[i]nference to the best explanation is obviously fallible; it can lead us from some true evidence E to a false theory T." (Williamson, 2016). However, as in the case of intuitions, let us only demand that IBE is reliable, something which is shown, I think, by the reliable use that is made of it in hard sciences like physics. In Williamson's words, "the central role of abduction in the success of the natural sciences provides good reason to think that it is a good method, even though we do

not fully understand why." (Williamson, 2016) Second, abductive arguments are not decisive. Again, when we look at it from our Bayesian perspective, we see that, depending on the prior of the epistemic agent, an abduction can be convincing for one agent, but not for the other. The third problem is about what data philosophers should explain. Let us grant Williamson that "[n]othing in this account requires the evidence propositions, the explananda, to be of some special kind. Any known truths will do." (Williamson, 2016). Williamson states that he thinks we should use our *total evidence*, that is, all our current knowledge.

But here, a crucial difference between Williamson's and my position emerges. I believe philosophical knowledge differs in kind from scientific knowledge. Scientific knowledge, although its standards are high, need not be beyond a reasonable doubt in the strict sense that I'm using it. This is why inductive and abductive inferences can ground scientifc truth. Furthermore, there is nothing mysterious about there being different standards. For instance, methods that are acceptable in physics, such as inference to the best explanation, are not acceptable in mathematics where different standards hold.

Williamson claims that abduction grounds the *first principles*, that is the axioms and the inference rules, of mathematics. Now, it is clear that we do no longer believe them to be *self-evident*. The axioms of ZFC, however plausible you think they are, are not self-evident. I still have the intuition that, were it not for the paradoxes, the principle of comprehension is self-evidently

true. However, like Shapiro and Potter, I favor postulatism. Otherwise, we cannot explain how two mutually inconsistent theories can both be a part of mathematics. Williamson might believe that pluralism must stop at the meta-level, or at least at some meta-level, but, following Shapiro, I see no reason to think that this is so.

Now since these inferences take place in a philosophical context, I demand that these known truths, as any philosophical truth, are *beyond a reasonable doubt.* We have already put the evidentiality of philosophical intuitions to question. If they are not reliable enough to ground or refute a theory, then there is no reason to believe that they are reliable enough to serve as a base for a philosophical inference to the best explanation. Why not provide a psychological explanation for them instead?

The fourth problem has to do with the term 'best' in 'best explanation'. Williamson seems to assume that there is some objective measure by which to determine which of many possible explanations is ultimately best, i.e. that there is one set of theoretical virtues that ranks all the theories. But such an evaluation of a theory is value laden. Depending on one's values, one could rank the theories differently.

It might be true that simple explanations are less likely to overfit the data (Forster & Sober, 1994). Williamson uses this to argue that one need not necessarily become a pragmatist. There might be objective reasons for all of the theoretical virtues. However, there are other ways of preventing overfitting, for instance by way of a lasso or a ridge regularization which add

regularization terms. If one works as a data scientist, one actually becomes quite pragmatic about these measures.

The reader might grant that most intuitions are not reliable, but surely no reasonable being would deny the moral intuition that it is wrong to torture sentient beings for fun. This might indeed *seem* so to the reader. However, there are philosophers called *error theorists* who would grant that it does indeed *seem* to them that it is wrong to torture sentient beings for fun, but who nevertheless insist that moral propositions such as these are false, because moral qualities do not exist.[25]

Nevertheless, it seems that intuitions play a crucial role in philosophical methodology (pace Capellen). Could philosophical intuitions perhaps have another function to fulfill? I think yes. I agree with (Cappelen, 2012) that talk of intuitions flags that a proposition is considered to be a part of the *common ground*, that is, to flag that it is a proposition that is not under discussion. A similar view is defended by (Nagel, 2014), who claims that the term 'intuitively' is used in a dialectical context to flag claims that the intuiters believes will be shared with others. We find the same view defended by (Ichikawa, 2014), who thinks intuitions signal a dialectical move. In the words of De Cruz, "[p]hilosophers may not be expert intuiters *per se*, but they are expert elicitors of such intuitions in dialectical contexts." (De Cruz, 2015, p. 247).

---

[25]The classical account of an error theory is (Mackie, 1977). For a more recent defense of an error theory, see (Streumer, 2017).

4.4.2   There are no Knockdown Arguments

In the previous section, I argued that intuitions are unreliable. As such, they cannot be used as a foundation nor as a testing bed for philosophical theories. In this section, I want to look at the more general issue of knockdown arguments in philosophy.

For the purpose of this dissertation, a knockdown argument is an argument that definitively refutes a philosophical thesis. The main question concerning knockdown arguments in philosophy is whether they exist. Peter van Inwagen (van Inwagen, 2006, ch. 3), for instance, has argued that there are no knockdown arguments against substantive philosophical conclusions. Similarly, David Lewis, has claimed that

> [w]hether or not it would be nice to knock disagreeing philosophers down by sheer force of argument, it cannot be done. Philosophical theories are never refuted conclusively. (Lewis, 1983, p. x)

It is widely believed that knockdown arguments exist, and in this dissertation, I will simply assume this too. For instance, even for non-physicists like us, being shown a feather and a hammer fall in a vacuum definitively refutes any Arisotelian principle we might have held that denser objects fall faster than less dense objects.[26] Van Inwagen also maintains that there are

---

[26]For the classical demonstration, performed on the moon by David Scott, that a feather and a hammer fall at equal speed in a vacuum, see `http://nssdc.gsfc.nasa.gov/planetary/image/featherdrop_sound.mov`. For a more recent demonstration of

knockdown arguments in the (other) sciences (van Inwagen, 2009, p. 10).

A good example of a purported knockdown argument in philosophy is Gettiers argument against the JTB-account of knowledge that has served as a running example in this dissertation. But we have already seen in the previous section that intuitions, in Gettier-cases in the form of a counter-intuitive consequence, cannot definitively refute a philosophical thesis because they are not sufficiently reliable. In normal circumstances, we can trust our eyes when we see that feather and hammer hit the surface of the moon at the same time. In normal circumstances, we cannot, however, trust our judgment that a Gettier case is a case in which a subject has a justified true belief but not knowledge. For one, Gettier-cases can be disarmed if one is willing to adopt a more stringent account of justification. On those accounts, even though it might *seem* that Smith is justified in believing that Jones will get the job, he does not actually have sufficient evidence to justify his belief.

Other purported knockdown arguments in philosophy, such as Searle's *Chinese Room* (Searle, 1980), Block's *Chinese Nation* (Block, 1978), and Thomson's *Transplant Case* (Thomson, 1985) have all been rationally challenged in a similar fashion by professional philosophers. On the other hand, no professional physicist has ever rationally challenged David Scott's observational evidence that a hammer and a feather fall at the same time in a vacuum, and there seems to be no possible ground to do so.

---

the same thing in a vacuum chamber on earth, see `https://www.youtube.com/watch?v=E43-CfukEgs`. Note that, barring sufficient resources, there is nothing preventing a skeptic in principle from performing these experiments themselves.

The case for knockdown arguments in philosophy, therefore, looks bleak from the start. However, in a recent paper, Nathan Ballantyne (Ballantyne, 2014) has argued that if there are non-philosophical knockdown arguments, then there must also be knockdown arguments of philosophical theses.[27] Let us discuss his argument.

In his paper, Ballantyne writes

> Perhaps, for any putative knockdown argument, a thinker can understand it but reject one of its premises or assumptions, and thereby do nothing irrational by not accepting its conclusion. That may be defensible and I don't mean to dispute it here. I want to observe, though, that this reason to think there are no knockdown arguments in philosophy seems to count equally against knockdown arguments outside philosophy. (Ballantyne, 2014, p. 533)

Now remember that we have assumed that there are knockdown arguments outside philosophy. Ballantyne now argues that some of these knockdown arguments can also be used to refute certain philosophical positions. For instance, according to Ballantyne, the established scientific fact that the earth orbits the sun refutes the substantial philosophical view that there is no motion. Similarly, Ballantyne takes it that the fact that all atoms are consti-

---

[27]Ballantyne takes knockdown arguments to be the establishment of a thesis, thereby the refutation of its negation and the impossibility of a suspension of judgment. I take it this difference between how to conceive of knockdown arguments is inconsequental for my use of his argument.

tuted by electrons, neutrons, and protons refutes philosophical monism, the claim that there is only one thing.

Again, we can answer this charge by appealing to contextualism. Philosophical knowledge simply seems to have stricter standards than scientific knowledge. Some propositions, such that there is no motion, cannot be accepted in a scientific context. Similarly, the proposition that people are not responsible for their actions because free will does not exist cannot be accepted in a court of law. But that does not mean that there are no genuine philosophical questions here. If motion is indeed paradoxical, then there is a real philosophical problem, even if physicists can legitimately ignore this problem.[28] More generally, premises that are not problematic in an argument for ϕ in context A may very well be problematic in an argument for ψ in context B, even if ϕ entails ψ (Keller, 2015, p. 1208).

Ballantyne does try to answer this challenge of contextualism. He writes

> [E]ven if shifting standards prevent the entailed philosophical propositions from counting as established facts by strict philosophical standards, those propositions can still count as established facts by the relevant scientific standards. Consequently, some philosophical propositions can be supported by knockdown arguments in whatever sense an argument is knockdown in the sci-

---

[28]As we will see in the next section where I discuss *philosophical midwifery*, I do think that where philosophical problems truly manage to make connection with established scientific fact something interesting is going on. However, I don't think that Ballantyne's cases are examples of this.

ences. If different standards hold sway in philosophy, who cares? Again, some arguments for philosophical theses succeed by scientific standards—philosophers ought to celebrate. (Ballantyne, 2014, p. 539)

But at this point, the reader might be reminded of Samuel Johnson's purported refutation of Berkeley's Idealism by kicking a large stone and asserting, "I refute it thus." It misses the point. Even Parmenides granted that it appeared to be the case that there was motion, and monists like F.H. Bradley similarly granted that it appears to be the case that there are many things. In Kantian terms, although the world as we experience it might be adequately described by science, the world as it is in itself might still be fundamentally different even if we have no way of knowing it.

4.5 Philosophical Midwifery

Another way in which philosophy is sometimes said to be making progress is by philosophical *midwifery*; i.e. by giving birth to a new science. Standard examples of sciences that grew out of philosophical investigation are physics and psychology. Physics, which was still called "natural philosophy" up to the 19th century,[29] started to become independent from philosophy from the 17th century onward, when physicists started to emphasize the use of ob-

---

[29]For instance, Lord Kelvin and Peter Guthrie Tait's influential physics textbook from 1867 was published under the title *Treatise on Natural Philosophy* (Thomson & Guthrie Tait, 1867).

servation and experiment over the use of pure reason. Similarly, psychology became a science in its own right in the 19th century, when people started looking for ways to observe the workings of the human mind and developing systematic methods to explain and predict human behavior. One could argue that mathematics and astronomy, both of which were studied in Plato's Academy, already underwent a similar development during the time of the Ancient Greeks.

Here is the idea elegantly presented by Austin:

> In the history of human inquiry, philosophy has the place of the initial central sun, seminal and tumultuous: from time to time it throws off some portion of itself to take station as a science, a planet, cool and well regulated, progressing steadily towards a distant final state. (Austin, 1979, p. 232)

The main idea here is as follows. Remember that we defined philosophical problems in an earlier chapter as those that cannot be solved completely using empirical and formal methods, although empirical and mathematical considerations might indeed restrict the logical space of possible answers. Now, as better empirical and formal methods are developed, some of these problems might become completely solvable. When that happens, these problems stop being philosophical.[30]

---

[30]See, for instance, (Russell, 1912, p. 89) and (Searle, 2016, p. 544) for an exposition of this idea.

If this is indeed a fruitful way of thinking about philosophical progress, then we shouldn't be surprised by our earlier conclusion that there is no store of philosophical knowledge. That is, of course, because once philosophers have managed to find systematic ways to demonstrate or refute a philosophical thesis, future philosophers will not be able to claim that victory for philosophy. Rather, that victory would be awarded to the relevant science. It would, for instance, be strange to claim the refutation of the perfect nature of the heavens by e.g. spotting mountains and craters on the moon as a victory for philosophy. Something similar holds for the refutation of vitalism through methods of modern biology.

The reader might be reminded of Auguste Comte's *law of three stages* when thinking about philosophical progress in this sense. Remember that, according to Comte, scientific development goes through three stages: the theological stage, the metaphysical stage, and the positive stage. It is, indeed, tempting to identify the metaphysical stage with the philosophical phase of a problem and the positive stage with its scientific phase. Although it would go too far to discuss Comte's law of three stages in detail here, note that this resemblance might only be superficial since, according to Comte, the positive stage is characterized by abandoning the search for the causes of a phenomenon in favor of formulating laws that govern its behavior (Bourdeau, 2014). Such a commitment to scientific positivism is, however, not presupposed in the idea of philosophical midwifery.

Although I do indeed believe that some problems that we now deem

philosophical could be answered by future developments within the sciences, it is important to note that there is still a difference in kind between philosophical problems and scientific problems, even if it is sometimes hard to see now whether a problem is indeed a philosophical problem rather than an extremely hard scientific one.

First, note that not all philosophical problems can be solved using scientific methods. Because of their normative nature, ethical problems, and problems from political philosophy, cannot be turned into an empirical or formal science (see also (Hacker, 2009) for this argument). No amount of observational data will finally turn ethics into a science, nor will any formal proof show how to implement a just society. Although these issues might be empirically and formally informed, they cannot be solved by finding out what is actually the case. Similarly, general conceptual questions cannot be answered by finding out how things are. For instance, no experimental result will solve problems regarding vague boundaries in nature such as finding a precise delineation between things that are alive from things that are not. Scientists might, of course, fruitfully work on e.g. trying to understand the nature of vagueness in categorization, but by our definition, these are philosophical issues, not scientific ones.

Second, we need not be afraid that we will eventually exhaust philosophy as more and more issues become amenable to scientific scrutiny. Although Austin might have thought that there were only about a thousand philosophical problems left (Searle, 2016, p. 545), there is no reason to believe that

scientific developments will lead to a reduction of problems that are deemed to be philosophical. The birth of a new scientific discipline, in general, tends to give rise to more philosophical problems, not fewer. This is because certain philosophical problems within that science become more salient as it matures (Hacker, 2009, p. 132). For instance, with the birth and maturation of physics also came the birth and maturation of a philosophy of physics, and the development of modern physics (especially quantum mechanics) has made it more rather then less mysterious how to conceive of causation. Similarly, with the rise of cognitive science we now have more, rather then fewer, ways in which we can conceive of cognition.

Third, although some philosophical questions that were posed by the Ancient Greeks, such as those concerning the nature of the heavens and the nature of the infinite, have become the subject-matter of empirical and formal sciences, many other questions that e.g. Aristotle posed are still the subject of philosophical deliberation today. Questions concerning the nature of substances and the nature of universals are good examples of this. There is, therefore, no historical reason to think that every problem will eventually become the subject-matter of science.

The above discussion suggests that there are indeed two kinds of issues that are considered to be philosophical; those that only differ in degree from scientific problems in the sense that they seem hard or even impossible to answer by current scientific methods, and those that differ in kind. Those that differ in kind will never be solved using methods from the empirical and

formal sciences (although, again, they might be informed by these sciences), and there seems to be every reason to believe these problems are proliferating. Where there is indeed a difference in kind, we do not lack philosophical knowledge because philosophical questions are so much more difficult to answer, but because such questions cannot be definitively dealt with at all. This is why I agree with Russell (Russell, 1912, p. 86), that the aim of philosophy regarding such problems cannot be to finally develop scientific knowledge. What we can do is is provide criticisms of philosophical views.

Chapter 5

The Epistemic Value of Philosophy

In the previous chapter, I defended *Philosophical Skepticism* on the basis of a dialectical model of philosophical practice that I introduced in chapter 3. I argued that, in principle, no philosophical knowledge is possible in the sense that no fully explicated argument for a substantial philosophical thesis brought forth by a proponent will ever convince a universal audience. At the same time, no opponent will be able to refute all possible explications of a philosophical thesis either. The reason for this is that we cannot assume the existence of a universal common ground in philosophy, and any purported refutation can always either be avoided by choosing different inference rules or be accepted by "biting the bullet". Different rational choices of immediate propositions and logical principles in philosophy are always possible in

169

principle.

But if philosophical skepticism is correct, then what could the epistemic purpose of engaging in philosophical practice be? The claim that philosophy has no epistemic value at all is hard to accept (especially for a philosopher). As Clinton Golding expresses this sentiment, such a claim "is perplexing given how many people pursue philosophy, both formally and informally, and who seem legitimately satisfied by what is achieved." (Golding, 2011, p. 200) Although I agree with Golding that such a claim would be perplexing given the rich history of philosophy, I do need to point out that even if it is the case that many people satisfactorily engage in a practice, this does not necessarily mean that such a practice has epistemic value. Many people, for instance, engage in astrology, homeopathy, and ufology, and many of these practitioners seem genuinely satisfied with what they achieve. Still, these practices offer nothing of epistemic value. If we therefore want to maintain that philosophy has some epistemic value after all, and if this value is not due to the production of knowledge, then we need to articulate what this value consists of. That is the purpose of this final chapter.

5.1   Reflective Equilibrium

In previous chapters, we have seen that, because of the high epistemic context of philosophical practices, no *common ground* can be found in a universal audience on the immediate propositions and the rules of inference. It is

always possible that a neutral rational agent will refuse to accept anything that the proponent proposes as common ground, although, at the same time, he also cannot demonstrate that it is irrational to hold the propositions and principles that the proponent claims as *bedrock*.

This means that, in the limit of playing the Game of Giving and Asking for Reasons, the process of philosophical practice leads to a stable system of propositions where every ground and consequence is made explicit. Every ground has been claimed as bedrock by the proponent, and every consequence of the logical principles that the proponent has put forward has been deemed acceptable. However, since an arbitrary opponent does not need to accept the bedrock propositions, the universal audience is not convinced. They might not be able to refute the position put forward, but nothing forces them to accept it either. Only the particular audiences that have granted the bedrock propositions and the logical principles that the proponent has put forward must necessarily grant the truth of the thesis under discussion, on pain of irrationality.

The reader has, no doubt, recognized the above state where there is a coherent system of propositions as a product of philosophical practice. It is usually called *reflective equilibrium* in the philosophical literature. The basic idea of the method of reflective equilibrium is that a system of judgments and principles is justified if it is the product of a process of adjustments after which all the judgments and principles are in equilibrium with one another.[1]

---

[1] I will explicate what it could mean for a system of propositions and principles to

This method was first popularized by John Rawls, who gave the method its name (Rawls, 1971), although the use of the method itself can already be found in, e.g. (Goodman, 1955).[2] For a more detailed account of this method, the reader may consult (Elgin, 1996) and (Elgin, 2005).

As Georg Brun has pointed out, the literature on reflective equilibrium often leaves it undetermined what the differences are between judgments and principles (Brun, 2014, p. 239). He is right that the fairly standard idea that general propositions are principles while particular propositions are judgments doesn't work because judgments can also be general, as was also observed by Rawls himself (Rawls, 1999, p. 289). However, in our case, the differences between the two are clear. Judgments are expressed by propositions, while principles are expressed by rules. I prefer this characterization over Brun's idea that principles are propositions within the system—either grounds or consequences—while judgments are propositions outside the system that the proponent feels committed to in some degree which are in tension with it. That is because Brun's explication obscures the need for both propositions and rules, a lesson that we learned earlier from Achilles and the Tortoise in (Carroll, 1895).

Reaching a total reflective equilibrium as an end product is, of course,

---

be in equilibrium below. For now, the reader may simply assume that it means logical coherence.

[2](Goodman, 1955), in particular, defended the view that the method of reflective equilibrium could be used to inductively justify the logical principles that one uses. In recent years, we have seen this methodology resurface as an attempt to answer the problem of logical pluralism.

an ideal situation. As of today, no philosophical position is completely without problems. And such an ideal state need, of course, never be actually reached. It should rather be seen as a regulative ideal that participants of philosophical practices strive towards. Furthermore, while being in equilibrium or not is a categorical matter, there is an important matter of degree here. Some philosophical positions are more problematic then others, i.e. some positions are further away from equilibrium than others (however this should be measured). It is, however, nevertheless theoretically useful to think of reflective equilibria as the best epistemic result that philosophical practice can achieve.[3]

5.1.1 Equilibrium, Justification, and Knowledge

We have seen that the aim of the method of equilibrium is to justify a system of propositions. Let us grant that achieving a reflective equilibrium is indeed such a justification. Why does granting this not refute my claim that philosophical knowledge is impossible? It seems that, if the propositions of the system are indeed true and the method of reflective equilibrium does indeed confer justification on the system, philosophical knowledge is possible after all.[4]

It is, however, important to distinguish between *doxastic* justification

---

[3]It is important to note that, while I grant that reflective equilibrium can only be achieved in ideal cases, I have argued in chapter 4 that philosophical knowledge cannot be achieved even in the ideal case.

[4]See below for some remarks on the coherence theory of truth.

and *personal* justification here (Littlejohn, 2009). Doxastic justification is a justification of a proposition. It demonstrates *that* a proposition is true (and in the ideal case, it also explains *why* that proposition is true). Personal justification, on the other hand, is justification for *holding a belief.* It explains *why* an epistemic agent is rationally permitted to hold a certain proposition to be true. When I grant that the method of reflective equilibrium confers some form of justification, what I mean is that it confers personal justification, not doxastic justification.

I take it that the reader will simply grant me that reaching a reflective equilibrium grants personal justification for holding the position true. It is, of course, not a problem for me if a reader refuses to grant me this. Although I do indeed believe that philosophers can have personal justifications for their views, none of my arguments in favor of philosophical skepticism depend on this.

However, personal justification is not the right kind of justification for knowledge. What is needed is doxastic justification. To see this, suppose that there are two philosophers Smith and Jones. Smith has achieved a reflective equilibrium with respect to a philosophical thesis T, while Jones has achieved a reflective equilibrium with respect to the negation of that thesis. For concreteness, let us say that Smith is an idealist while Jones is a realist. Given that they are both in equilibrium, neither Smith nor Jones is epistemically culpable for the position that they hold. They are both rationally permitted to hold the views that they do. That is, both of them

are personally justified in holding their views. However, it is also obvious that neither of them knows that e.g. idealism is true or false. That is because, in this case, it would just be a matter of luck who happens to hold true beliefs and who happens to hold false ones. Such a form of epistemic luck is widely acknowledged to be incompatible with knowledge (for a detailed discussion of epistemic luck see e.g. (Pritchard, 2005)).[5]

What I need to show is that achieving a reflective equilibrium does not confer doxastic justification. The strategy that I will pursue here is to show that a set of beliefs and principles being in reflective equilibrium is not sufficient for it to be more probably true.

In the previous chapter, we have already rejected foundationalism in philosophy. There is simply no set of immediate propositions and principles that can serve as the universal common ground to build a system of philosophy on. As such, when one has reached reflective equilibrium, one cannot doxastically justify the equilibrium by claiming that justification is transferred through the intuitions and logic that anchor the view. As we saw in the previous chapter, intuitions and logical principles simply lack such doxastic justification themselves, and therefore, cannot transfer it.[6]

If foundationalism is out, what about coherentism? Coherentists claim

---

[5]For readers who are experts on the topic of epistemic luck, this is a case of *veritic luck*, i.e. the kind that is problematic for knowledge, because it is a matter of luck that e.g. Smith beliefs the truth given her evidence while Jones does not. It could just as easily have been the other way around.

[6]I do grant that one can be personally justified in holding intuitions and principles by designating them as bedrock. As such, personal justification can transfer through the system.

that justification is not conferred *atomically*, i.e. by particular beliefs, but holistically, i.e. by the way a set of beliefs coheres together.[7] That is, justification transfers from the whole to its parts. One way to try to justify a philosophical view is therefore to insist that being in a reflective equilibrium doxastically justifies both the philosophical view and the intuitions and principles that anchor it.

In order to investigate this claim, one should spell out precisely what it means for a set of propositions and principles to cohere. As always in philosophy, one can find many different notions of coherence in the literature.[8] For my purpuses, it is sufficient to follow Lehrer's explication that a system of beliefs is coherent when all objections of the Skeptic are met (see (Lehrer, 2000) and (Lehrer, 2003) for a detailed exposition of this view. Note that Lehrer uses slightly different terminology). This view dovetails nicely with our dialectical model of philosophical practice.

Again, I do not doubt the intuition that coherence may confer *personal* justification. Remember, the claim to be defended here is only that coherence does not confer *doxastic* justification. We explicate this by saying that the coherence of a philosophical view does not guarantee a raise in the likelihood of that view. It could perfectly be the case that, as the philosophical view becomes more coherent, it becomes less likely to be true.

---

[7]For a classical presentation of coherentism, see (BonJour, 1985).

[8]For a selection, one can see (Ewing, 1934) for a classical notion of coherence as consistency, (Lewis, 1946) for a classical notion of coherence as congruence, and (BonJour, 1985) or (Thagard, 2007) for a more modern conception of coherence as a complex notion.

To see this, let us look at the argument presented by Peter Klein and Ted Warfield in their classic paper (Klein & Warfield, 1994). In that paper, Klein and Warfield argued that, in general, the coherence of a set can be raised by adding information that supports the information that is already in the set. However, given that the added propositions are not completely certain themselves, this will decrease the probability of the whole set being true. This simply follows from the well studied inverse relationship between probability and informational content. Although things turned out to be more problematic than Klein and Warfield expected, recent impossibility results have indeed shown that no plausible measure of coherence is truth conducive even assuming a weak *ceteris paribus* sense (for a presentation of these results see, e.g. (Bovens & Hartmann, 2003) or (Olsson, 2005, Appendix B)). Although this debate is far from over since the impossibility results depend, as always, on assumptions that can be denied or weakened, it is in no way straightforward to defend the view that having a coherent system of beliefs entails that that system is more likely to be true.

5.2 The Underdetermination of Reflective Equilibria

In the previous section, I argued that the best a philosopher can do is reach a *reflective equilibirium* concerning a substantive philosophical thesis and that, while such equilibria do confer personal justification, they do not confer

doxastic justification. As such, by achieving a reflective equilibrium one does not gain the right kind of justification to be able to claim knowledge. In this section, I wish to give a second argument often found in the literature that basically makes the same point. Because philosophical equilibria are underdetermined by their evidence, they are not free from personal values.

It is clear that reflective equilibria are not necessarily unique, even given the same starting evidence. Different choices can be made to stabilize a web of beliefs that is out of balance. This problem becomes especially clear when we also incorporate the problem of *logical pluralism* into this discussion; not only can different choices be made concerning which propositions are ultimately considered bedrock and which consequences ultimately need to be accepted (albeit by gritting the teeth); different rational agents may also choose different rules of consequence that they deem to be correct. Even Rawls himself, although originally defending the idea that theories in equilibrium converge to the truth (Rawls, 1971), had to change his mind and admitted that the method of reflective equilibrium ends in a plurality of justifiable views (Rawls, 1993).

This is a problem, because this means that the method of reflective equilibrium does not necessarily lead to a narrowing of the logical space with respect to a philosophical problem, and can even result in a widening of it. Of course, if one wants to reach knowledge, this is exactly the opposite of what one needs. Progress through knowledge demands, at least, that a theory converges on truth over time. But what the method of reflective equilibrium

is giving us is a divergence of justifiable views each time a choice is made to restore equilibrium.

We should, of course, not be too surprised by this result, since it is nothing but the problem of the *underdetermination of scientific theories* made famous by Pierre Duhem (Duhem, 1954) and Willard van Orman Quine (Quine, 1951). However, philosophy is much more vulnerable to the epistemological problems that are related to this underdetermination due to the nature of philosophical evidence.

When we think about the underdetermination of theories, it is important to distinguish *holist* from *contrastive* forms of underdetermination (Stanford, 2013). A theory is holistically underdetermined when there is room for a choice about how to accommodate an unacceptable consequence. Reflective equilibria are underdetermined in this sense. A quite standard critique of the method of reflective equilibrium is that it does not tell us what to save and what to revise when we find ourselves out of equilibrium. A proponent of a philosophical view, for instance, has a lot of freedom to "fix a problem" that has come up with her theory in any way she sees fit. Usually, in the philosophy of science, this idea is captured by the observation that there are no such things as "crucial experiments". In our specific case of philosophy, this is a more general form of the particular problem I discussed earlier that there are no "knockdown-arguments" in philosophy. There is even an additional complication in the case of philosophy, since, apart from the problem of what to change in the web of belief when confronted with an unaccept-

able consequence, philosophers can usually even challenge the judgment that the consequence is unacceptable to start with itself. They usually do so by stating the infuriating sentence that "one person's *modus ponens* is another one's *modus tollens*".

A theory is contrastively underdetermined if different theories can be defended on the basis of the same evidence. The standard example to illustrate this is that an infinite number of curves can always be drawn through the same set of data points. In the case of the method of reflective equilibrium, that would mean that different equilibria can be established from the same starting point. Usually the idea that this is always a theoretical possibility is already sufficient to make this problem pressing, even if the alternatives do not look very promising. For instance, a disjunctive theory that states that all objects in motion continue in motion with the same velocity and direction unless acted upon by an unbalanced force until January 1st, 2020, after which all objects in motion will start to continuously accelerate by a constant factor, is a theory that is consistent with our current data, although nobody thinks that this theory is plausible.[9] In philosophy, however, there is not just the mere theoretical possibility that different theories can be defended on the basis of the same intuitions, but the fact that alternative theories

---

[9]This is, of course, nothing but Goodman's *New Riddle of Induction* (Goodman, 1955). Incidentally, Goodman is also famous for holding the controversial view known as *irrealism*, the philosophical view that the world dissolves into world-versions. For our purposes, it is sufficient to think of a world-version as an empirically adequate theory. According to Goodman, every world-version corresponds to an actual world. The interested reader should take a look at (Goodman, 1978). One could think of the view that I present here as a form of *philosophical irrealism.*

with respect to the same intuitions are defended by choosing whether certain intuitions need to be accepted or explained away. An example of this from the history of philosophy is how to respond to Gettier-cases; one can reject the justified true belief analysis of knowledge, one can try to repair the analysis by adding a fourth condition that aims to rule out lucky cases, or one can save the analysis by trying to explain away the intuition that people in Gettier-cases are justified in their beliefs.[10]

From the above it becomes clear that, not only do philosophical theories suffer from both holistic and contrastive underdetermination, they suffer from it more intensely because philosophers have the additional option of explaining away recalcitrant evidence, and alternative theories based on the same evidence are not only a theoretical possibility, but a reality.

5.3 The Role of Values

Of course, that reflective equilibria are underdetermined by their evidence base does not mean that there are no principles or heuristics that philosophers can use to guide themselves in choosing how to restore equilibrium. In response to the problem of underdetermination, both contrastive and holistic, it is quite usual to point to certain epistemic values that can act as a guide. Standard examples in the literature are that simple theories should be preferred over more complex ones, that explanatory theories should be

---

[10]Again, the interested reader could look at Jenkins Ichikawa & Steup (2012) for an introduction to the analysis of knowledge.

favored over non-explanatory ones, and that fruitful theories should be pre-
ferred over theories that have fewer applications or do not stimulate further
research.

However, we also have the intuition that epistemic practices should be
as objective as possible, that is, as free from the subjective values of its
practitioners as it can be. Furthermore, not all values that influence epistemic
practices are considered to be equally unproblematic. Choosing to accept a
theory because it aligns better with your political ideology or religious beliefs,
for instance, is considered to be significantly more problematic for the ideal of
epistemic objectivity than choosing to accept a theory because it postulates
fewer theoretical entities. It is therefore important to distinguish between two
kinds of values when we are thinking about the role of values in epistemic
practices: epistemic values and contextual values.[11]

Epistemic values are taken to be those values that indicate good scien-
tific practice. The values mentioned in the first paragraph of this section
are good examples of such epistemic values. It is now generally believed that
epistemic practices cannot be completely independent from such values.[12] At
the same time, such values are not generally considered to be a problem for
the objectivity of scientific practices. Contextual values, on the other hand,
are moral, personal, social, political and cultural values that threaten the ob-
jectivity of a scientific practice, and should, therefore, not play a constitutive

---

[11]This section has been heavily based on (Reisch & Sprenger, 2014). See that work for
more details.

[12]See, for instance, (Douglas, 2013) for a detailed discussion.

part in these practices. Our example above of accepting a theory because it aligns with one's political ideology is a case in point. Similarly, one should not choose to defend idealism over realism solely because defending a more controversial view is good for one's philosophical career.

Some philosophers argue that there is actually no sharp distinction between epistemic values and contextual values, because there are no purely epistemic values. Helen Longino (Longino, 1996), for instance, argues that epistemic values are always partly determined by contextual values, such as political and social values. One could, for example, argue that in our modern society, the fruitfulness of a theory is partly valued because we are committed to the political ideal of human progress, the desire to improve the quality of life for everyone. A society that does not aim at human progress, such as the Ancient Greeks who had a much more static world-view, will value the practical applicability of a theory to a far lesser degree then we do (as was indeed the case). I do not, however, wish to look at Longino's argument in detail here because it spells trouble for the objectivity of science in general, a topic which is beyond the scope of this dissertation. Let us therefore, for the sake of the argument, accept a sharp distinction between purely epistemic values and contextual values.

Even given this distinction, the influence of clearly contextual values on epistemic practices is not always considered to be problematic either. For instance, it seems perfectly fine for a philosopher to let her interests and what topics are in fashion determine what she is going to work on in the coming

year. Other forms of influence on epistemic practices are not considered to
be so benign.

Influence from contextual values are considered to be particularly prob-
lematic in two cases:

- in the determination of the nature of the evidence

- in theory choice

Unfortunately, we see both of these problematic influences of values on prac-
tice in philosophy.

In philosophy, contextual values play a role both in the problem of holis-
tic underdetermination of philosophical theories and in the problem of con-
trastive underdetermination.

With respect to the problem of holistic underdetermination, the issue of
what to revise when a theory is out of equilibrium, we see that philosophers
are free to change their stance towards the nature of the evidence. That
is, philosophers decide for themselves which propositions they consider to
be bedrock and what they considered to be an acceptable or unacceptable
consequence of their theory. A strong case can be made that philosophers
often exhibit a problematic preference bias when trying to resolve the problem
of holistic underdetermination; they cherry-pick intuitions that support their
favorite theories and explain away any intuitions that are in conflict with it.

With respect to contrastive underdetermination, there is the problem that
every acceptance or rejection of a theory comes with risk, since no theory is

ever confirmed or refuted beyond a reasonable doubt. Personal value comes in with respect to the amount of epistemic risk one is willing to take when one accepts or rejects a theory (Reisch & Sprenger, 2014). Although this could be conventionally institutionalised, like the p-value threshold of 0.05 in many of the social sciences, it is clear that if a society would value a different p-value, research in the social sciences would have looked different. One would think such an attitude toward risk plays a role in the statistical sciences, that depends on inductive inferences, but not in philosophy. However, in philosophy this corresponds to how revisionary people are willing to be towards their intuitions.[13]

Furthermore, apart from these two problematic cases, the subject-matter of a philosophical investigation itself might be irreducibly value laden with contextual values. For instance, it seems impossible for parts of philosophy, ethics and political philosophy in particular, to eliminate what Hilary Putnam has called *thick ethical concepts*. These are concepts that have both descriptive and normative content (see (Putnam, 2002) for more details). An example is the concept of *cruelty*. Designating someone as cruel, say a university professor, implies that she behaves in a certain way towards her students. But it also implicitly contains a value judgment that this behavior is unacceptable. There are no morally acceptable ways to act cruelly, although we might differ in opinion whether a certain way of acting was cruel

---

[13]In (Dutilh Novaes & Geerdink, 2017), Catarina Dutilh Novaes and I have argued that the distinction between conservative and revisionary attitudes towards philosophical intuitions played a particularly important role in the history of analytic philosophy.

or not. Similarly, arguing that a certain way of arranging a society is just cannot be done without also implicitly taking in the political position that we should change society to reflect it.

It therefore seems as if the acceptance of a philosophical position is not free from contextual values. As such, when a philosopher accepts a philosophical viewpoint, this says as much about what she believes on the basis of the evidence she has, as it does about the values that she holds.

5.4   Philosophical Pluralism and the Worry that Anything Goes

In the previous section, I argued that developing a reflective equilibrium is the best philosophers can do, and that there is always a plurality of such equilibria possible. No such philosophical position can ever be definitively established or refuted, and choosing among these equilibria is determined by what one values. Furthermore, engaging in philosophical analysis does not seem to lead to a reduction in defensible philosophical viewpoints. Rather, it seems that, over time, more and more defensible positions are developed.

So what is all this philosophizing good for? Isn't engaging in such an epistemic practice where no knowledge is possible just a waste of time and resources? Indeed,

> many men, under the influence of science or of practical affairs,

are inclined to doubt whether philosophy is anything better than innocent but useless trifling, hair-splitting distinctions, and controversies on matters concerning which knowledge is impossible.

(Russell, 1912, p. 89)

If we take the idea of philosophy as a game of giving and asking for reasons seriously, we can make a vivid analogy between studying all the possible moves regarding a philosophical problem and studying all the possible moves in a chess position (Dennett, 2013). Chess positions also usually allow for different possible continuations,[14] and different players might prefer different end positions because e.g. the one values a positional advantage while the other prefers material advantage. Depending on the position, finding good plays in chess might be extremely hard, just as finding good argumentative moves in philosophy can be.

However, one should note that studying a structure that is cognitively difficult is not enough to make this study worthwhile. Dennett warns us that the only reason studying chess is worthwhile is because chess matters to real people. We believe that chess itself has value. One could also study a game that looks like chess, but where the king can move two spaces rather than one. Dennett calls this game chmess. Finding good moves in chmess might be just as hard as finding good moves in chess. However, dedicating one's life to studying chmess, according to Dennett, has no merit because nobody plays this game. Quoting Donald Hebb, famous for his contribution to the

---

[14]To see this, just remember that there is no single best opening to play in chess.

development of *neural networks*, Dennett posits the principle that

> [i]f it isn't worth doing, it isn't worth doing well. (Dennett, 2013, p. 421)

The worry is now that, when philosophers investigate certain philosophical viewpoints, they are basically engaging in developing a philosophical viewpoint akin to developing the theory of chmess. For example, although it might show philosophical ingenuity, is it really worth dedicating one's life to defending a philosophical position such as solipsism, the idea that you are the only thinking thing in existence, from objections that other philosophers consider definitive? Similarly, what good does it achieve to defend the claim that there are true contradictions (Priest, 2006)? What distinguishes defending these kinds of positions from defending conspiracy theories?

There are some things to say here. First, it might be the case that some philosophical investigations are more like investigating chess, even though others are more like investigating chmess. For instance, how to conceive of causes is of interest not just to the philosopher, but also to the physicist. Similarly, many people are interested in the question of how the mind relates to the brain. How to lead a meaningful life is another example. There seems to be a real need among us humans, even if this need is only psychological, to try and answer certain philosophical questions. On the other hand, only philosophers seem to be interested in the metaphysical question of how many objects I see when I look at a cloud outside.[15]   Dennett believes that a

---

[15]"Think of a cloud—just one cloud, and around it a clear blue sky. Seen from the

broad interest in a philosophical problem is a good sign that a problem is worth exploring (Dennett, 2013, p. 421). Even if knowledge on philosophical problems is ultimately impossible, excellent philosophy that attempts to deal with problems of general interest might have value similar to a good novel or movie. It elevates the minds of many. This consideration alone should lead us to try to make good philosophy as accessible as possible (something Dennett himself is, of course, dedicated to).

Although I believe that the above considerations are sound, it is also important to note that the analogy we have been using is slightly misleading. As we have seen in our discussion of philosophy as a game of giving and asking for reasons, the rules of the game are not fixed. A better analogy would be to think of philosophical investigation as trying to find out whether a certain game is interesting to play or not. We also did not know which variant of chess was the best to play *a priori* from inspecting the rules. This only became apparent after the game was played for a certain time. For instance, the rule that a pawn can advance two square on its first move was only added

---

ground, the cloud may seem to have a sharp boundary. Not so. The cloud is a swarm of water droplets. At the outskirts of the cloud, the density of the droplets falls off. Eventually they are so few and far between that we may hesitate to say that the outlying droplets are still part of the cloud at all; perhaps we might better say only that they are near the cloud. But the transition is gradual. Many surfaces are equally good candidates to be the boundary of the cloud. Therefore many aggregates of droplets, some more inclusive and some less inclusive (and some inclusive in different ways than others), are equally good candidates to be the cloud. Since they have equal claim, how can we say that the cloud is one of these aggregates rather than another? But if all of them count as clouds, then we have many clouds rather than one. And if none of them count, each one being ruled out because of the competition from the others, then we have no cloud. How is it, then, that we have just one cloud?" (Lewis, 1993, p. 164)

in the 15th century, roughly a thousand years after the first versions of what is recognizably chess were played in India, as was the *en passant* capturing rule. Similarly, the rule that a game ends in a draw if, after fifty moves, no piece has been captured and no pawn has been moved has been modified several times in the 20th century after certain mating nets were discovered with the help of computers that took more than a hundred moves.[16] Similarly, it might take creativity and persistence to find philosophical problems that are of broad interest. Promising looking research programs might very well end up to have been for naught, and programs that nobody found interesting when they were carried out might turn out to define the philosophical agenda for decades.[17]

Dennett, of course, agrees with these points. Since it is impossible to discover *a priori* which problems can fruitfully be explored and which cannot, Dennett exclaims

> [. . .] let a thousand flowers bloom, I say. But just remember: if
> you let a thousand flowers bloom, count on 995 of them to wilt.
> (Dennett, 2013, p. 421)

---

[16]Since such positions are extremely rare in real play and often involve seemingly random moves that defy human comprehension the FIDE fully reinstated the 50 move rule in 2001.

[17]A case in point is Frege's investigation into the foundation of mathematics. In his own time, Frege lamented that his work mostly received unfavorable reviews, and he died believing that his work was not appreciated. Luckily, Frege seems to have been convinced of the value of his own work. In bequeathing his work to his adopted son Alfred, he wrote: "I believe there are things here which will one day be prized much more highly than they are now. Take care that nothing gets lost" (Klement, 2010, p. 861). Frege turned out to be right. Although much of his work was destroyed during a bombing raid of Münster, where Frege's *Nachlass* had ended up in the hands of Heinrich Scholz, Frege is now lauded as one of the greatest philosophical innovators of all time.

Chapter 6

Conclusion

In this dissertation, I have set out to answer the question of whether philosophy makes any progress, and if so, in what sense.

It seemed to me that such an investigation was needed, because there are strong *prima facie* reasons to doubt there is any epistemic progress in philosophy at all. We do not seem to have secured any substantial philosophical knowledge in the past centuries, and it is doubtful that our philosophical views are closer to the truth then those of our philosophical ancestors.

For one, we still treat our philosophical ancestors as authorities on philosophical issues that we struggle with. There are, for instance, still Aristotelians, Kantians, and Wittgensteinians. That makes it seem like we might have learned little in the intervening time. It is true that arguments and counter-arguments have sometimes become more sophisticated, but it is equally true that we often look back to the past masters for guidance. Some-

times, it looks like we know little more, philosophically speaking, than Plato or Aristotle. This is definitely not true for many other scientific disciplines. There are no Newtonian physicists, Gaussian mathematicians or Darwinian biologists,[1] and most high-school students know more about e.g. physics, mathematics, and biology than any person in the ancient world ever did. Aristotle might have been a genius, but he still believed that women had fewer teeth then men.[2]

Furthermore, the pervasive disagreement in philosophy with regards to every substantial issue is suspicious since consensus among practitioners is usually seen as an important signs of scientific knowledge. If philosophers do not agree on any substantial philosophical issue, past or present, then in what sense can we claim that philosophy has made progress in the past years? Even if some lucky individual philosopher has gained knowledge because they just happened to have defended the right view, if they couldn't transfer that knowledge by convincing their peers, then what good does that do to philosophy as a whole?

Lastly, every substantial philosophical position that has been defended in the past has turned out to be problematic (although, equally importantly,

---

[1]They are simply called biologists.

[2]Interestingly, Aristotle might have thought that men and women have a different number of teeth on the basis of a hasty induction. The standard joke that he could have known better if only he had looked in his wife's mouth does not seem to hold up. The relevant quote from *The History of Animals* reads:

> Males have more teeth than females in the case of men, sheep, goats, and swine; in the case of other animals observations have not yet been made. (501b20)

no substantial philosophical position has ever been conclusively refuted, not even solipsism). If no philosophical defense has held up in the past, then isn't that a good reason to believe that the defenses of substantial philosophical positions that we are currently developing today will prove equally unconvincing in fifty years? Why even bother trying to answer questions we can be confident to get wrong?

Although I think the above is a strong case against philosophical progress, as a philosopher myself, I also strongly feel that it is actually worthwhile to engage in philosophy. Philosophy is not a degenerative research program. It is a good thing that philosophy is still being studied and taught at universities, and, like the other humanities, this is where philosophy belongs. Of course, these two beliefs are in tension with one another. This study was the result of an attempt to relieve this tension.

As an analytic philosopher who has a soft spot for Hegel, I am pleased to find that, at the end of my study, my conclusion is a synthesis of both starting points. I do find myself convinced of *philosophical skepticism*, the view that substantial philosophical knowledge is impossible, while on the other hand I also found that philosophical methods allow for personally justified philosophical positions on questions that demand to be answered.

I have become a philosophical skeptic, because, as I have argued in this dissertation, there is simply too big of a problem with the quality of the data that philosophers start out with when they start theorizing. Intuitions, whether in the from of common sense "truths" or "truths" of meaning, cannot

play the foundational role nor the testing role that they must if substantial philosophical knowledge is to be possible. Intuition conceived of as a special sense for *a priori* truths makes its nature mysterious, and experimental evidence strongly suggests that, if intuition is conceived of as a Type 1 capacity, then that capacity isn't reliable. Furthermore, if the analogy with perception and observation is to hold up, we need some kind of causal story that links rational intuitions with *a priori* states of affairs.[3] Methodologically speaking, there are simply too few "truths" in philosophy that can't be further questioned, and there is simply too much leeway in how to deal with recalcitrant intuitions. The problem of logical pluralism that I set out in chapter four, that it is not determined what exactly follows from a set of propositions, just makes matters worse.

Philosophers, and analytic philosophers in particular, had, of course, already started to shy away from the idea that it was possible to build a philosophical system. I think this study indicates that the convergence route through hypothesis, confirmation, and refutation is closed to philosophy as a path to progress as well. Progress in philosophy does not consist of acquiring truth or convergence on truth.

What about progress as puzzle solving? The problem with this conception of progress is that, due to the underdetermination problem, there is simply no unique way in philosophy to solve a puzzle, and there are too many

---

[3]Although this is also a problem for mathematical knowledge, postulatism (or ifthenism) and the applicability of mathematics in the real world mitigate these problems in the case of mathematics.

paradigms operating simultaneously within philosophy at any given time. That is, not only do philosophers disagree on how to solve a philosophical puzzle, they also disagree on what constitutes the solution of a puzzle. In Kuhnian terms, philosophy is in a perpetual state of revolution with no periods of stable normal science in between. This is a symptom of a problem that I identified in the final chapter of this dissertation; that contextual values play a constitutive role in philosophical practice.

As I have argued, this does not mean that philosophical practice has no epistemic function. Even though knowledge of philosophical truths or, at least, doxastic justification of them is out of reach, what philosophy can aim at is personal justification of a philosophical view. In acquiring such personal justification, one not only finds out, clarifies, and revises what one believes, but also what one values. That is, the product of philosophy is not knowledge of the world, but knowledge of oneself.[4]

---

[4]Of course, in so far as one can transfer one's reasoning to others, one can teach them one way of how they can fruitfully think about certain subject-matters. Philosophers can act as guides to one another in a landscape that is difficult to traverse. But ultimately, every philosophical journey is one's own, and sometime one needs to explore new ground or go where others refuse to follow.

Chapter 7

Samenvatting in het Nederlands

In het afgelopen decennium zijn filosofen zich steeds actiever gaan bezighouden met het nadenken over de aard van de filosofie zelf. Deze tak van de filosofie wordt vaak de filosofie van de filosofie, of ook wel de metafilosofie, genoemd. Een van de belangrijkste vragen die omtrent de filosofie zelf steeds meer in het licht is komen te staan is de vraag of de filosofie wel vooruitgang maakt, en zo ja, wat voor soort vooruitgang dit dan is. In dit proefschrift staat deze vraag centraal. Om deze vraag te kunnen beantwoorden moeten we natuurlijk duidelijk maken wat we onder vooruitgang verstaan. In dit proefschrift interpreteer ik de notie van vooruitgang epistemisch; d.w.z. in termen van het vergroten van de verzameling van filosofische kennis of het rechtvaardigen van filosofische posities.

Het idee dat filosofie epistemische vooruitgang maakt, zoals dat bij andere wetenschappen het geval is, is op het eerste gezicht problematisch. Dit blijkt uit een aantal punten. Ten eerste lijkt het zo te zijn dat oude filosofen, zoals Aristoteles en Kant, nauwelijks hun filosofische positie zouden hoeven te herzien wanneer ze naar de huidige tijd zouden worden getransporteerd. Sterker nog, hun filosofische systemen vormen nog altijd de inspiratie voor vele hedendaagse filosofen. Hedendaagse debatten in de metafysica en de ethiek worden o.a. gevoerd door neo-Aristotelianen en neo-Kantianen. Ter vergelijking; er bestaan geen neo-Aristotelianen in de biologie of neo-Kantianen in de kosmologie. Hoewel Aristoteles en Kant in hun tijd grote wetenschappelijke bijdragen hebben geleverd buiten de filosofie, zijn deze wetenschappelijke ideeën heden ten dage volledig verouderd.

Een tweede probleem met het idee dat filosofie epistemische vooruitgang zou boeken is dat uit empirische onderzoek blijkt dat er nauwelijks consensus is in de filosofie over substantiële filosofische stellingen. Schertsend kan men wellicht zelfs stellen dat er onder filosofen waarschijnlijk meer overeenstemming is dat klimaatverandering plaatsvindt dan dat de wereld werkelijk bestaat in metafysische zin. Dit is een probleem omdat zo'n gebrek aan consensus onmiddellijk vertaald kan worden in een sceptisch argument. Als filosofen het onder elkaar al niet eens kunnen worden over wat metafysisch of ethisch het geval is, hoe kunnen zij dan beweren dat ze weten wat het geval is?

Een derde probleem is wat ook wel de pessimistische meta-inductie wordt

genoemd. Tot nu toe is nog geen enkele filosofische theorie uit het verleden door de filosofische gemeenschap algemeen geaccepteerd. Men zou kunnen stellen dat filosofische theorievorming zich constant in een revolutiefase heeft bevonden, om in termen van Kuhns visie op wetenschapsontwikkeling te spreken. Dit stelt pessimistisch voor de toekomst. Welke redenen hebben we om aan te nemen dat we in de toekomst uit deze toestand van permanente revolutie zouden kunnen komen? Dat zo'n doorbraak aan de horizon ligt is vaak beloofd; denk aan Kants Copernicaanse wending, de Weense nadruk op formele methoden, of de Britse belofte dat taalanalyse alle filosofische problemen zal gaan oplossen in de toekomst. Allemaal zeer interessante ont-wikkelingen binnen de filosofie, maar geen van deze wendingen heeft geleid tot de doorbraak die ons steeds was beloofd. Waarom zouden de hedendaagse beloften van neo-Carnapianen en experimentele filosofen anders zijn?

Om beter grip te krijgen op de metafilosofische problematiek rond filosofi-sche vooruitgang analyseer ik de filosofische praktijk zelf en presenteer ik een explicatie van deze praktijk in termen van Robert Brandoms spel van "giving and asking for reasons". In dit spel staat het geven van (en het vragen naar) redenen centraal.

Deze explicatie is noodzakelijk om helder te krijgen waar we het precies over hebben. Net als voor de wetenschap in het algemeen geldt ook voor de filosofie in het bijzonder een demarcatieprobleem. Er lijken geen noodzake-lijke en voldoende voorwaarden te zijn die exact definiëren wat filosofisch is en wat niet. Filosofische praktijken zijn simpelweg te heterogeen. Dat betekent

gelukkig niet dat filosofische praktijken geen familie vormen met eigenschappen die ze geneigd zijn te hebben. In mijn proefschrift bespreek ik drie belangrijke eigenschappen die typisch zijn voor filosofische praktijken: filosofen onderbouwen hun stellingen met argumenten en redenen; filosofen zijn ideaal gesproken epistemisch autonoom; en filosofische praktijken spelen zich af in een context van hoge epistemische standaarden. Ik beargumenteer dat deze drie familie-eigenschappen verklaren waarom deductieve argumenten in de filosofie een bijzondere plaats innemen. Ik gebruik vervolgens deze analyse om een vorm van filosofisch scepticisme te verdedigen, de stelling dat substantiële filosofische kennis onmogelijk is. Het bedrijven van de filosofische praktijk leidt niet tot het vergroten van onze verzameling van substantiële filosofische kennis noch tot een objectieve en universeel accepteerbare rechtvaardiging van filosofische posities. De mogelijkheid tot substantiële kennis hangt namelijk af van twee aannames die in de filosofie niet vervuld zijn: ten eerste de aanname dat we weten welke proposities logisch uit welke andere proposities volgen; ten tweede de aanname dat er een verzameling is van betrouwbare proposities die als fundament kan dienen voor een filosofisch systeem of als toetssteen van filosofische hypothesen.

De eerste aanname is in de filosofie niet vervuld omdat ontwikkelingen in de logica laten zien dat het in filosofische contexten helemaal niet onomstreden is welke proposities precies uit welke anderen volgen. Er zijn momenteel verschillende logica's die allen claimen de enige ware logica te zijn. Geen van deze logica's wordt momenteel algemeen geaccepteerd als correct binnen de

filosofie. We weten momenteel dus niet wat de juiste logica is (en misschien komen we het ook nooit te weten). Daarnaast duidt onderzoek in de filosofie van de logica er steeds meer op dat er waarschijnlijk helemaal geen uniek correcte logica is. De situatie lijkt vermoedelijk meer op de situatie in de meetkunde, waar een oneindig pluralisme van meetkunden heerst die men al dan niet aanwendt afhankelijk van pragmatische overwegingen m.b.t. het probleem waaraan men werkt.

De tweede aanname is ook niet vervuld omdat er in de filosofie geen verzameling van betrouwbare onmiddellijke proposities is die als fundament of toetssteen gebruikt kan worden. Normaal wordt deze functie aan filosofische intuïties toebedeeld. Dit zijn filosofische oordelen die onmiddellijk als juist ervaren worden, oftewel, oordelen die zelfevident zijn. Empirisch onderzoek laat echter zien dat filosofische intuïties helemaal niet betrouwbaar zijn. Er is grote onenigheid over wat zelfevident is. Daarnaast worden filosofische intuïties sterk beïnvloed door irrelevante contextuele factoren. Ook zijn filosofische intuïties epistemologisch mysterieus. Normaal gesproken wordt de betrouwbaarheid van onmiddellijke oordelen gegrondvest door een causaal verband tussen het oordeel en datgene waarover geoordeeld wordt. Observatieoordelen zijn bijvoorbeeld betrouwbaar omdat ze normaal gesproken veroorzaakt worden door datgene dat geobserveerd wordt. Hetzelfde geldt voor introspectieoordelen, die normaal gesproken in een causale relatie staan tot een innerlijke toestand. In beide gevallen gaat het over een relatie tussen een oordeel en wat het geval is. Bij filosofische intuïties gaat het echter

vaak om normatieve of modale oordelen. Het is echter mysterieus hoe er een causaal verband kan zijn tussen een onmiddellijk oordeel en wat het geval zou moeten zijn of wat het geval zou kunnen zijn.

Hoewel ik in dit proefschrift filosofisch scepticisme verdedig, neem ik toch een slag om de arm. In uitzonderlijke gevallen kan de filosofie namelijk een nieuwe wetenschap voortbrengen wanneer het lukt om een filosofische theorie wel aan een verzameling van betrouwbare onmiddellijke proposities te koppelen. Ik beargumenteer echter dat dit als bijzondere gevallen gezien moeten worden. Er is geen indicatie dat filosofische problemen in het algemeen op deze manier opgelost kunnen worden.

Nadat ik heb beargumenteerd dat de filosofische praktijk in het algemeen geen systeem van kennis kan opbouwen, noch een verzameling van filosofische hypothesen kan bevestigen of weerleggen, ga ik in op de vraag wat het epistemische product van filosofische praktijken dan wel is. In het bijzonder beargumenteer ik dat, in het ideale geval, het product van een filosofische praktijk een reflectief equilibrium is, d.w.z. een coherente verzameling van stellingen waarin de algemene principes die zijn opgesteld perfect in balans zijn met de specifieke gevallen die worden waargenomen. Men zou kunnen denken dat wanneer een filosofische positie in reflectief equilibrium is, deze positie hierdoor ook gerechtvaardigd wordt. Ik beargumenteer echter dat men hier een onderscheid moet maken tussen persoonlijke rechtvaardiging en doxastische rechtvaardiging. Iemand is persoonlijk gerechtvaardigd in het accepteren van een filosofische positie wanneer zij epistemisch gesproken nie-

ts verkeerd heeft gedaan. Wanneer een filosoof een filosofische positie in equilibrium aanhangt, dan is zij inderdaad persoonlijk gerechtvaardigd te denken wat zij denkt; haar positie is coherent en haar epistemische keuzes zijn verdedigbaar. Dit betekent echter niet dat de positie zelf ook in objectieve zin, oftewel doxastisch, gerechtvaardigd is.

Dit komt omdat reflectieve equilibria ondergedetermineerd zijn. Er zijn altijd meerdere manieren om tot een equilibrium te komen. In het algemeen gesproken hebben filosofen veel vrijheid om tot verschillende equilibria te komen omdat ze vrij zijn om intuïties of logische principes te verdedigen of weg te verklaren wanneer deze problematisch blijken te zijn. Iemand kan dus persoonlijk gerechtvaardigd zijn in het aanhangen van een filosofische positie terwijl zij het fundamenteel oneens is met een andere filosoof die ook persoonlijk gerechtvaardigd is in wat hij denkt. Ik verwijs in dit verband ook naar recent onderzoek binnen de formele filosofie. De zogenaamde onmogelijkheidsresultaten, als eerste gepresenteerd door Luv Bovens en Stephan Hartmann, zijn een sterke indicatie dat er geen verband is tussen een verzameling van proposities in equilibrium en de waarschijnlijke waarheid van deze verzameling van proposities.

Dat een reflectief equilibrium in de filosofie altijd ondergedetermineerd is laat zien dat het al dan niet accepteren van een filosofische positie altijd waardebeladen is in problematische zin. De acceptatie van een filosofische theorie gebeurt niet alleen op basis van wat de theorie zegt, maar ook op basis van wat de persoon die de positie overdenkt belangrijk vindt. Welke

proposities geaccepteerd worden als zelfevident en welke niet hangt sterk af van de persoonlijke waardering van deze proposities. Sommige filosofen zijn bereid om klassieke logische principes op te geven om een filosofische stelling te verdedigen, anderen niet. Sommige filosofen denken dat "het is moreel onacceptabel om voor je plezier bewuste wezens te martelen" een zelfevidente waarheid is, anderen niet.

Het klopt natuurlijk dat andere wetenschappelijke praktijken ook waarde-beladen zijn. Ik betoog echter dat, in tegenstelling tot andere wetenschappen, zogenaamde contextuele waarden binnen de filosofie een essentiële rol spelen. Dit soort waarden worden in de wetenschapsfilosofie over het algemeen gezien als zeer problematisch omdat ze epistemisch irrelevant geacht worden. Het wetenschappelijke ideaal om dit soort irrelevante waarden uit de epistemische praktijk te elimineren kan in de filosofie echter niet bereikt worden omdat ze een constitutieve rol spelen.

Hoewel ik in dit proefschrift voor filosofisch scepticisme argumenteer, is dit werk geen oproep om te stoppen met filosofische theorievorming. Ten eerste worden we simpelweg in de dagelijkse praktijk geconfronteerd met filosofische vragen en het ligt in de menselijke aard te proberen deze vragen te beantwoorden. Dat we door deze zoektocht soms tot nieuwe wetenschappen komen, zoals boven besproken, zou rechtvaardiging genoeg moeten zijn om met filosofische vragen te blijven worstelen. Echter, we zien ook dat we door het articuleren van filosofische problemen helderder kunnen krijgen wat we boven alles waarderen en welke denkbeelden we bereid zijn om op te geven

ten gunste van andere. Het eindproduct van filosofische praktijken is dus misschien niet kennis van de wereld, maar wel kennis van onszelf. En dit is precies de soort van kennis waar Socrates, de filosoof bij uitstek, volgens Plato naar op zoek was.

Bibliography

Adamson, P. (2014). *Classical Philosophy: A History of Philosophy Without Any Gaps, vol I.* Oxford: Oxford University Press.

Adleberg, T., Thompson, M., & Nahmias, E. (2014). "Do men and women have different philosophical intuitions? Further data". *Philosophical Psychology* 28 (5): 615–641.

Aikin, S. (2011). *Epistemology and the Regress Problem.* New York: Routledge.

Alexander, J. & Weinberg, J.M. (2007). "Analytic epistemology and experimental philosophy". *Philosophy Compass* 2 (1): 56–80.

Andow, J. (forthcoming). *What is This Thing Called Philosophical Methodology.* Routledge.

Appel, K. & Haken, W. (1989). "Every Planar Map is Four Colorable". *Contemporary Mathematics* 98.

Arras, J. (2007). "The Way We Reason Now: Reflective Equilibrium in

Bioethics". In Steinbock, B. (ed.). *The Oxford Handbook of Bioethics.* New York: Oxford University Press. 46–71.

Armstrong, D.M. (1997). *A World of States of Affairs.* New York: Cambridge University Press.

ATLAS Collaboration. (2012). "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC". *Physics Letters B* 716 (1): 1–29.

Ashby, N. (2002). Relativity and the Global Positioning System. *Physics Today* 41.

Austin, J.L. (1979). "Ifs and Cans". In Urmson, J.O. & Warnock, J. (eds.). *Philosophical Papers.* Oxford: Oxford University Press.

Ayer, A. J. (1936). *Language, Truth and Logic.* London: V. Gollancz.

Ballantyne, N. (2014). "Knockdown Arguments". *Erkenntnis* 79 (3): 525–543.

Bauer, N. (2015). *How to Do Things with Pornography.* Harvard: Harvard University Press.

Bealer, G. (1998). "Intuition and the Autonomy of Philosophy". In DePaul, M. & Ramsey, W. (eds.). *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry.* Lanham, MD: Rowman & Littlefield. 201–240.

Bealer, G. (2000). "A Theory of the A Priori". *Pacific Philosophical Quarterly* 81 (1): 1–30.

Beall, J.C. & Restall, G. (2006). *Logical Pluralism.* Oxford: Clarendon Press.

Beaney, M. (2014). "Analysis". In Zalta, E.N. (ed.). *The Stanford Encyclopedia of Philosophy* (Summer 2016 Edition). URL = `https://plato.stanford.edu/archives/sum2016/entries/analysis/`.

Bendegem, J.P. van (2012). "A Defense of Strict Finitism". *Constructivist Foundations* 7 (2): 141–149.

Beth, E.W. (1965). *The Foundations of Mathematics. A Study in the Philosophy of Science.* (2nd ed.). Amsterdam: North-Holland Publishing Company.

Biggs, S. (2011). "Abduction and Modality". *Philosophy and Phenomenological Research* 83 (2): 283–326.

Biggs, S. & Wilson, J. (2017). "The A Priority of Abduction". *Philosophical Studies* 174 (3): 735–758.

Blackford, R. & Broderick, D. (2017). *Philosophy's Future; The Problem of Philosophical Progress.* Malden, MA: Wiley Blackwell.

Block, N. (1978). "Troubles with functionalism". *Minnesota Studies in the Philosophy of Science* 9: 261–325.

Block, N. & Stalnaker, R. (1999). "Conceptual analysis, dualism, and the explanatory gap". *Philosophical Review* 108 (1): 1–46.

BonJour, L. (1985). *The Structure of Empirical Knowledge.* Cambridge, Mass.: Harvard. University Press.

BonJour, L. (2010). "The Myth of Knowledge". *Philosophical Perspectives* 24 (1): 57–83.

Bourdeau, M. (2014). "Auguste Comte". In Zalta, E.N. (ed.). *The Stanford Encyclopedia of Philosophy* (Winter 2015 Edition). URL = `https://plato.stanford.edu/archives/win2015/entries/comte/`.

Bourget, D. & Chalmers, D.J. (2014). "What do philosophers believe?" *Philosophical Studies* 170 (3): 465–500.

Bovens, L. & Hartmann, S. (2003) *Bayesian Epistemology.* Oxford: Clarendon Press.

Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment.* Harvard University Press.

Brandom, R. (2000). *Articulating Reasons: An Introduction to Inferentialism.* Harvard University Press.

Brouwer, L.E.J. (1905). *Leven, kunst en mystiek.* Delft: Waltman.

Brun, G. (2014). "Reflective Equilibrium Without Intuitions?" *Ethic Theory Moral Practice* 17: 237–252.

Brun, G. (2016). "Explication as a Method of Conceptual Re-engineering". *Erkenntnis* 81 (6): 1211–1241.

Buckwalter, W. & Stich, S. (2014). "Gender and Philosophical Intuition". In Knobe, J. & Nichols, S. (eds.). *Experimental Philosophy.* (Vol.2). Oxford: Oxford University Press.

Cappelen, H. (2012). *Philosophy Without Any Intuitions.* Oxford: Oxford University Press.

Carnap, R. (1931). "Überwindung der Metaphysik durch Logische Analyse der Sprache". *Erkenntnis* 2: 219–241.

Carroll, L. (1895). "What the Tortoise Said to Achilles". *Mind* 104 (416): 691–693.

Carnap, R. (1934). *Logische Syntax der Sprache.* Wien: Springer.

Carnap, R. (1962). *Logical Foundations of Probability* (2nd ed.). Chicago: The University of Chicago Press.

Chalmers, D.J. (2002). "Does conceivability entail possibility?" In Gendler, T.S. & Hawthorne, J. (eds.). *Conceivability and Possibility.* Oxford: Oxford University Press. 145–200.

Chalmers, D.J. (2011). "Verbal Disputes". *The Philosophical Review* 120 (4): 515–566.

Chalmers, D. J. (2015). "Why Isn't There More Progress in Philosophy?" *Philosophy* 90 (1): 3–31.

Chalmers, D.J. & Jackson, F. (2001). "Conceptual analysis and reductive explanation". *Philosophical Review* 110 (3): 315–61.

Christensen, D. (2004). *Putting Logic in its Place: Formal Constraints on Rational Belief.* Oxford: Oxford University Press.

Christensen, D. (2007). "Epistemology of Disagreement: The Good News." *Philosophical Review* 116: 187–218.

Chomsky, N. (1965). *Aspects of the Theory of Syntax.* MIT Press.

CMS Collaboration. (2012). "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC". *Physics Letters B* 716 (1): 30–61.

Colaco, D. Buckwalter, W. Stich, S., & Machery, E. (2014). "Epistemic Intuitions in Fake-Barn Thought Experiments". *Episteme* 11 (2): 199–212.

Counihan, M. (2008). *Looking for Logic in All the Wrong Places: an Investigation of Language, Literacy and Logic in Reasoning.* URL = `http://www.illc.uva.nl/Research/Publications/Dissertations/DS-2008-10.text.pdf`

Cummins, R. (1998). "Reflections on Reflective Equilibrium". In DePaul, M. & Ramsey, W. (eds.), *Rethinking Intuition; the Psychology of Intuition*

and its Role in Philosophical Inquiry. Lanham: Rowman & Littlefield. pp. 113-128.

van Dalen, D. (2001). *L.E.J. Brouwer; een biografie.* Amsterdam: Uitgeverij Bert Bakker.

Daniels, N. (1996). *Justice and Justification: Reflective Equilibrium in Theory and Practice.* New York: Cambridge University Press.

De Cruz, H. (2015). "Where Philosophical Intuitions Come From." *Australasian Journal of Philosophy* 93: 233–249.

Dehaene, S. (2011). *The Number Sense: How the Mind Creates Mathematics* (Revised and Updated Edition). Oxford: Oxford University Press.

Dennett, D.C. (1991). *Consciousness Explained.* London: Penguin

Dennett, D.C. (2013). *Intuition Pumps and Other Tools for Thinking.* W. W. Norton & Company.

Deutsch, M. (2009). "Experimental philosophy and the theory of reference". *Mind and Language* 24 (4): 445–466.

Deutsch, M. (2010). "Intuitions, counter-examples, and experimental philosophy". *Review of Philosophy and Psychology* 1 (3): 447–460.

Deutsch, M. (2015). *The Myth of the Intuitive.* Cambridge MA: The MIT Press.

Devitt, M. (2006). "Intuitions in linguistics". *British Journal for the Philosophy of Science* 57 (3): 481–513.

Dietrich, E. (2011). "There Is No Progress in Philosophy". *Essays in Philosophy* 12 (2): 329-44.

Do Carmo, J.S. (ed.) (2015). *A Companion to Naturalism.* Pelotas: Dissertatio Filosofia.

D'Oro, G. & Overgaard, S. (2017). *The Cambridge Companion to Philosophical Methodology.* Cambridge: Cambridge University Press.

Douglas, H. (2013). "The Value of Cognitive Values". *Philosophy of Science* 80: 796–806.

Douven, I. (2010). "Simulating peer disagreements". *Studies in History and Philosophy of Science Part A* 41 (2): 148–157.

Dretske, F. (1981), "The Pragmatic Dimension of Knowledge". *Philosophical Studies* 40: 363–378.

Duhem, P. (1954). *The Aim and Structure of Physical Theory.* (Wiener, P.W. trans.). Princeton: Princeton University Press.

Dutilh Novaes, C. & Reck, E. (2017). "Carnapian explication, formalisms as cognitive tools, and the paradox of adequate formalization". *Synthese* 194 (1): 195–215.

Dutilh Novaes, C. & Geerdink, L.M. (2017). "The Dissonant Origins of Analytic Philosophy: Common Sense in Philosophical Methodology". In Lapointe, S. & Pincock, C. (eds.). *Innovations in the History of Analytical Philosophy.* London: Palgrave Macmillan.

Easwaran, K. (2009). "Probabilistic proofs and transferability". *Philosophia Mathematica* 17 (3): 341–362.

Easwaran, K. (2015). "Rebutting and undercutting in mathematics". *Philosophical Perspectives* 29 (1): 146–162.

Elga, A. (2007). "Reflection and Disagreement". *Noûs* 41: 478–502.

Elgin, C.Z. (1996). *Considered Judgment.* Princeton: Princeton University Press.

Elgin, C.Z. (2005). "Non-foundationalist Epistemology. Holism, coherence, and tenability". In Steup M. & Sosa E. (eds.). *Contemporary debates in epistemology.* Malden: Blackwell. pp. 156–167.

Elkin, L. & Wheeler, G. (forthcoming). "Resolving Peer Disagreements Through Imprecise Probabilities". *Noûs.*

van Eemeren, F.H., Garssen, B., Krabbe, E.C.W., Snoeck Henkemans, F.A., Verheij, B. & Wagemans, J.H.M. (2014). *Handbook of Argumentation Theory.* Dordrecht: Springer.

Evans, J. S. B. T. (1984). "Heuristic and analytic processes in reasoning". *British Journal of Psychology* 75: 451–468.

Ewing, A.C. (1934). *Idealism: A Critical Survey.* London: Methuen.

Fahrbach, L. (2009). "Pessimistic meta-induction and the exponential growth of science". In Hieke, A. & Leitgeb, H. (eds.). *Reduction, Abstraction, Analysis.* Ontos Verlag. 11–95

Feldman, R. (2014). "Evidence of Evidence is Evidence". In Matheson, J. and Vitz, R. (eds.). *The Ethics of Belief: Individual and Social.* Oxford: Oxford University Press.

Feltz, A., Cokely, E.T., & Nelson, B. (2016). "Experimental philosophy needs to matter: Reply to Andow and Cova". *Philosophical Psychology* 29 (4): 567–569.

Feynman, R.P., Leighton, R.B., Sands, M. (1963). *The Feynman Lectures on Physics. Mainly Mechanics, Radiation, and Heat.* Reading, MA: Addison-Wesley Publishing Company, INC.

Field, H. (2009). "Pluralism in Logic". *The Review of Symbolic Logic* 2 (2): 342–359.

Fitelson, B. (2012). "Evidence of evidence is not (necessarily) evidence". Analysis 72 (1): 85–88.

Foley, R. (1993). *Working Without a Net.* Oxford: Oxford University Press.

Foley, R. (2001). *Intellectual Trust in Oneself and Others.* Cambridge: Cambridge University Press.

Forster, M. & Sober, E. (1994). "How to tell when simpler, more unified, or less ad hoc theories will provide more accurate predictions". *British Journal for the Philosophy of Science* 45 (1): 1–35.

van Fraassen, B. C. (1977). "The pragmatics of explanation". *American Philosophical Quarterly* 14 (2): 143–150.

Frances, B. (2005). "When a skeptical hypothesis is live". *Noûs* 39 (4): 559–595.

Francis, B. (2010). "The Reflective Epistemic Renegade". *Philosophy and Phenomenological Research* 81 (2): 419 – 463.

Frege, G. (1879). *Begriffschrift.* Halle: Louis Nebert.

Fricker, E. (2006). "Testimony and epistemic autonomy". In Lackey, J. & Sosa, E. (eds.). *The Epistemology of Testimony.* Oxford: Oxford University Press. 225–253.

Gasiunas, G., Barrangou, R., Horvath, P., & Siksnys, V. (2012). "Cas9–crRNA ribonucleoprotein complex mediates specific DNA cleavage for adaptive immunity in bacteria". *Pnas* 109. E2579–E2586.

Gettier, E. "Is Justified True Belief Knowledge?" *Analysis* 23 (6): 121–123.

Gödel, K. (1930). "Die Vollständigkeit der Axiome des logischen Funktionenkalküls". *Monatshefte für Mathematik*. 37 (1): 349–360.

Gödel, K. (1931). "Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme, I". *Monatshefte für Mathematik und Physik* 38 (1): 173–198.

Goldberg, S. (2009). "Reliabilism in Philosophy". *Philosophical Studies* 142: 105–117.

Golding, C. (2011). "A Conception of Philosophical Progress". *Essays in Philosophy* 12 (2): 200–223.

Goodman, N. (1955). *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.

Goodman, N. (1978). *Ways of Worldmaking*. Indianapolis: Hackett.

Hacker, P.M.S. (2009). "Philosophy: A Contribution, not to Human Knowledge, but to Human Understanding". *Royal Institute of Philosophy Supplement* 65: 129–153.

Hansson, S.O. (2013). "Defining pseudoscience and science". In Pigliucci, M. & Boudry, M. (eds.). *Philosophy of Pseudoscience. Reconsidering the demarcation problem*. Chicago: Chicago University Press. 61–77.

Hansson, S. O. (2014). "Science and Pseudo-Science". In Zalta, E.N. (ed.). *The Stanford Encyclopedia of Philosophy* (Spring 2015 Edition).

URL = `https://plato.stanford.edu/archives/spr2015/entries/pseudo-science/`.

Harman, G. (1965). "The inference to the best explanation". *Philosophical Review* 74: 88–95.

Harman, G. (1986). *Change in View: Principles of Reasoning.* Cambridge MA: The MIT Press.

Haslanger, S. (2000). "Gender and Race: (What) Are They? (What) Do We Want Them To Be?" *Noûs* 34 (1): 31–55.

Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique.* Oxford: Oxford University Press.

Hawking, S. & Mlodinow, L. (2010). *The Grand Design.* Bantam Books.

Hawley, K. (2006). "Science as a Guide to Metaphysics?" *Synthese* 149 (3): 451–470.

Hawthorne, J. (2012). "Inductive Logic". In Zalta, E.N. (ed.). *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition). URL = `https://plato.stanford.edu/archives/spr2017/entries/logic-inductive/`.

Heath, T.L. (1910). *Diophantus of Alexandria* (2nd ed.). Cambridge: Cambridge University Press.

Heilbron, J.L. (2015). *From Quintessence to Quarks; Physics a Short History.* Oxford: Oxford University Press.

Hintikka, J. (1999). "The Emperor's New Intuitions". *Journal of Philosophy* 96 (3): 127–147.

Hjortland, O. T. (2017). "Anti-Exceptionalism About Logic". *Philosophical Studies* 174 (3): 631–658.

Hodges, W. (2001). *Logic; an Introduction to Elementary Logic.* (2nd ed.). London: Penguin.

Hodgkinson, G.P., Langan-Fox J., & Sadler-Smith, E. (2008). "Intuition: a fundamental bridging construct in the behavioural sciences." *British Journal of Psychology* 99 (1): 1–27.

Howson, C. & Urbach, P. (2006). *Scientific Reasoning; The Bayesian Approach.* Chicago: Open Court.

Ichikawa, J. (2014). "Who needs intuitions? Two Experimentalist Critiques". In Booth, A. and Rowbottom, D. (eds.). *Intuitions.* Oxford: Oxford University Press. 232–256.

van Inwagen, P. (1996). "Problems In Philosophy". *Philosophical Review* 105 (2): 253–256.

van Inwagen, P. (2004). "Freedom to Break the Laws". In *Midwest Studies in Philosophy* 28: 334–350.

van Inwagen, P. (2006). *The Problem of Evil.* Oxford: Clarendon Press.

van Inwagen, P. (2009). *Metaphysics.* (3rd ed.). Boulder, CO: Westview Press.

Jaynes, E.T. (1968). "Prior Probabilities". *Institute of Electrical and Electronic Engineers Transactions on Systems Science and Cybernetics* SSC-4. 227–241

Jinek, M., Chylinski, K., Fonfara, I., Hauer, M., Doudna, J.A., and Charpentier, E. (2012). "A programmable dual-RNA-guided DNA endonuclease in adaptive bacterial immunity". *Science* 337. 816–821.

Jenkins, C.S.I. (2014). " 'Intuition', Intuition, Concepts and the A Priori". In Booth, A. and Rowbottom, D. (eds.). *Intuitions*. Oxford: Oxford University Press. 91–115.

Jenkins Ichikawa, J. & Jarvis, B.W. (2013). *The Rules of Thought*. Oxford: Oxford University Press.

Jenkins Ichikawa, J. & Steup, M. (2012). "The Analysis of Knowledge". *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition). Zalta, E.N. (ed.). URL = `https://plato.stanford.edu/archives/win2016/entries/knowledge-analysis/`.

Johnson, R. & Blair, J.A. (2006). *Logical Self-Defense*. New York: International Debate Education Association.

de Jong, W.R. & Betti, A. (2010). "The Classical Model of Science: a Millennia-old Model of Scientific Rationality". *Synthese*: 174 (2): 185–203.

Kahneman, D. (2011). *Thinking Fast and Slow*. Penguin Books.

Kant, I. (1784). "Beantwortung der Frage: Was ist Aufklärung?" In Buchenau, A. & Cassirer, B. (1913). *Immanuel Kants Werke. Band IV. Schriften von 1783–1788.* Berlin. 167–176.

Kant, I. (1787). *Kritik der Reinen Vernunft.* Riga: Johann Friedrich Hartknoch.

Kant, I. (1992). *Lectures on Logic.* (Young, J.M., Trans.). Cambridge: Cambridge University Press.

Kekes, J. (1980). *The Nature of Philosophy.* Rowman & Littlefield.

Keller, J.A. (2015). "On Knockdown Arguments". *Erkenntnis* 80 (6): 1205–1215.

Kelly, T. (2005). "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology* 1. (Gendler, T. & Hawthorne, J., eds.). Oxford: Oxford University Press. 167–196.

Kelly, T. (2014). "Evidence Can Be Permissive". In Steup, M. Turri, J. and Sosa, E. (eds.). *Contemporary Debates in Epistemology.* (2nd ed.). Oxford: Blackwell Publishers. 298–313.

Kennard, F. (2015). *Thought Experiments. Popular Thought Experiments in Philosophy, Physics, Ethics, Computer Science, & Mathematics.*

Kim, M. & Yuan, Y. (2015). "No cross-cultural differences in the Gettier car

case intuition: A replication study of Weinberg et al. 2001". *Episteme* 12 (3): 355–361.

King, N. (2011). "Disagreement: What's the Problem? Or a Good Peer Is Hard to Find." *Philosophy and Phenomenological Research* 85 (2): 249–272.

Klein, P. & Warfield, T. A. (1994). 'What Price Coherence?" *Analysis* 54: 129–132.

Klement, K.C. (2010). "Gottlob Fege". In Moyar, D. (ed.). *The Routledge Companion to Nineteenth Century Philosophy.* Abingdon: Routledge. 858–886.

Kitcher, P. (1993). *The Advancement of Science: Science Without Legend, Objectivity Without Illusions.* Oxford: Oxford University Press.

Koriat, A. (2008). "Subjective Confidence in One's Answers: The Consensuality Principle". *Journal of Experimental Psychology: Learning, Memory, and Cognition* 34 (4): 945–959.

Kornblith, H. (2013). "Is Philosophical Knowledge Possible?" In Machuca, D.E. (ed). *Disagreement and Skepticism.* New York: Routledge. 260–276.

Kripke, S.A. (1980). *Naming and Necessity.* Cambridge, MA: Harvard University Press.

Kripke, S.A. (1982). *Wittgenstein on Rules and Private Language.* Harvard University Press.

Kuipers, T.A.F. (2001). *Structures in Science. Heuristic Patterns Based on Cognitive Sctructures.* Dordrecht: Kluwer Academic Publishers.

Kuhn, T. (1962). *The Structure of Scientific Revolutions.* Chicago: University of Chicago Press.

Kusch, M., (1995). *Psychologism: A Case Study in the Sociology of Philosophical Knowledge.* London: Routledge.

Lackey, J. (2010a). "A Justificationalist View of Disagreement's Epistemic Significance." In *Social Epistemology.* (Haddock, A., Millar, A. & Pritchard, D., eds.). Oxford: Oxford University Press. 298–325.

Lackey, J. (2010b). "What Should We Do When We Disagree?" In *Oxford Studies in Epistemology.* (T. Szabo Gendler, T. & Hawthorne, J., eds.). Oxford: Oxford University Press. 274–293.

Ladyman, J. & Ross, D. (2007). *Everything Must Go: Metaphysics Naturalized.* Oxford: Oxford University Press.

Lakatos, I. (1970). "Falsification and the Methodology of Scientific Research Programmes." In Lakotos, I. & Musgrave, A. (eds.). *Criticism and the Growth of Knowledge.* Cambridge: Cambridge University Press.

Lange, M. (2002). "Baseball, Pessimistic Inductions and the Turnover Fallacy". *Analysis* 62: 281–285.

Laudan, L. (1981). "A Confutation of Convergent Realism". *Philosophy of Science* 48: 19–48.

Laudan, L. (1983). "The Demise of the Demarcation Problem". In Cohen, R.S. & Laudan, L. (eds.). *Physics, Philosophy and Psychoanalysis: Essays in Honor of Adolf Grünbaum*. D. Reidel. 111–127.

Lehrer, K. (2000). *Theory of Knowledge* (2nd ed.). Boulder: Westview Press.

Lehrer, K. (2003). "Coherence, Circularity and Consistency: Lehrer Replies". In Olsson, E.J. (ed.). *The Epistemology of Keith Lehrer*. Dordrecht: Kluwer. 309–356.

Leitgeb, H. (2014). "The Stability Theory of Belief". *Philosophical Review* 123 (2): 131–171.

Lewis, C.I. (1946). *An Analysis of Knowledge and Valuation*. LaSalle: Open Court.

Lewis, D.K. (1979). "Scorekeeping in a language game". *Journal of Philosophical Logic* 8 (1): 339–359.

Lewis, D.K. (1983). *Philosophical Papers*. Oxford: Oxford University Press.

Lewis, D.K. (1986). *On The Plurality of Worlds*. Oxford: Blackwell.

Lewis, D.K. (1993). "Many, but Almost One". In Bacon, J. (ed.). *Ontology, Causality and Mind: Essays in Honour of D M Armstrong.* New York: Cambridge University Press.

Lewis, D.K. (1996). "Elusive Knowledge". *Australasian Journal of Philosophy* 74(4): 549–567.

Littlejohn, C. (2009). "The Externalist's Demon". *Canadian Journal of Philosophy* 39 (3): 399–434.

Longino, H. (1996). "Cognitive and Non-Cognitive Values in Science: Rethinking the Dichotomy". In Nelson, L.H. & Nelson, J. (eds). *Feminism, Science and the Philosophy of Science.* Dordrecht: Kluwer. pp. 39–58.

Losee, J. (2003). *Theories of Scientific Progress.* New York: Routledge.

Love, A.C. (2013). "Experiments, Intuitions and Images of Philosophy and Science". *Analysis* 73 (4): 785–797.

MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and Its Applications.* Oxford: Oxford University Press.

Machery, E., Mallon, R., Nichols, S., & Stich, S.P. (2004). "Semantics, cross-cultural style". *Cognition* 92 (3): 1–12.

Mackenzie, M.M. (1988). "Heraclitus and the Art of Paradox". *Oxford Studies in Ancient Philosophy* 6: 1–37.

Mackonis, A. (2013). "Inference to the best explanation, coherence and other explanatory virtues." *Synthese* 190 (6): 975–995.

Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong.* Harmondsworth: Penguin.

Maddy, P. (2014). *The Logical Must.* Oxford: Oxford University Press.

Magnus, P.D. & C. Callender (2004). "Realist Ennui and the Base Rate Fallacy". *Philosophy of Science* 71: 320–338.

Mahner, M. (2013). "Science and pseudoscience. How to demarcate after the (alleged) demise of the demarcation problem." In Pigliucci, M. & Boudry, M. (eds.). *Philosophy of Pseudoscience. Reconsidering the demarcation problem.* Chicago: Chicago University Press. 29–43.

Makinson, D.C. (1965). "The Paradox of the Preface." *Analysis* 25 (6): 205–207.

Mancosu, P. (2011). "Explanation in Mathematics". In Zalta, E.N. (ed.). *The Stanford Encyclopedia of Philosophy* (Summer 2015 Edition). URL = `https://plato.stanford.edu/archives/sum2015/entries/mathematics-explanation/`.

Marx, K. (1888). *Karl Marx über Feuerbach.* In Engels, F. *Ludwig Feuerbach und der Ausgang der Klassischen Deutschen Philosophie.* Stuttgart: Verlag von J.H.W. Dietz.

Matheson, J. (2009). "Conciliatory Views of Disagreement and Higher-Order Evidence". *Episteme* 6 (3): 269–279.

Matheson, J. (2015). "Disagreement and Epistemic Peers". In *Oxford Handbooks Online*. URL = `http://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199935314.001.0001/oxfordhb-9780199935314-e-13`.

McCauley, R.N. (2011). *Why Religion is Natural and Science is Not*. Oxford: Oxford University Press.

Meiland, J.W. (1989). "Argument as inquiry and argument as persuasion." *Argumentation* 3: 185–196.

Merzbach, U.C. & Boyer, C.B. (2011). *A History of Mathematics*. (3rd ed.). New Jersey: John Wiley & Sons.

Moody, T. (1986). "Progress in Philosophy". *American Philosophical Quarterly* 23 (1): 35–46.

Mukherjee, S. (2011). *The Emperor of all Maladies; a biography of cancer*. London: Fourth Estate.

Müller, F. (2015). "The Pessimistic Meta-induction: Obsolete Through Scientific Progress?" *International Studies in the Philosophy of Science* 29 (4): 393–412.

Müller-Lyer, F.C. (1889). "Optische Urteilstäuschungen". *Archiv für Physiologie Suppl.* 263–270.

Nado, J. (2016). "The intuition deniers". *Philosophical Studies* 173 (3): 781–800.

Nagel, J. (2014). "Intuition, Reflection, and the Command of Knowledge". *Aristotelian Society Supplementary Volume* 88 (1): 219–241.

Nagel, J., Valerie San Juan, V., & Mar, R.A. (2013). "Lay Denial of Knowledge for Justified True Beliefs". *Cognition* 129: 652–661.

Neale, S. (1990). *Descriptions.* Cambridge: The MIT Press.

Open Science Collaboration (2015). "Estimating the reproducibility of Psychological Science". *Science* 349 (6251): aac4716.

O'Connor, P.T. (2010). *Woe is I. The Grammarphobe's Guide to Better English in Plain English.* (3rd ed.). Riverhead Books.

Olsson, E. J. (2005). *Against Coherence: Truth, Probability, and Justification.* Oxford: Clarendon Press.

Overgaard, S., Gilbert, P. & Burwood, S. (2013). *An Introduction to Metaphilosophy.* Cambridge: Cambridge University Press.

Park, S. (2011). "A Confutation of the Pessimistic Induction". *Journal for General Philosophy of Science* 42 (1): 75–84.

Peijnenburg, J. (2007). "Infinitism Regained". *Mind* 116 (463): 597–602.

Pereleman, C. & Olbrechts-Tyteca, L. (1969). *The New Rhetoric: A Treatise on Argumentation.* Notre Dame: University of Notre Dame Press.

Pew Research Center for the People & the Press. (2009). *Public Praises Science; Scientists Fault Public, Media.* URL = `http://www.people-press.org/files/legacy-pdf/528.pdf`.

Pigliucci, M. (2013). "The demarcation problem. A (belated) response to Laudan". In Pigliucci, M. & Boudry, M. (eds.). *Philosophy of Pseudoscience. Reconsidering the demarcation problem.* Chicago: Chicago University Press. 9–28.

Plantinga, A. (1976). "Actualism and Possible Worlds". *Theoria* 42: 139–60.

Poincaré, H. (1905). *Science and Hypothesis.* New York: The Walter Scott Publishing Co., LTD.

Pollock, J.L. & Cruz, J. (1999). *Contemporary Theories of Knowledge.* (2nd ed.). Lanham: Rowman & Littlefield Publishers, INC.

Potter, M. (2004). *Set Theory and its Philosophy: A Critical Introduction.* Oxford: Oxford University Press.

Prawitz, D. (2005). "Logical Consequence: A Constructivist View". In Shapiro, S. (ed.). *The Oxford Handbook of Philosophy of Mathematics and Logic.* Oxford: Oxford University Press. 671–695.

Priest, G. (2006). *In Contradiction: A Study of the Transconsistent.* Oxford: Oxford University Press.

Priest, G. (2014). "Revising logic". In Rush, P. (ed.). *The Metaphysics of Logic.* Cambridge: Cambridge University Press. 211–223.

Pritchard, D. (2005). *Epistemic Luck.* Oxford: Oxford University Press.

Pritchard, D. (2014). "The Modal Account of Luck". *Metaphilosophy* 45 (4-5): 594–619.

Psillos, S. (1999). *Scientific Realism: How Science Tracks Truth.* London: Routledge.

Pust, J. (2004). "On Explaining Knowledge of Necessity". *Dialectica* 58 (1): 71–87.

Pust, J. (2017). "Intuition". In Zalta, E.D. (ed.). *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition). URL = `https://plato.stanford.edu/archives/sum2017/entries/intuition/`.

Putnam, H. (1975). *Methamatics, Matter and Method.* Cambridge: Cambridge University Press.

Putnam, H. (1981). *Reason, Truth, and History.* Cambridge: Cambridge University Press.

Putnam, H. (1992). *Renewing Philosophy.* Harvard: Harvard University Press.

Putnam, H. (2002). *The Collapse of the Fact/Value Dichotomy and Other Essays.* Cambridge, MA: Harvard University Press.

Quine, W. V. O. (1951). "Two Dogmas of Empiricism". In *From a Logical Point of View.* Cambridge, MA: Harvard University Press. pp. 20–46.

Quine, W.v.O.(1960). *Word and Object.* Cambridge MA: The MIT Press.

Quine, W.v.O. (1975). "A Letter to Mr. Ostermann". In Føllesdal, D. & Quine, D.B. (2008). *Quine in Dialogue.* Cambridge MA: Harvard University Press.

Rayo, A. (2013). *The Construction of Logical Space.* Oxford: Oxford University Press.

Rawls, J. (1971). *A Theory of Justice.* Cambridge, MA: Harvard University Press.

Rawls, J. (1993). *Political Liberalism.* New York: Columbia University Press.

Rawls, J. (1999). "The independence of moral theory". In Rawls, J. *Collected papers.* Cambridge, MA: Harvard University Press. pp 286–302.

Reiss, J. & Sprenger, J. (2014). "Scientific Objectivity". In Zalta, E.N. (ed.). *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition). URL = `https://plato.stanford.edu/archives/sum2017/entries/scientific-objectivity/`.

Rescher, N. (2006). *Philosophical Dialectics: An Essay on Metaphilosophy.* State University of New York Press.

Restall, G. (2005). "Multiple Conclusions". In Hájek, Valdés-Villanueva, L., & Dag Westerståhl (eds.). *Logic, Methodology and Philosophy of Science.* College Publications.

Rorty, R. (1979). *Philosophy and the Mirror of Nature.* Princeton University Press.

Russell, B. (1903). *The Principles of Mathematics.* Cambridge University Press.

Russell, B. (1912). *The Problems of Philosophy.* Oxford: Oxford University Press.

Russell, G. (2008). "One True Logic?". *Journal of Philosophical Logic* 37 (6): 593–611.

Russell, G. (2013). "Logical Pluralism". In Zalta, E.N. (ed.). *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition). URL = `https://plato.stanford.edu/archives/win2016/entries/logical-pluralism/`.

Ryle, G. (1953). "Ordinary Language". *The Philosophical Review* 62 (2): 167–186.

Rysiew, P. (2016). "Epistemic Contextualism". In Zalta, E.N. (ed.). *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition).

URL = https://plato.stanford.edu/archives/win2016/entries/
contextualism-epistemology/.

Saatsi, J. (2005). "On the pessimistic induction and two fallacies". *Philosophy of Science* 72 (5): 1088–1098.

Sartwell, C. (1992). "Why knowledge is merely true belief." *Journal of Philosophy* 89 (4): 167–180.

Schiffer, S.R. (2003). *The Things We Mean.* Oxford: Oxford University Press.

Scholz, H. (1930). "Die Axiomatik der Alten". *Blätter für Deutsche Philosophie* IV: 259–278.

Searle, J.R. (1980). "Minds, Brains and Programs". *Behavioral and Brain Sciences* 3 (3): 417–457.

Searle, J. R. (2016). "The Future of Philosophy". *Nova et Vetera* 14 (2): 543–558.

Sellars, W. (1963). "Philosophy and the Scientific Image of Man". In Colodny, R. (ed.). *Science, Perception, and Reality.* Ridgeview. pp. 35–78.

Shapiro, S. (2005). "Logical Consequence, Proof Theory, and Model Theory". In Shapiro, S. (ed.). *The Oxford Handbook of Philosophy of Mathematics and Logic.* Oxford: Oxford University Press. 651–670.

Shapiro, S. (2014). *Varieties of Logic.* Oxford: Oxford University Press.

Sider, T. (2009). "Ontological realism". In Chalmers, D.J., Manley, D., & Wasserman, R. (eds.). *Metametaphysics: New Essays on the Foundations of Ontology.* Oxford: Oxford University Press. 384–423.

Singer, P. (2009). *The Life You Can Save: Acting Now to End World Poverty.* New York: Random House.

Smith, P. (2003). *An Introduction to Formal Logic.* Cambridge: Cambridge University Press.

Smith, T. (2007). *Ayn Rand's Normative Ethics; The Virtuous Egoist.* Cambridge: Cambridge University Press.

Soames, S. (2010). *Philosophy of Language.* Princeton: Princeton University Press.

Sosa, E. (1998). "Minimal Intuition". In DePaul, M. & Ramsey, W. (eds.). *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry.* Lanham, MD: Rowman & Littlefield. 257–269.

Sosa, E. (2007). "Experimental Philosophy and Philosophical Intuition". *Philosophical Studies* 132 (1): 99–107.

Stalnaker, R. (2002). "Common ground". *Linguistics and Philosophy* 25 (5-6): 701–721.

Stanford, K. (2013). "Underdetermination of Scientific Theory". In Zalta, E.N. (ed.). *The Stanford Encyclopedia of Philosophy* (Spring 2016

Edition). URL = `https://plato.stanford.edu/archives/spr2016/`
`entries/scientific-underdetermination/`.

Stich, S. (1990). "Building belief: Some queries about representation, indication, and function". *Philosophy and Phenomenological Research* 50 (4): 801–806.

Streumer, B. (2017). *Unbelievable Errors: An Error Theory About All Normative Judgements*. Oxford: Oxford University Press.

Talbott, W. (2008). "Bayesian Epistemology". In Zalta, E.N. (ed.). *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition). URL = `https://plato.stanford.edu/archives/win2016/entries/epistemology-bayesian/`.

Tartaglia, J. (2011). "Philosophy between Religion and Science". *Essays in Philosophy* 12 (2): 224–241.

Thagard, P. (2007). "Coherence, Truth, and the Development of Scientific Knowledge". *Philosophy of Science* 74: 28–47.

Thagard, P. (2009). "Why Cognitive Science Needs Philosophy and Vice Versa". *Topics in Cognitive Science*, 1 (2): 237–254.

Thomasson, A.L. (2015). *Ontology Made Easy*. Oxford: Oxford University Press.

Thomson, L. & Guthrie Tait, P. (1867). *Treatise on Natural Philosophy.* Oxford: Clarendon Press.

Thomson, J.J. (1985). "The Trolley Problem". *Yale Law Journal* 94: 1395–1415.

Vaesen, K., Peterson, M., & van Bezooijen, B. (2013). "The Reliability of Armchair Intuitions". *Metaphilosophy* 44 (5): 559–578.

Walton, D. (1997). *Appeal to Expert Opinion: Arguments From Authority.* Pennsylvania State University Press.

Weber, Z. (2011). "Issue Introduction". *Essays in Philosophy* 12 (2). 195–199.

Weinberg, J.M., Nichols, S., & Stich, S. (2001). "Normativity and epistemic intuitions". *Philosophical Topics* 29 (1-2): 429–460.

Whitehead, A.N. & Russell, B. (1910). *Principia Mathematica.* Cambridge: Cambridge University Press.

Wiles, A. (1995). "Modular elliptic curves and Fermat's Last Theorem". *Annals of Mathematics* 142 (3): 443–551.

Williamson, T. (1994). *Vagueness.* Routledge.

Williamson, T. (2007). *The Philosophy of Philosophy.* Malden MA: Blackwell Publishing.

Williamson, T. (2010). *Knowledge and its Limits.* Oxford: Oxford University Press.

Williamson, T. (2015). "Semantic paradoxes and abductive methdology". In Armour-Garb, B. (ed.). *The Relevance of the Liar.* OUP, Oxford. Forthcoming.

Williamson, T. (2016). "Abductive Philosophy". *Philosophical Forum* 47 (3-4): 263–280.

Wittgenstein, L. (1922). *Tractatus Logico-Philosophicus.* London: Kegan Paul, Trench, Trubner & Co., LTD.

Wittgenstein, L. (1953). *Logical Investigations.* (Anscombe, G.E. (trans.)). Oxford: Blackwell.

Worrall J. (2009). *Miracles, pessimism and scientific realism.* URL = `https://philpapers.org/archive/WORMPA.pdf`

Worsnip, A. (2014). "Disagreement about Disagreement? What Disagreement about Disagreement?" *Philosophers' Imprint* 14 (18): 1–20.

Zhang, J., Liu, H., Luo, S., Lu, Z., Chávez-Badiola, A., Liu, Z., Yang, M., Merhi, Z., Silber, S.J., Munné, S., Konstantinidis, M., Wells, D., Tang, J.J., & Huang, T. (2017). "Live birth derived from oocyte spindle transfer to prevent mitochondrial disease". *Reproductive Biomedicine Online* 34 (4): 361–368.

Zamora Bonilla, J. (2006). "Science as a Persuasion Game: An Inferentialist Approach" *Episteme* 2 (3): 189–201.

Curriculum Vitae

Leon Geerdink was born in Gouda on the 9th of June 1981. He finished his Bachelor's (2010) and Master's (2012) degree *cum laude* at the University of Utrecht, under the supervision of Albert Visser and Paul Ziche. In 2012, he moved to the University of Groningen to work as a PhD candidate in Catarina Dutilh Novaes' NWO project *the Roots of Deduction*. In 2015, he went to Munich for a semester to visit the Munich Center for Mathematical Philosophy at the Ludwig Maximilian University. In 2017, he left academia to work in the tech industry. He now lives in Munich with his wife and daughter, where he works as a Data Scientist.