PHILOSOPHY AND THE AIM FOR TRUTH

Tom Kaspers

A Thesis Submitted for the Degree of PhD
at the
University of St Andrews

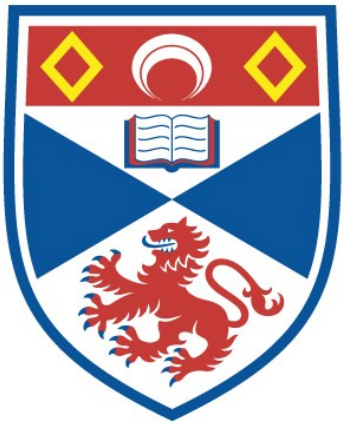

2023

Philosophy and the Aim for Truth

Tom Kaspers

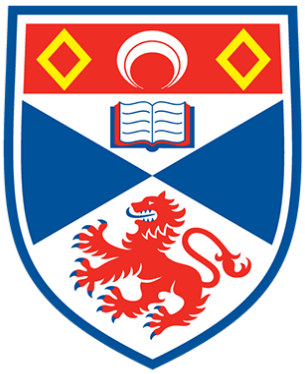

University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of

Doctor of Philosophy (PhD)

at the University of St Andrews

February 2023

Candidate's declaration

I, Tom Kaspers, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 70,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree. I confirm that any appendices included in my thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

I was admitted as a research student at the University of St Andrews in September 2019.

I confirm that no funding was received for this work.

Date 12 February 2023 Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree. I confirm that any appendices included in the thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

Date 12 February 2023 Signature of supervisor

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Tom Kaspers, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date 12 February 2023 Signature of candidate

Date 12 February 2023 Signature of supervisor

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Tom Kaspers, hereby certify that no requirements to deposit original research data or digital outputs apply to this thesis and that, where appropriate, secondary data used have been referenced in the full text of my thesis.

Date 12 February 2023 Signature of candidate

Philosophy and the Aim for Truth

Abstract

This thesis seeks to answer the question: Does philosophy aim for truth? To address this question, I must first ask a more general question: What is it to aim for truth? My approach to answering this latter question uses *pragmatism* as its frame of reference by exploring the practical significance of the aim for truth. Huw Price argues that this aim provides a common purpose that helps us to engage with each other, turning our inquiry into a shared enterprise. (Price 2003) Price's theory gives an apt description of the practical role of aiming for truth. However, I argue that this aim, conceived as such, is not advantageous to all of our inquiries. This thesis formulates a distinction between *shared* inquiry that aims for truth — such as scientific inquiry — and *personal* inquiry that lacks this aim — such as the inquiry into matters of taste. While this distinction is practical, it has many theoretical reverberations that this thesis tries to trace. It cuts across many topics of philosophical concern, such as the mechanics of disagreement, the nature of truth, objectivity, rationality, and knowledge. After dressing up the practical distinction between personal and shared inquiry in these theoretical layers, I apply it to the inquiry to which this thesis belongs. I argue against the popular belief that philosophical inquiry is scientific and thus shared. An analysis of philosophical practices shows that philosophical inquiry is personal. As a result, my thesis concludes that *philosophy does not aim for truth*.

Acknowledgments

I am incredibly grateful to a great many people. To my fellow students in St Andrews and, more recently, in New York, with whom I have discussed the topics of this thesis endlessly: Duccio Calosi, Enrico Galvagni, Ilja Kaspers, Jacob LiBrizzi, Jordina Francès de Mas, Lauren Somers, Luca Stroppa, Matt Moore, Mattia Cecchinato, Nick Allen, Niccolò Pisano, Nicola Bonatti, Soren Schlassa, and Yoichi Kobe. To Francesco Berto, who helped me pick my research topic as my academic advisor during the MLitt. To Andrew Howat, who taught me a great deal about pragmatism. To Kevin Scharp, who served as my secondary supervisor for the first half of my PhD. To Gillian Russell, who served as my secondary supervisor for the second half. And, of course, I am extremely indebted to Crispin Wright, my primary supervisor and the supervisor of my MLitt thesis before that. His guidance and enormous breadth of expertise gave me the confidence to try to answer big philosophical questions. Every single one of our meetings has been incredibly rewarding and intellectually stimulating.

I want to thank my partner, Sophie Lafayette, who came with me to St Andrews and got a backbreaking job as a nurse assistant at the local hospital, through the pandemic no less, while she probably could've gotten a better job or the opportunity to continue working on her credentials, outside of Scotland. She, more than anyone, turned my time in St Andrews into a fond and treasured memory. Lastly, I want to thank my parents, who have helped me in every possible way. I can be tough on myself, but they have always been proud of me and have thereby allowed me to be proud of myself as well. Regardless of what is to come, I am proud of the achievement that is this thesis. Therefore, it is only fitting that I dedicate this thesis to them.

Table of Contents

'Tis not solely in poetry and music, we must follow our taste and sentiment, but likewise in philosophy.

David Hume[1]

[1] In *The Letters of David Hume, Vol. 1: 1727-1765* edited by J. Y. T. Greig (1932). Oxford: Oxford University Press: p. 113.

1. Introduction

1. Research Question

The research question of this thesis is as follows:

Research question: Does philosophy aim for truth?

My thesis argues for a negative answer to this question: no philosophy does not aim for truth. One of the upshots of this result is that philosophy is less like scientific inquiry than popular belief will have it and that the attempts to emulate scientific practices are misguided.

Therefore, this thesis is not scientific. If it were, I would continue with a neatly outlined methodology for answering my research question. The methodology section of a scientific thesis is supposed to be neutral ground; it serves as a way for other scientists to check whether the research satisfies the standards of scientific inquiry. Yet, this is a philosophical thesis, and there is no neutral ground in philosophy. The question, "What are the standards of philosophical inquiry?", is itself a philosophical question. To say anything about how philosophy should be carried out is to say something philosophical. This is particularly pertinent to my thesis because it concerns itself with the nature of philosophical inquiry and thus with how philosophy should be done. As a result, while this thesis contains 200 or so pages of *philosophy*, it is only at the very end, in the conclusion, that I can say exactly what I have been doing, for it is only then that I say what philosophical inquiry is like, and perhaps more significantly, what it is *not* like. So, I cannot start with the methodology. In some sense, the methodology *is* the result. It shall appear in the conclusion, like the reveal of a magic trick: look, this is what I've been doing all along!

Nevertheless, I must start with some methodological background commitments. Every research project needs an angle. To answer my research question, I must analyze what it is to aim for truth. My thesis *could have been* a thesis in which I develop *metaphysical* arguments for why truth is correspondence and why aiming for truth amounts to aiming for the accurate representation of reality. If it were such a thesis, I could give further metaphysical arguments to show that philosophical properties do not exist, or at least that they do not exist in such a way that we can discover and represent them. Finally, I could have argued that to avoid an error-theoretic interpretation of philosophical inquiry, we must give up on the idea that philosophical inquiry aims for truth at all.

Yet, this is not how I shall approach these questions. Instead, when asking what truth is and what it is to aim for truth, I shall ask what use we make of truth *in practice* and what *practical* difference it makes to aim for truth. My mode of inquiry is thus a *pragmatist* mode of inquiry. Hence, before I proceed, I need to say some things to qualify this starting point that is my commitment to pragmatism.

What makes matters difficult for me is that this thesis ends up defending a couple of views that are usually thought of as being in opposition to pragmatism. I have already mentioned one of these views, namely the rejection of the idea that philosophy must be scientific. The other controversial view is the idea that truth is correspondence. Given that pragmatism is so interwoven with the rejection of the correspondence theory of truth as well as with the idea that we should apply the scientific method to philosophy, I will need to say a lot in the section below to disentangle pragmatism from these views, so please bear with me.

2. Pragmatism

The pragmatism that I endorse is, in its very essence, not a position on truth or metaphysics but a metaphilosophical principle. The principle that I have in mind is closely related to Peirce's Pragmatic Maxim:

> Consider what effects, which might conceivably have practical bearings, we conceive the object of our conception to have. Then, our conception of these effects is the whole of our conception of the object. (Peirce 1878, 293)

However, this maxim is broader and more contentious than I would like. It is too broad because it does not concern philosophical inquiry specifically. It is contentious because it is rather strongly put. Sure, a pragmatist must believe that conceptual distinctions without practical bearings — such as, perhaps, the metaphysical distinction between sparse and abundant entities — are somehow deficient, underdeveloped, confused, inconsequential, empty, or otherwise lacking. Yet, this does not mean that if something has no practical bearings, we cannot conceive of it at all.

When the above quote is seen in its proper context, it becomes clear that Peirce himself does not really believe so either. Peirce argues that there are three grades of clarity. The first two grades he derives from Descartes's criteria of clearness and distinctness: concepts seem clear to us if we have a firm unreflective grasp of them, and they are distinct if they "have nothing unclear about them." (Peirce 1878, 287) Peirce argues that the Leibnizian ideal for the second grade of clarity is "the clear apprehension of everything contained in the definition." (Peirce 1878, 288) The classical Cartesian and Leibnizian conception of philosophical analysis characterizes it as the project of gaining this second grade of clarity. Yet, Peirce argues there is a third, pragmatist, grade of clarity, and the quote above is the formulation of this grade. Thus, it is possible to grasp concepts without understanding their practical bearings, but it is

the philosopher's task to elevate this understanding; philosophical analysis is a *pragmatic elucidation* of our concepts.

I believe this thought is captured by the following formulation, which I consider to be the metaphilosophical principle that is the very essence of my commitment to pragmatism:

Pragmatism: To give a philosophical analysis of a concept is to show its practical bearings.

This principle has far-reaching consequences for what is to follow. Traditional philosophical inquiry sees differences in our practices as an explanandum, i.e., as something that needs explaining in, for instance, metaphysical terms. For example, a metaphysician might try to explain the fact that we are more willing to agree to disagree when it comes to matters of taste than when it comes to mathematics by saying that mathematical entities are more *real* than aesthetic properties or that mathematical matters are *objective* while matters of taste are *subjective*. However, the pragmatist would point out that such metaphysical distinctions must themselves be explained in terms of the difference they make to our practices, so the metaphysician gets things the wrong way around. Practical distinctions do not require metaphysical underpinnings, but metaphysical distinctions *do* require practical bearings.

2.1 Pragmatist Conceptions of Truth

This thesis defends a metaphysically inflated interpretation of the idea that truth is correspondence. I argue that my commitment to pragmatism does not conflict with this metaphysical commitment. However, it does require that this metaphysical substance is explanatorily downstream from the story of how we use the concept of truth in practice. I argue, in Chapter 4, for my metaphysically inflated interpretation of correspondence by showing how

the aim for this kind of truth is practically advantageous for, what I call, *shared inquiry*, such as scientific inquiry.

Nevertheless, my acceptance of the correspondence theory of truth for shared inquiry marks a substantial difference between the theory of truth developed in this thesis and the pragmatist tradition. It might be enough to make me an outlaw in the eyes of those belonging to this tradition. As Douglas McDermid notes in his book *The Varieties of Pragmatism* (2006), one of the few things that seems to bind pragmatists of all stripes is their animosity towards the correspondence theory of truth. Cheryl Misak also expresses this thought:

> It should be clear that pragmatism, of any stripe, will be set against versions of the correspondence theory of truth, on which a statement is true if it gets right or mirrors the human-independent world. (Misak 2018, 283)

As always in philosophy, there are exceptions, so I do have allies, such as Andrew Howat (2020), who argues that a pragmatist conception of truth and the correspondence conception of truth are not at odds. However, most pragmatists remain dead set against the correspondence theory of truth.

One of the main concerns with truth as correspondence to a mind-independent world, shared by almost all pragmatists, is that it would be impossible to step outside of our own minds to compare our representations of the world to the world itself. It would be impossible to know whether our best theories actually correspond to reality. (McDermid 2006, 11) How could it ever be practically advantageous to aim for a target when it would be impossible to tell whether we have hit or missed this target?

In Chapter 4 of this thesis, I try to face this challenge head on by exposing the practical bearings of the aim for an evidence-transcendent standard. However, pragmatists usually take the above question to be rhetorical: of course there is no practical value to aiming for a target

we cannot even see. Therefore, they have advocated for replacing the correspondence theory with a notion of truth that is within our reach, such as an epistemic notion of truth, a utility-based notion of truth, or, more recently, a deflationist notion of truth. I shall briefly explain these different views below.

Misak argues for a distinction between two kinds of pragmatism: Peircean pragmatism, which she supports herself, and Jamesian pragmatism. According to Misak, the Peircean pragmatist trades the absolutist correspondence conception of truth for an epistemic notion of truth, which offers an interpretation of truth as a standard of correctness that does not transcend our practices. The Jamesian pragmatist believes that once we let go of the idea that we aim to correspond to a mind-independent reality, all that remains is the idea that we aim for beliefs that work, i.e., that are optimally practically beneficial. Misak sees this contrast emerge time and again throughout the history of pragmatism. For instance, Rorty is a Jamesian pragmatist, whereas Putnam is a Peircean pragmatist. Misak puts the distinction as follows:

> On the one side of the divide we have Rorty and his classical predecessors (James and Dewey) holding that there is no truth at which we might aim—only agreement within a community or what works for an individual or what is found to solve a problem. In some moods Rorty goes as far as claiming that truth and objectivity are nothing more than what our peers will let us get away with saying. On the other side of the divide, we have those who think of pragmatism as rejecting an ahistorical, transcendental, or metaphysical theory of truth, but nonetheless being committed to doing justice to the objective dimension of human inquiry — to the fact that those engaged in deliberation and investigation take themselves to be aiming at getting things right, avoiding mistakes, and improving their beliefs and theories. On this more objective kind of pragmatism, which emanates from [Chauncey] Wright and Peirce, the fact that our inquiries are historically situated does not entail that they lack objectivity. (Misak 2013, 3)

2.2 Jamesian Pragmatism

The Jamesian pragmatist argues against the correspondence conception of truth by arguing that we have no use for "an inert static relation" between our ideas and mind-independent reality. (James 1907, 142) If sense is to be made of the idea that we aim for truth, then truth must be seen as a human value, a practical value. The opposition to the "intellectualist" notion of truth is captured very vividly by James's contemporary F. C. S. Schiller:

> The intellectualist philosopher [...] has dreamt a wondrous dream of a truth that shall be absolutely true, self-testing, and self-dependent, icily exercising an unrestricted sway over a submissive world, whose adoration it requites with no services, and scouting as blasphemy all allusion to use or application. (Schiller 1907, 9)

James and Schiller both believe that their commitment to pragmatism makes this idea of truth, as detached from any practical values, untenable. James says:

> Our obligation to seek truth is part of our general obligation to do what pays. The payments true ideas bring are the sole why of our duty to follow them. (James 1907, 230)

James turns this idea into a kind of revisionist idea of truth: true beliefs just are those beliefs that pay. He argues that it makes no difference whether we say of a truth that "it is useful because it is true" or that "it is true because it is useful" since "these phrases mean exactly the same thing". (James 1907, 143) I call it revisionist because it clashes with some of our core intuitions about truth, including the equivalence schema: ⌜ *p* ⌝ *is true* if and only if *p.* There could be many beliefs that are merely useful fictions. In my article (2022), I give the example of a counterfactual community that believes that strawberries are poisonous, even though they aren't. If they were ever to pick strawberries, they would set a process of ecological changes

in motion that would be much too complicated for them ever to comprehend, and that would bring forth an ecological disaster that would wipe out the whole community. Thus, the belief that strawberries are poisonous would certainly pay, and by James's standards, it would thereby be true, even though strawberries are not poisonous. I take it that there are countless such examples that one could think of.

James also believes that truths are not discovered but *made* by our inquiry. (James 1907, 142) Truth, as a property, is constructed by ourselves for practical purposes. John Dewey argues for a similar meta-alethic view in *The Quest for Certainty* (1929).

More recently, Richard Rorty has defended a view on truth in a Jamesian spirit. Yet, while James equates true beliefs with those that work, Rorty argues that we should leave truth behind altogether. (Rorty 1990, 127) Truth is not something we can meaningfully aim for because this aim has no practical bearings. We aim for beliefs that serve us well, and that is the only aim that a pragmatist could account for. (Rorty 1995, 281) Rather than equating truth to beliefs that are useful (thereby violating core intuitions about the concept of truth), Rorty argues that truth is merely an expressive device, a way of showing agreement or of generalizing, e.g., "All Rorty said is true". This is the view known as deflationism. However, as we shall see in Chapter 4, not all deflationists deny that we aim for truth.

Rorty extends the Jamesian tradition to concepts in the vicinity of truth, such as knowledge, rationality, and objectivity. He treats these concepts as relics of Cartesianism — the doctrine centered around the idea that our thoughts can mirror a mind-independent reality — which a thoroughgoing pragmatism ought to dispense with. (Rorty 1979, 11) Rorty is also weary of what has been called the philosophical practice of "bifurcating" our language, which is the practice of drawing philosophical distinctions between various kinds of discourses, e.g., realist and anti-realist, representationalist and expressivist, objective and subjective, descriptive and normative, etc. (Kraut 1990, 180)

I, however, shall do precisely this in my thesis when I introduce the distinction between personal and shared inquiry. Moreover, I end up giving pragmatist interpretations of all of these concepts Rorty despises. Chapter 2 argues that shared inquiry aims for truth. Chapter 4 shows that the truth it aims for is correspondence and that shared inquiry is realist and objective. Chapter 5 then shows that though personal inquiry lacks these features, it nevertheless deals in reasons and aims for knowledge.

2.3 Peircean Pragmatism

Given my fondness for concepts such as knowledge, objectivity, and a truth we can aim for, it is tempting to categorize my views as being in line with Peircean pragmatism. After all, Misak characterizes this pragmatism in the quote above as the "more objective kind of pragmatism". (Misak 2013, 3) The Peircean pragmatist rejects James's view that we only aim for truth insofar as we aim to do what pays. Misak uses examples similar to my strawberry example above to show the inadequacy of a utility-based notion of truth, and she argues against the Rortyan deflationist by making the case that "truth plays a role in our essential cognitive or epistemic practices." (Misak 2015, 263) She then argues that a more sensible way to proceed from the rejection of the correspondence theory of truth is not to equate truth to utility but to a kind of indefeasibility; true beliefs are those that cannot be improved upon.

However, it is not immediately clear how this differs from the conception that truth is what works best. To draw this distinction, more needs to be said about the standards that give substance to the sense of *improvement* in play. If one belief counts as an improvement upon another if and only if it works better than the other belief, then clearly, the two conceptions of truth are the same.

Misak argues that though there must be some connection between beliefs that are true and those that are useful, utility should not be allowed to have free rein in determining which beliefs are true. Beliefs could be useful in all sorts of ways that have nothing to do with *the facts*. This is how Misak puts this thought:

> The pragmatist should not identify '*p* is true' with '*p* is useful'. For the belief that *p* will be useful *only if p*. [...] It is useful because of what Ramsey calls 'objective' factors [...] We have seen that Peirce tries to get at something very similar: a belief must not be settled by a method extraneous to the facts. (Misak 2015, 268)

The thought behind these claims is, I take it, that when we aim for truth, we do not merely aim for beliefs that work best, but we aim for beliefs that satisfy certain epistemic norms — beliefs that account for our experiences and respond to reasons in the neatest way. It is the satisfaction of these epistemic norms that gives the dimension of objectivity to our inquiry.

But how could this be? How could a pragmatist motivate this course of action? Why should we not just aim for beliefs that work best? A possible answer is that these additional standards ensure that we gain beliefs that work best not just here and now but in the long run. Yet, this would neither be enough to distinguish the view from Jamesian pragmatism — since the utility-based notion of truth is usually already understood as involving a sense of general utility in the long run — nor would it be enough to save the view from counterexamples such as my strawberry example.

Misak seems to suggest that these epistemic standards are *not* grounded (or, rather, not wholly grounded) by the aim for beliefs that work best. Instead, they are grounded by the aim to "get things right". (Misak 2015, 263) Yet this is just the same as saying we aim for beliefs that are true. Thus, what Misak is suggesting is that there is an intrinsic, irreducible value to having beliefs that are not merely useful but that are *true*.

This intrinsic value of truth is a posit that is not in any way less mysterious than any other philosophical doctrine that pragmatists have criticized. Misak tries to defend it by saying that *in practice*, we try to get things right; we don't just try to find beliefs that work. This may well be the case, but as it stands, this is merely an appeal to intuition. What are the practical bearings of this aim? To say, as a pragmatist, that we aim to get things right instead of merely aiming for beliefs that work, we have to show what difference this would make to our practices. This is the challenge raised by Rorty and overlooked by the Peircean pragmatist. The posit that truth is intrinsically valuable bars us from finding an answer to this question.

Whenever I discuss the aim for truth in this thesis, I shall take the Jamesian motto mentioned above as my leitmotif:

> Our obligation to seek truth is part of our general obligation to do what pays. The payments true ideas bring are the sole why of our duty to follow them. (James 1907, 230)

For example, in Chapter 2, I argue that some of our inquiries aim for truth by showing the practical advantages of this aim for these inquiries. In Chapter 4, I attempt, like the Peircean pragmatist, to do justice to the objective dimension of inquiry, but, unlike the Peircean pragmatist, I shall stay true to the above Jamesian principle. I shall argue that this principle leaves room for a notion of truth that does not reduce to utility by showing that the aim for beliefs that work best sometimes requires of us that we aim to correspond to reality.

2.3 Naturalism

The founding fathers of pragmatism tried to move away from the mainstream theistic philosophical tradition of the generation that preceded theirs, and in doing so, they drew

inspiration from the scientific developments of their time. This resulted in a commitment to *naturalism.* In the quote below, Misak sketches the cultural, scientific, and philosophical backdrop of the founding of pragmatism:

> The members of The Metaphysical Club found themselves at a critical juncture in the history of American thought. They were the first generation of philosophers to put some distance between philosophy and religion. Protestantism, in one variety or another, dominated America when the pragmatists began their careers. The college philosopher tended to be also the college minister or moral tutor and appointments in philosophy were reserved for the religiously orthodox. This was about to change. [...] [T]he scientific method was being brought to philosophy with a vengeance. In 1740 Hume titled his masterpiece *A Treatise of Human Nature: An Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects* and a century later the kind of naturalism he favored had taken root. When The Metaphysical Club started to meet, the implications of Comte's positivism and Darwin's *On the Origin of Species* were being hotly debated in universities and in drawing rooms. Science seemed to entail the abandonment of the world-view that had God and religious absolutes at its centre. (Misak 2013, ix)

Naturalism can mean a couple of things. Ontological naturalism is the rejection of supernatural entities and the view that the entities that really exist are only the ones that can be studied by science. Methodological naturalism is the acceptance of a scientific methodology for philosophical inquiry. It is the idea that philosophy must turn away from the method of *a priori* reflection and that the methods of theoretical physics and mathematics should be extended into the philosophical domains. These theses may go hand in hand, but the

pragmatists were mainly concerned with the latter thesis, not the former. In what follows, I shall assume the second interpretation of naturalism whenever I use the term.

The founders of pragmatism were themselves working scientists. William James was a renowned psychologist, Chauncey Wright lectured in mathematical physics, and Charles Sanders Peirce was employed as a geodesist, using pendulums to measure the local strength of the gravitational field. They did not see their philosophical activities as belonging to a different domain of inquiry but merely as a different subfield belonging to the same domain.

Naturalism might have grown to be more popular than pragmatism. It became a mainstream view outside of pragmatism due to Quine's received influence on philosophical methodology, and especially his metaphysical methodology. Contemporary metaphysicians see Quine as their hero who saved metaphysics from the attacks of Carnap and other logical empiricists.

Carnap was also a naturalist, but he believed metaphysics to be anti-scientific. He defended the position that metaphysical theorizing has less to do with reality than with the practical utility of conceptual frameworks. According to Carnap, science operates within a particular conceptual framework to make object-level claims about reality. Metaphysics, on the other hand, operates outside of any particular conceptual framework and, instead of making object-level claims about reality, it deals in meta-level claims about how best to conceive of reality, i.e., about which conceptual framework would serve us best. (Carnap 1959) Quine argued that this distinction between science and metaphysics is untenable since it mistakenly presupposes that a neat distinction could be drawn between synthetic claims — claims about the world — and analytic claims — claims about our concepts. He maintained instead that metaphysics must be on a par with science. Hence a breed of scientific metaphysics was born.

A watered-down version of this naturalistic "Quinean" methodology is described by the contemporary metaphysician Theodore Sider, whose work focuses on "joint-carving" and fundamentality:

> [A]s a general epistemology of metaphysics I prefer the vague, vaguely Quinean, thought that metaphysics is continuous with science. We employ many of the same criteria […] for theory choice within metaphysics that we employ outside of metaphysics. Admittedly, those criteria give less clear guidance in metaphysics than elsewhere; but there's no harm in following this argument where it leads: metaphysical inquiry is by its nature comparatively speculative and uncertain. (Sider 2011, 12)

The way Sider and many of his contemporaries see it is that there might be a difference in degree but certainly not a difference in nature between theoretical philosophy and science.

However, Sider's metametaphysical stance is not at all pragmatist; it is realist. He believes that the goal of metaphysics is to represent the world as it is in and of itself. (Sider 2009) His realist view of metaphysics relies on a realist, rather than pragmatist, view of scientific inquiry, which entails that the scientific method allows us to not only figure out which conception of reality works best for creatures like us but also which conception of reality accurately represents the way things actually are.

Interestingly, this "Quinean" view rests on a misinterpretation of Quine's argument against Carnap. As Huw Price argues in "Carnap, Quine and the Fate of Metaphysics" (1997), by putting metaphysics on a par with science, Quine did not elevate metaphysics, but he demoted science. He was certainly no scientific realist. Instead, he was a pragmatist about science: science is about figuring out which theory works best. This is best brought out by this quote by Quine:

> Carnap maintains that ontological questions [...] are questions not of fact but of choosing a convenient scheme or framework for science; and with this I agree only if the same be conceded for every scientific hypothesis. (Quine 1951, 72)

Never mind its mistaken origins, the kind of realist naturalism that Sider endorses has been off and running for a while now, and it has arguably eclipsed the pragmatist naturalism of philosophers like Price. This is why it is important to note their differences. Whereas Sider's realist naturalism presupposes scientific realism, pragmatist naturalism is incompatible with a realist view of scientific inquiry. The reason pragmatists adopt a scientific methodology for their philosophical inquiry is not that it allows them to say what reality is really like, but because it connects intricate philosophical debates to testable practical differences. It enables us to decide between philosophical views by asking which work best in the field, so to speak.

I argue, in this thesis, for a realist view of scientific inquiry. I do so without violating my pragmatist principle, for I show that this conception of scientific inquiry is practically advantageous. However, this is not the case for philosophical inquiry. Thus, I maintain that pragmatist naturalism is indeed incompatible with a realist view of scientific inquiry. However, whereas other pragmatists accept the former and reject the latter, I accept the latter, and must therefore reject the former. This does not mean, of course, that I take philosophy to be anti-scientific. Philosophy must respect the results of scientific inquiry. It's just that philosophy, for the most part, operates at a different level than science. The mechanics of philosophical inquiry do not and should not mimic the mechanics of scientific inquiry. And, whereas scientific inquiry aims for truth, I argue that philosophical inquiry does not.

2.4 Pragmatism Without Naturalism

What would pragmatism look like without naturalism? The prospects might seem a bit bleak at first. Pragmatists are dead set against the idea that philosophy is the *a priori* analysis of concepts. They argue instead that to analyze a concept, like truth or knowledge, is to show how we use this concept in practice. Many pragmatists are heavily influenced by the evolution theory, which leads them to believe that our concepts, and how we use them, are contingent products of evolution; we conceptualize the world in one way rather than another because this way of conceptualizing the world happens to be more useful for creatures like us. Therefore, the fact of what truth or knowledge is like is a contingent fact that has to do with how we happen to use such concepts in practice. Hence, our concepts are best analyzed using empirical science to dissect these practices.

This is, I take it, the view of Price, who argues that many matters of philosophy of language are empirical issues that are best studied "from the outside" by biologists or anthropologists. (Price 2011, 110) A corollary of maintaining that science should reign over philosophy is a commitment to the idea that our conceptual framework can be explained by appealing to natural entities. In Price's words, "the natural entities themselves get cast in explanatory roles". (Price 2011, 31)

For instance, Dewey tries to give a natural history account of philosophical concepts, such as knowledge, by showing how their function in our practices has aided, and continues to aid, our survival as a species. He calls this the "genetic method" and sees his naturalism as a construction of a new Darwinian philosophy. (Misak 2013, 116) However, evolutionary forces act on our social and cognitive lives only indirectly by acting on our genetic makeup. Thus, a Darwinian philosopher is committed to the idea that scientific facts about our genes ultimately explain our philosophical concepts.

The same holds for negative arguments that appeal to evolution, such as Richard Rorty's argument that the aim for our beliefs to correspond to a mind-independent reality is inconsistent with the teachings of evolutionary theory. Our conceptual framework has evolved for the practical purpose of survival, and, Rorty argues, it is highly improbable that evolution has blessed us with a capacity that presumably all other animals lack: the capacity to figure out not just which conception of reality works best for creatures like us, but which conception of reality accurately represents the way things actually are. (Rorty 1995, 282) The argument here is thus that our genetic makeup is such that it could not accommodate the aim to correspond to reality.

A possible worry with this argumentative strategy is that it sounds pretty realist to say that there is an empirical fact of nature as to which conception of, say, truth or knowledge is correct. Price deals with this worry by arguing that this kind of claim should be interpreted as an object-level scientific claim instead of as a meta-level philosophical claim. At the meta-level, we could still say that science and philosophy are both language games that are not about representing a mind-independent reality but about developing a worldview that works best for creatures like us. (Price 2011, 142)

But what if there is a conflict between a philosophical thesis about, for example, which conception of truth is most useful and a scientific claim about what truth *actually is*? Suppose that one believes that, *as a matter of fact*, we do not aim to have our beliefs correspond to reality, even though it turns out to be practically efficacious, as I argue in Chapter 4, to *think* of ourselves as being involved with the project of trying to correspond to reality. In that case, Rorty's argument wouldn't hold up anyways since it undermines the premise in his argument that the aim for correspondence is practically inefficacious. However, imagine an anthropologist who has analyzed our *actual* usage of "is true" or a neuroscientist who somehow is able to translate patterns of neuron firings into an empirical theory about our concept of truth.

While Price has argued that this is consistent with pragmatism about science, would it also be consistent with pragmatism *about philosophy* to say that, though it might be most beneficial to *think* that we aim for correspondence, this is not *actually* the case?

The issue reminds me of Hume's treatment of causality as a psychological principle rather than a necessary connection in nature; it *pays to think* of events in the world as causally connected. It is a controversial issue as to how to interpret Hume's point, and perhaps he meant to say that causality does not *actually* exist, even though it is useful *to act as if* it did. Yet, this is certainly not how I see my own project. Once we know that it pays to think of ourselves as aiming for correspondence, the question as to whether we *actually* aim for correspondence becomes irrelevant and simply misguided. My kind of pragmatist about philosophy rejects the distinction between which philosophical theory *works best* and which theory is *actually correct*.

I am inclined to think that this is a necessary condition for any metaphilosophical view to count as pragmatist, but trying to argue for this claim, and thereby against pragmatist naturalism in general, would distract too much from the discussion of my own brand of pragmatism. I do not need to show that my pragmatism without naturalism is the only coherent form of pragmatism; I just need to argue that there is space for an approach like mine.

In rejecting the naturalist approach, I am rejecting an empirical approach to philosophy, and I am thus committed to the idea that philosophy is an activity of *a priori* reflection. There is no way around this. Am I thereby committed to the deeply anti-pragmatist, non-naturalist view that the *a priori* method can put us in unmediated contact with the inner workings of reality? Or should I otherwise defer to the highly restrictive view that philosophy must only concern itself with analytic matters? I do not think so. I argue in Chapter 6 that philosophical claims can be interpreted as normative claims that constitute a worldview, i.e., a theory about how we *ought* to conceive of the world and our place in it. My kind of pragmatism, without naturalism, thus considers its philosophical theses to be normative.

There is a connection between this approach and C. I. Lewis's pragmatic conception of the *a priori*. (Lewis 1923) Lewis argues that our *a priori* reflections tell us how we must conceive of reality. This sounds rather Kantian, but Lewis gives it a pragmatic twist by interpreting this "must" not as something categorical but as a practical must. For instance, we must conceive of reality as adhering to the laws of logic, but this does not mean that there is only one possible set of logical laws that structures our experiences and that we could not begin to doubt. Instead, there are many systems of logic, and their utility depends on our practical purposes, so we must adjudicate between them on practical grounds. Misak connects Lewis's thoughts on the *a priori* to Peirce's thoughts on regulative assumptions. A regulative assumption is something we *ought* to believe for us to carry out our practices in the way we believe we should. (Misak 2011, 268) My kind of pragmatist takes their philosophical theses to be of this nature: they are normative claims as to how we must conceive of things to make the best sense of our practices, where the "best" conception is to be the most useful one. They are akin to assumptions, not in the sense that they are unmotivated, but in the sense that they are not verifiable statements of facts. Thus, the philosophical views that work best are *ipso facto* the views that are correct. There is no need for an appeal to evolution to show why the philosophical concepts that work best are the ones we actually possess.

This does not mean that science cannot aid us in determining which philosophical views would work best for us. For example, in Chapter 3, I appeal to research that uses agent-based computer models to show the ways in which different possible responses toward peer disagreement would alter how we carry out our inquiry. (Douven 2010) I do so to determine which response to peer disagreement would be most useful in which situations. This kind of scientific research is very helpful for deciding which approach to peer disagreement to side with. However, though scientific research can help us determine which response to peer disagreement is most useful, it could not answer our philosophical question about peer

disagreement on its own because this is a normative question. The agent-based models show that one approach to peer disagreement, the steadfast approach, helps us gain fairly adequate beliefs quickly, whereas the other approach, the conciliatory approach, trades speed for accuracy. It is up to the philosopher to argue which one we *ought* to prefer. The scientist might lay out the options for us, telling us that the road to the left is short but undulating, whereas the road to the right is long but level, but the philosopher must ultimately argue which road we ought to take.

If the practice of philosophy is to develop a worldview — a unified theory that tells us how we ought to conceive of the world and of our place in it — then we must be mindful of facts about what the world *is actually like*, but once those facts are discovered the project would be entirely normative. An anthropologist could teach us invaluable things about our moral conduct, but they could not tell us whether lying is wrong, nor could they tell us what the nature of morality is. Likewise, they could show us how "is true" is used in different societies, which would again be of great value to us, but, on its own, this does not tell us what truth is. According to my metaphilosophical framework, a philosophical thesis about truth is a thesis that tells us what we *ought* to take truth to be for us to carry out our practices most advantageously. This is why the philosopher's reflections are, in some sense, *a priori*, armchair reflections and why the philosopher's duties cannot simply be carried out by anthropologists or other scientists.

2.4 A Strange Brand of Pragmatism

Let me restate the principle that encapsulates my commitment to pragmatism:

> Pragmatism: To give a philosophical analysis of a concept is to show its practical bearings.

This principle serves as a starting point that I share with other pragmatists. For instance, to show what truth is, we must show what use we make of it in practice. As shown above, some deflationists, like Rorty, believe that we have no use for truth at all except for its role as an expressive linguistic device. Others, like Misak or Price, argue that truth has a much more important role as the aim of inquiry.

However, I argue in Chapter 4 that we could only conceive of truth as having this role if we conceive of it as consisting in correspondence, which is why, if we want to carry out our inquiry in the way we think we should, we ought to take truth to be correspondence. This represents the first point of departure from the pragmatist tradition. It leads to another significant departure from the tradition. I argue that this realist conception of truth makes sense for scientific inquiry but does not make sense for philosophical inquiry. Hence, these two inquiries are substantially different. The former is shared, whereas the latter is personal. The former aims for truth, whereas the latter does not. These conclusions make pragmatist naturalism an untenable position.

Thus, whereas my project starts from the same place as other pragmatist theories, it ends up somewhere else entirely. The fact that both naturalism and the rejection of the correspondence theory are so intertwined with pragmatism makes it difficult to say to which extent I am still a pragmatist by the end of this thesis. However, I do hope that pragmatists who are convinced by my arguments see some promise in my brand of philosophy, regardless of the name it bears.

3. Summary of the Chapters

This thesis argues that philosophy does not aim for truth. Every chapter builds towards this claim. Chapter 2 asks what it is to aim for truth, focusing on what difference doing so makes

to our practices. It then argues that not all of our inquiries aim for truth, and it introduces a distinction between shared inquiry, which aims for truth, and personal inquiry, which lacks this aim. The next chapters add more layers to this distinction. Chapter 3 shows that, in the case of shared inquiry, it is generally rational to conciliate in the face of peer disagreement, whereas, in the case of personal inquiry, we can rationally be entitled to remain steadfast when we face peer disagreement. I then give an overview of how various forms of "unscientific" inquiry fit the mold of personal inquiry, such as aesthetic inquiry, moral inquiry, and even spiritual inquiry. Chapter 4 explores the metaphysical differences between shared and personal inquiry, arguing that shared inquiry is guided by an evidence-transcendent standard that makes its subject matter realist and objective. I explain how these metaphysical commitments could be seen as explanatorily downstream from practical commitments, which is why pragmatists need not reject them. Chapter 5 argues against a noncognitivist interpretation of personal inquiry by showing how we can separate rationality, and the aim for knowledge, from the aim for truth. By doing so, I introduce a contrast between shared and personal rationality, as well as a similar contrast for knowledge. Finally, Chapter 6 applies this framework to the case of philosophy by arguing that philosophical inquiry is personal. It thereby shows that philosophy does not aim for truth and offers an alternative interpretation of the nature and goals of philosophical inquiry.

While these chapters all build up to the conclusion that philosophy does not aim for truth, there are also aspects of the framework that I develop that, I hope, can be appreciated independently of its application to philosophical inquiry. For example, in Chapter 3, I argue that the discussion between the conciliatory and steadfast views on peer disagreement can be mediated by connecting these views to the presence or absence of the aim for truth. Also, in Chapter 4, I develop a pragmatist version of pluralism about truth, which I call pluralism *at the level of use*. I shall now give a chapter-by-chapter summary of the thesis that shows both the individual achievements of the chapters as well as the way in which the chapters hang together.

Chapter 2: The Aim for Truth and Its Practical Bearings

This chapter tries to expose the practical bearings of the aim for truth by engaging with the work of Huw Price. The Pricean view is that truth is a coordination device and that aiming for truth is practically beneficial insofar as aiming to converge towards the same beliefs is practically beneficial. I argue that this theory can be used to distinguish the aim for truth from the aim for beliefs that work best. Ultimately the former is in service of the latter. However, there is a distinction to be made between beliefs that work best for me and beliefs that work best for everyone. It can be useful to assume that these two coincide. In those cases, it is useful to pool our resources and treat inquiry as a shared enterprise. I call such an inquiry a "shared inquiry". Thus, in the case of shared inquiry, it is useful to aim toward convergence and, therefore, to aim for truth.

However, it is not always useful to assume that beliefs that work best for me coincide with beliefs that work best for everyone. For instance, in the case of beliefs about matters of basic taste, it seems very impractical to make this assumption. It would be more practical to maintain that different beliefs might work best for different people, e.g., the belief that pistachio ice cream is tasty might work great for me but not for you. This point might seem obvious, but it is easily overlooked when talking about the use of truth and disagreement *in general*. I argue that even without the aim to converge towards the same beliefs, it would still be productive to engage in disagreements because this allows us to structure our beliefs and keep each other in check. This is how we must understand disagreement about matters of basic taste. I argue that, therefore, the inquiry into such matters does not aim for truth. As a result, there is no pressure to coordinate our beliefs, which is why the inquiry is not shared but personal. I call this kind of inquiry that lacks the aim for truth a "personal inquiry".

Chapter 3: Peer Disagreement and the Aim for Truth

This chapter applies the distinction between shared and personal inquiry to the discussion on peer disagreement. The philosophical debate on this topic is concerned with figuring out which kind of response to peer disagreement is rational. There are roughly two sides to this debate: the conciliatory view and the steadfast view. The former argues that the mere presence of peer disagreement — that is, the presence of individuals who are presumably just as capable and well acquainted with the facts and who have nevertheless reached a verdict about the matter at hand that is incompatible with our own verdict — should have a negative impact on the confidence with which we hold our beliefs. The latter argues that we should, at times, be rationally entitled to remain steadfast in our beliefs; if we have sufficiently engaged with our peer's arguments and we deem them flawed and unconvincing, the disagreement should not have a diminishing effect on our confidence.

I show that this debate turns on what I call the *assumption of sameness*, which is the assumption that others generally experience and process things the same way as we do. If we maintain this assumption, it becomes viable to believe that, for all instances of disagreement, there is only one correct doxastic state to be had; we must either believe the proposition that is the topic of disagreement, disbelieve the proposition, or suspend judgment. This would, in turn, validate the conciliatory approach to peer disagreement. Yet, the assumption of sameness is also an assumption that is central to the question of whether our inquiry is shared or personal. If this assumption is generally beneficial, then our inquiry should be shared, which means it should aim for truth. If not, it should be personal and lack the aim for truth. Therefore, I argue that there isn't a single approach to peer disagreement that is rational across the board. The conciliatory view holds for shared inquiry, while the steadfast view holds for personal inquiry. The chapter concludes by showing in some detail why science is a paragon of shared inquiry

and by subsequently discussing different forms of "unscientific" inquiry — e.g., aesthetic, moral, and spiritual inquiry — which are shown to fit the mold of personal inquiry.

Chapter 4: Truth and Objectivity

This chapter is concerned with developing the metaphysical credentials of the distinction between shared and personal inquiry. I argue that deflationism about truth is incompatible with this distinction. Deflationism is committed to the idea that the nature of truth is captured, in full, by the equivalence schema: ⌜ *p* ⌝ *is true* if and only if *p.* This is why their expression of the truth norm ("It is correct to assert *p* if and only if ⌜ *p* ⌝ is true") must be equivalent to the assertion norm ("It is correct to assert *p* if and only if *p*). However, the assertion norm ranges over all discourse, including those areas of discourse that personal inquiry belongs to. If the truth norm is to play the role of a coordination device, then it must entail more than what is expressed by the assertion norm. I then consider the proposal that truth is epistemic, and I argue that this won't suffice. I draw on Crispin Wright's notion of Cognitive Command (Wright 1992, 144) to argue that truth for shared inquiry must be correspondence to reality and that the subject matter of this inquiry must be seen as realist and as objective. However, I give a pragmatist interpretation of these metaphysical commitments. These commitments are explained by our practical commitment to the assumption of sameness.

I then turn to the question of whether my account of truth is a pluralistic — or rather dualistic — account of truth. I argue for correspondence as the only truth property and for a unified conception of truth as a linguistic device for expressing agreement. However, I do show that this device can be used in two distinct ways and that the metaphysical differences described above are due to this difference in how we use the concept of truth in personal and shared inquiry. Therefore, my account of truth is pluralistic, but only *at the level of use*.

Chapter 5: Rationality and Knowledge

Whereas the previous chapter concerns itself with the metaphysical credentials of the distinction between shared and personal inquiry, this chapter is concerned with its epistemological credentials. It might be tempting to think of personal inquiry in a kind of noncognitivist fashion as concerned with sentiments instead of beliefs. This chapter shows that this view would be mistaken. It argues that personal inquiry is involved in games of giving and asking for reasons, which shows that its contents must be such that they could be offered as, or stand in need of, reasons. As Robert Brandom (1994) argues, this is precisely what it is for an attitude to count as a belief.

However, what could rationality for personal inquiry amount to? I show that it is rather difficult to find an account of (theoretical) rationality that does not presuppose that we aim to be rational in order to get to the truth. Personal inquiry does not aim for truth, which is why this presupposition would be problematic. The chapter explores the connection between reasons and rules, and Kripke's Wittgenstein's (1982) thoughts on rule-following. It then develops an understanding of rationality in a rule-following fashion, and it shows that in the case of personal inquiry, we are allowed, to a certain extent, to follow our own rules, whereas in the case of shared inquiry, we must aim to follow the same rules.

The chapter then turns to knowledge and asks whether personal inquiry could aim for knowledge. *Prima facie*, the answer would be no. Truth and knowledge are related concepts, and it seems difficult to aim for knowledge without thereby aiming for truth. Nevertheless, I argue that this is possible. I give an account of the *personal* aim for knowledge as the aim for practical certainty. Then, I interpret knowledge attributions as a way of expressing shared doxastic commitments. This is related to the function of truth as a device for expressing agreement. I show that the *shared* aim for knowledge involves the aim for convergence toward the same beliefs, just like in the case of the aim for truth.

Finally, I show how these two senses of aiming for knowledge explain the intuitive differences in the mechanics of testimony between shared and personal inquiry. In the case of shared inquiry, we are rationally allowed to adopt beliefs through the means of testimony, even if we do not understand the reasons behind these beliefs. This is much rarer in the case of personal inquiry, in which case understanding often seems to be a requirement for knowledge.

Chapter 6: Philosophy as Personal Inquiry

In this chapter, I finally apply the framework of the thesis to the case of philosophy by asking whether philosophical inquiry should count as personal or shared inquiry. I show that, though philosophical inquiry shows some similarities to scientific inquiry, many of the philosopher's practices, such as their hiring decisions, their attitude to peer reviewing, and their teaching, to name just a few, show far greater tolerance for disagreement than what is common to scientific practices. I contemplate the suggestion that this diversity of opinions might be in service of our aim for truth, but I argue that it couldn't be.

The chapter then proceeds to show how we can conceive of philosophical inquiry as personal. I respond to the worry that this conception of philosophy turns all our philosophical disagreements into mere matters of taste. I draw some comparisons between the development and organization of academic philosophy and that of the artworld. The philosopher's core task, I argue, is to develop a coherent worldview, which is not a description of what the world is like, but rather a view of how we ought to conceive of the world and our place in it. This conception of philosophy draws no principled distinction between practical and theoretical philosophy.

Finally, I argue that once we conceive of philosophical inquiry as personal rather than shared, we can resist the thought that there is no, or only very little, philosophical progress. There are many widespread disagreements in philosophy, and given the presupposition that it

is generally rational to conciliate in the face of peer disagreement, these disagreements would undercut our justification for our philosophical beliefs, which means that we would not be rationally entitled to believe our philosophical views. Insofar as progress is constituted by the accumulation of knowledge, philosophical progress would thus seem out of reach. However, if philosophical inquiry is personal, the presupposition of the conciliatory response to peer disagreement can be rejected, which means that there could be room for justified belief and even knowledge in spite of the presence of widespread disagreement.

This shows that once we let go of the idea that philosophy aims for truth, we no longer need to think of philosophical inquiry as the disorganized and second-rate cousin of scientific inquiry. Instead, we can conceive of our philosophical practices as thriving in spite of, or perhaps even thanks to, our propensity for entrenchment in the face of disagreement.

2. The Aim for Truth and Its Practical Bearings

1. Introduction

The research question of my thesis is: Does philosophy aim for truth? To answer this question, it is paramount that we get clear on what it is to aim for truth and what it is to lack this aim. As I explained in the introductory chapter, I endorse a pragmatist methodology that dictates that "getting clear on what it is to aim for truth" should be interpreted as exposing the practical bearings of the aim for truth. This is the goal of the present chapter. I aim to realize this goal by engaging with the work of Huw Price.

In section 2, I introduce Price's theory and show why it is practically beneficial to aim for truth rather than to aim (directly) for beliefs that work best. Price argues that truth serves as a coordination device that helps us resolve our disagreements and converge toward the same beliefs. I show that this convergence is practically beneficial. However, in section 3, I argue that aiming for truth is not practically advantageous for all domains of discourse. I draw a distinction between shared inquiry, which is guided by the aim for truth, and personal inquiry, which lacks this aim. I argue that inquiry into matters of basic taste is the exemplar of the latter kind of inquiry, whereas scientific inquiry is a paragon of the former. Section 4 discusses Price's arguments for why disagreements about matters of basic taste must involve the aim for truth, and it tries to counter these arguments. Section 5 explains that the function of such disagreements is not to facilitate the coordination of beliefs but that it is instead to get a grip on each other's dispositions to act and to keep each other in check through, what Robert Brandom (1994) calls, the game of giving and asking for reasons. Finally, section 6 uses Lionel Shapiro's (2021) account of the role of truth in this game of giving and asking for reasons to show that truth has a practical role to play in our discourse, even in those cases in which we do not aim for it.

2. Why Aim for Truth?

The question that I hope to answer is why we should take ourselves to be aiming for truth. Why not simply aim for beliefs that work? From a pragmatist perspective, it seems that all that is needed for our inquiry to be successful is to gain beliefs that we do not have to doubt. If our beliefs work — if they aid us effectively in satisfying our desires — then we have no active reason for doubting them.

A slightly different question is asked by Richard Rorty. He asks what the practical difference is between aiming for truth and aiming for justification. He argues that there is none and that, therefore, the aim for truth has no practical bearings. He then concludes that, as pragmatists, we must reject that truth is something we can aim for. Instead, truth is merely an expressive device for endorsement and generalization. It is devoid of normativity.[2] Rorty puts the point as follows:

> Pragmatists think that if something makes no difference to practice, it should make no difference to philosophy. This conviction makes them suspicious of the philosopher's emphasis on the difference between justification and truth. For that difference makes no difference to my decisions about what to do. If I have concrete, specific, doubts about whether one of my beliefs is true, I can only resolve those doubts by asking if it is adequately justified […] Assessment of truth and assessment of justification are, when the question is what I should believe now […] the same activity. (Rorty 1995, 281)

Rorty talks about "what I should believe now" because he is mindful that the aim for truth and the aim for justification can come apart retrospectively. Suppose that yesterday I asserted *p* and that I was justified in doing so, but today I find out that this warrant was based on misleading

[2] In the next chapter, I shall go into more detail about the deflationary conception of truth and its relation to the normativity of truth. For now, I shall focus solely on the question of whether we aim for truth.

and incomplete evidence and that *p* is not actually true. If my aim is truth, then I would retrospectively take myself to have been at fault in spite of having been justified. Yet, if I only care about justification, I would alter my beliefs without blaming myself for my earlier misgivings.

In any case, it is this problem, raised by Rorty, that Huw Price responds to in "Truth as Convenient Friction" (2003). In this article, Price argues that Rorty has overlooked an important practical use of the aim for truth. He shows that the aim for truth is distinct from the aim for justification by exploring what would be lost if we were to merely aim for justification. Price calls an assertion that aims merely for justification a "merely-opiniated assertion (MOA, for short)", and he calls the counterfactual community whose assertoric practices are guided only by justification the "Mo'ans". (Price 1998, 247)

The Mo'ans are still allowed to use a deflationary notion of truth to register agreement and disagreement, but truth for them has no normative force that is distinct from the normativity of justification. According to Price, the only norms the Mo'ans are guided by are those of "subjective assertibility" and "objective assertibility". (Price 1998, 246) The former norm tells us that one should assert that *p* only if one believes that *p*, whereas the latter norm states that one should assert that *p* only if one has good evidence for *p*. Price maintains that the Mo'ans treat differences of opinion in the same way as mere differences of preference:

> Mo'ans use linguistic utterances to express their beliefs (as well as other psychological states, such as preferences and desires). Where they differ from us is in the fact that they do not take a disagreement between two speakers in this belief-expressing linguistic dimension to indicate that one or other speaker must be at fault. They recognise the possibility of fault consisting in failure to observe one of the two norms of subjective or objective assertibility, but lack the idea of the third norm, that of truth itself. This shows up in the fact that by default, disagreements

> tend to be of a no-fault kind (in the way that expression of different preferences often are for us). (Price 1998, 250)

Price thus argues that the assertoric practices of the Mo'ans are quite distinct from ours, for they can do nothing more than express opinions. They only have reason to engage in disagreement if they suspect that someone is asserting insincerely or without justification.

By aiming for truth instead of justification, we treat truth as a norm distinct from the norms of sincerity and justification. What the truth norm does, according to Price, is that it adds a kind of *a priori* fault to disagreements:

> [The truth norm] is a norm which speakers immediately assume to be breached by someone with whom they disagree, *independently of any diagnosis of the source of the disagreement.* Indeed, this is the very essence of the norm of truth, in my view. (Price 2003, 164)

The fault added to disagreement by aiming for truth does not reduce to the violation of the norms of justification or sincerity because if it did, we would only take each other to be at fault once the source of our disagreement has been identified. Instead, the truth norm adds an *a priori alethic* fault to disagreement: one of us is in the wrong merely because of asserting a falsehood, even if one is fully justified in doing so.

Why would this additional fault be useful? Price argues that its utility lies in its coordinative role. He conceives of the truth norm as a coordination device and points out the *prima facie* utility of coordinating our opinions in light of reasoned argument:

> Without the [truth] norm, difference of opinion would simply slide past one another. Differences of opinion would seem as inconsequential as differences of preference. […] The [truth] norm makes what would otherwise be no-fault disagreements into unstable social

> situations, whose instability is only resolved by argument and consequent agreement. […] If reasoned argument is generally beneficial — beneficial in some long-run sense — then a community […] who adopt this practice will tend to prosper, compared to a community who do not. (Price 2003, 180-181)

It sounds intuitive that we end up with better opinions if we think of our inquiry as a shared concern with finding the truth rather than an individual concern with finding justification.

However, as Price understands, which opinions are *better* must not be measured in terms of representational accuracy, for then the motivation for aiming for truth would be circular. Sure enough, aiming for truth probably renders our opinions more accurate and, thus, more likely to be true, but whether our opinions are accurate (true) or not should be of no concern to us unless truth is what we aim for. Why should the Mo'ans adopt our practices, given that they, by postulation, do not value truth? The answer must be that aiming for truth makes their opinions better, not according to our standard (of truth), but according to their own standard. What needs to be shown is that aiming for truth indirectly makes our opinions stand up better to current or future evidence. What must be shown is that, by coordinating our opinions, we end up with beliefs that have better *practical* qualifications.

Let me repeat the Jamesian motto from the introduction:

> Our obligation to seek truth is part of our general obligation to do what pays. The payments true ideas bring are the sole why of our duty to follow them. (James 1907, 230)

Price argues that we seek truth because we aim for convergence. Since aiming for convergence pays, our aim for truth can be seen as a part of our obligation to do what pays. But we have not explicitly shown why aiming for convergence pays, and, as it turns out, it does not always pay. This is very significant because if it were always beneficial to aim for convergence, it would

be obvious to have this aim be a part of our obligation to do what pays, which would make truth a very shallow and platitudinous notion. It would be like saying that since I aim to get a drink from the fridge, I aim to get up from my chair. However, if I aim to see a friend who lives in a different town, and there are multiple ways of getting there, by train or by car, then it becomes more interesting to know if I should aim to go by train. This, I shall argue, is what it is like in the case of truth. Our obligation to do what pays is personal: I ought to do what works best *for me*. Whether that involves the kind of collaboration facilitated by the aim for truth is an open question. It is useful to aim for convergence if and only if what works best *for me* is also what works best *for everyone else*.

Price argues that, when it comes to speech acts and their corresponding mental states, this is only the case for assertions and beliefs. This is why we only apply the concept of truth to assertions and their contents. Truth is a device for expressing agreement, and since agreement is what we are after in our assertoric practices, it is useful to have this device in place. We should aim to converge toward the same beliefs because it is presumably beneficial to assume that the beliefs that have the best chance of working for me are those that also work for everyone else. The following quote from an early paper of his should provide some textual evidence for this interpretation of Price's theory:

> [In] a dispute, rival, incompatible views are exposed to common scrutiny. Ideally the more well-justified prevails, and one speaker recants, accepting the view of the other. Plausibly, there is enough of a general advantage in such dispute behaviour to explain the existence of a powerful linguistic device to facilitate it (i.e., the use of 'true' and 'false'). This advantage will be explained in terms of the behavioural consequences of particular views, and the consequent benefits of basing one's views on as wide a body of experience as possible. For the deliberately vague term 'views' here, it is natural to read 'beliefs'. There would seem too little point in such a "dispute" with respect to utterances expressing, say, desires, which could reasonably vary

> from speaker to speaker, even in the face of the same evidence as to matters of fact. In other words, the suggestion is this: some utterances (call them 'assertions') characteristically express states of mind ('beliefs') with respect to which there is reason to seek agreement between speakers. For these states of mind, "two heads are better than one"; there is a general advantage in exhibiting differences between speakers in this respect, so that less well-justified beliefs may be replaced by more well-justified ones. This explains why language has developed a general means of indicating such agreements and disagreements, in the application of the terms 'true' and 'false' to the associated utterances. Utterances such as questions, commands and requests, on the other hand, characteristically result from states of mind for which no such reason for unanimity exists. Different speakers can reasonably hold conflicting such states of mind (conflicting in the sense that no one person could hold them concurrently), even if fully acquainted with each other's viewpoint. Appropriately, ordinary usage does not apply 'true' and 'false' to the types of utterance which express, or result from, these states of mind. (Price 1983, 356-357)

Once we apply the Jamesian motto to Price's story, it will state that, in the case of assertions, the aim for convergence, and hence the aim for truth, is a part of our general obligation to do what pays. Price does not explain here why assertions and the beliefs they express are such that it is useful to pool our evidence and aim for agreement. He *stipulates* that assertions are the speech acts for which this holds and that the views they express are beliefs.

Price argues that the concept of truth, as a device for expressing agreement, is only useful insofar as it is useful to aim for agreement, which is why it is only applicable to assertions. He also argues that this aim marks the difference between a subjective expression of personal preference, e.g., "I like pistachio ice cream", and an assertion, e.g., "Pistachio ice cream is tasty". There is no friction when I say "I like pistachio ice cream" and you say "I don't", but there is friction when I say "Pistachio ice cream is tasty" and you say "No, it isn't".

Thus, I use the former way of putting it when I merely want you to be aware of my ice cream preference, whereas I use the latter if I aim for you to agree with me. (Price 2022b, 27)

3. Shared and Personal Inquiry

Price stipulates that beliefs, i.e., the contents of assertions, are such that it is useful to aim to converge and, therefore, that it is useful to aim for truth, whereas this does not hold for other mental states, such as desires. In this section, I argue against this stipulation by showing that it is not useful for all of our inquiries to assume that the beliefs that work best for us are the beliefs that work best for everyone. I call those inquiries for which this assumption is not useful "personal inquiries", since they consist in the search for beliefs that work best *for me*. I define "shared inquiries" as inquiries that are after beliefs that work best *for everyone*. Given the practical role of the aim for truth described in the previous section, it turns out that personal inquiries do not aim for truth.

It might seem evident that reasoned argument and subsequent convergence lead to better beliefs. As Price puts it in the quote above, two heads are often better than one. However, as I have argued, this does rely on the assumption that the beliefs that have the best chance of working for me are those that also work for everyone else. The following example shows a case in which this assumption seems useful:

> **Horse Races:** I have been to the horse races five times, and won every single time. I do not know much about horse races, and I've never bothered looking into the strengths and weaknesses of individual horses. Instead, I've always put my money on the third leftmost horse. I have come to believe that the third leftmost horse is the surest bet. I have shared this belief with my friends: "Look, don't bother doing your research; the

third leftmost horse always wins." My friends have tried the strategy, and it has failed them; they didn't manage to win a single time. I go back to the races to see for myself whether my strategy still works, so I bet on the third leftmost horse, and I win again.

The question is: should I still believe that the third leftmost horse is always the surest bet? Common sense dictates that I shouldn't and that in evaluating this belief, I should pool the evidence instead of considering only my personal success. However, what I am ultimately after is a belief that works best *for me*, and my belief has yet to fail me. Could it not be the case that different beliefs work best for different people? Well, no it couldn't because whichever horse wins the race is the same for everyone. If we were to go to every horse race together, we would find out soon enough that, insofar as our different beliefs lead to different bets, it is simply not the case that different beliefs work for different people.

The general idea is that there can only be one best belief for everyone because there is only one way things really are. Thus, whenever we have conflicting beliefs, one of us must be mistaken. This is what the aim for truth consists in on the practical level: the *a priori* commitment to taking disagreements to imply faults.

Pragmatists are unlikely to accept this order of explanation, though. If anything, the fact that it is useful to pool our evidence and to engage with each other's beliefs should be seen as explanatorily primary. It is this fact about our practices that gives rise to the metaphysical idea that there is only one way things really are.

However, this idea about the utility of pooling our evidence does not hold up for all domains of discourse. Let's consider another example:

Restaurant: I have been to the Yellow Pigtail, a neighborhood restaurant, five times, and each time the food was terrific. I have come to adopt the belief that the food of the

> Yellow Pigtail is excellent. I have told my friends: "If you're looking for a fantastic meal, the Yellow Pigtail is your surest bet." My friends have tried it and have all been disappointed. I return to the restaurant and try the food, which is once again excellent.

In this case, it would be useful *not* to pool the evidence and to stick with my belief. Why is that? Well, because gustatory taste is personal, and different beliefs work best for different people. If we were to always go to the same restaurants and pick the same meals off the menu, we would figure this out pretty quickly.

Price assumes that it is always beneficial to pool our resources and aim for convergence, but this does not seem to be the case for the inquiry into matters of basic taste. It is only useful insofar as it is useful to assume that beliefs that work best for me are beliefs that work best for everyone. This assumption is motivated by the assumption that we experience and process things in the same way. Let's call this the *assumption of sameness*. Scientific inquiry is built on this assumption. Scientific experiments and procedures are standardized in such a way that they allow us to build off each other's observations and computations. It is the archetype of what I call "shared inquiry".

Yet, why would it be wise to go by my own tastes but foolish to go by my own science? Suppose I disregard others' opinions and live by my own scientific theories. In that case, chances are that my theories would serve me well enough and that my justification for these theories would never be undermined. After all, I don't have the time or money to go to the moon to see how my gravitational laws will hold up there. However, I would risk missing out on theories that could serve me much better. I could isolate myself from the collective enterprise of modern science and live out my days in a shoddy shed I built using my primitive laws of physics. Perhaps from a moral perspective, we would all be better off living in shoddy sheds far away from modern science because current technology can do more harm than good.

Yet, from the perspective of the individual, this would be a dangerously impractical attitude to have. It would be foolish, for instance, for me not to get vaccinated because the science of vaccination is not compatible with my personal medical science, even if my personal medical science has, so far, kept me in good health.

By treating science as a shared enterprise, we can develop much more intricate and sophisticated theories with a much greater predictive accuracy, which could ultimately serve us better in practice. To engage in science in this shared fashion, we must pay heed to the beliefs of our peers. In Chapter 4, I argue that this requires conciliation in the face of disagreement. We should not rationally be allowed to remain steadfast in our belief when we find out that our peers disagree with us. This would be a similar rational failing as that of the superstitious gambler who keeps betting on the third leftmost horse even after discovering that this strategy has failed their friends many times. We need to believe that something is amiss when we disagree, and that, whenever our beliefs are incompatible, they cannot both be perfectly in order. This deficiency is described by the alethic fault introduced to our disagreements when we aim for truth. Therefore, it is the aim for truth that is characteristic of a *shared inquiry.*

To treat an inquiry as a shared inquiry is a gamble in and of itself. It would be easier to go about our inquiry by ourselves without paying any heed to the beliefs of others. This would not prevent us from gaining beliefs that are practically adequate. To treat an inquiry as a shared inquiry is to risk losing perfectly adequate beliefs to gain beliefs that will probably serve us even better. Even the Mo'ans should be moved by this gamble. They could hold onto the justification of their scientific beliefs by being indifferent to the fact that their peers disagree with them — just as we do in the case of our opinions on pistachio ice cream — and it might just be that they will be able to live by them well enough never to have to put them in doubt. Yet, by treating their scientific opinions like a matter of personal taste instead of engaging in

science as a shared enterprise, they would most likely miss out on beliefs that could serve them much better.

However, in the case of basic taste, this gamble would not pay off. If we were to doubt our beliefs as a result of disagreement, then we would probably gain precious few beliefs while losing out on a great many that would have been perfectly adequate. Thus, when it comes to taste, the Mo'ans have not gone extinct at all. They live among us. Perhaps this is why we commonly treat disagreements about basic taste as *faultless*. There is no practical need to try to reach an agreement on such matters. Instead, we can agree to disagree.

Insofar as the aim for truth is the aim for convergence, it would make no sense to aim for truth when it comes to our inquiry into matters of basic taste. It would thus make no sense to posit that our disagreements contain an alethic fault. However, this does not mean our disagreements are always entirely faultless. There are plenty of other non-alethic standards. I shall discuss these standards in Chapter 6. For now, it suffices to say that the rejection of the *a priori* commitment to *all* disagreements containing faults does not imply that *no* disagreements contain *any* faults whatsoever. It does imply, however, that whenever two inquirers carry out their inquiry impeccably and arrive at incompatible beliefs — incompatible in the sense that a single inquirer could not rationally endorse both of their beliefs —they need not presume that this incompatibility is indicative of a fault. Instead, they would both be entitled to their respective beliefs. This is what makes their inquiry a *personal inquiry*.

One way of characterizing the distinction between personal and shared inquiry is to say that our inquiry into matters of taste aims for *personal* justification, whereas our inquiry into scientific matters aims for *interpersonal* justification. That is, scientific inquiry aims for beliefs that *everyone* would be justified in having. There could only be such beliefs once we settle our disagreements. We have to be careful though, not to see this aim for interpersonal justification as able to explain why we settle our disagreements. Instead, it is the practical value of

coordination that explains why we aim for interpersonal justification in the first place. Insofar as truth plays the role of a coordination device, it is the aim for truth that explains why we aim for interpersonal justification.

Now we can finally answer the question, posed a while back, of why the Mo'ans should join us in aiming for truth if what they ultimately care about is not to gain opinions that accurately represent reality but to gain opinions that are practically serviceable. Yet, while they should join us in the case of science, for there are beliefs to be gained that could serve them better, in the case of taste, *we* should join *them*. Given the conceptual differences between the two cases mentioned above, I take it that, in the case of taste, we already belong to their clan.

4. The Phenomenology of Disagreement

As we have seen, Price argues that all assertions aim for truth, whereas other speech acts do not. According to Price, this is the distinction between a no-fault disagreement, like when we express different preferences — "I like pistachio ice cream", "I don't" — and a *substantive* disagreement of assertions, such as "Pistachio ice cream is tasty", "No it isn't". I have argued against the claim that an assertion like "Pistachio ice cream is tasty" aims for truth, and I have noted in passing that this kind of disagreement about basic taste is generally considered to be faultless as well. I have shown that it is practically unsound for the inquiry into matters of basic taste to aim for truth since it is not useful to assume that beliefs that work best for me are those that would work best for everyone else as well.

Yet, if what I have argued is true, why do we have this difference between "I like pistachio ice cream" and "Pistachio ice cream is tasty"? Price argues that even in the case of taste, it can sometimes be useful to engage in disagreement, but this utility is defeasible, which is why we can so easily transition from the latter claim to the former. Ordinarily, we would not

characterize our judgments on the tastiness of pistachio ice cream as having a shared aim. Instead, our inquiry into the tastiness of pistachio ice cream seems to be a purely personal one. Yet, Price argues that assertions like the assertion that pistachio ice cream is tasty do have a shared aim — a "bull's-eye" in the words of Simon Blackburn (Blackburn 2005) — but that the purported objectivity of the bull's-eye can easily be canceled out. (Price 2022a, 2022b)

He argues that, though all assertoric practices aim for the bull's-eye, the assertoric practices that seem more subjective have built-in escape hatches, which cancel out the normativity of the assertion. For instance, when I say "This pie is delicious!" and my assertion is met with a lot of resistance, I could cancel out the claim to objectivity implicit in the assertion by saying something overtly subjective, like "Well, I like it at any rate". Such escape hatches seem less appropriate for other domains of discourse. For instance, the weakening of an assertion like "14 + 13 = 27" to some overtly subjective claim like "To me, "14 + 13 = 27" seems quite ridiculous. This is not to say that there are no escape hatches at all for discourse on arithmetic, but only that there are not as many and that they are not used as frequently as in the case of taste. Price says:

> The phenomenology of subjectivity arises from the fact that use of the escape hatches cancels the normal objectivity of the bull's-eye – the normal presumption that there is a common goal – but we only see this properly for what it is when we understand the source of the latter. Proper play of the game – required for it to serve its coordinative function – requires that the escape hatches initially be ignored. (Price 2022b, 43)

There must be something to this idea that ordinary assertions about basic taste are different from the overtly subjective claims that are the escape hatches. Why else do we have two ways of expressing our preference: "I like pistachio ice cream" and "Pistachio ice cream is tasty"? The latter allows for disagreement in a way in which the former does not.

However, truth need not be invoked to explain this contrast, for it could be explained by appealing to Price's own distinction between subjective and objective assertibility. A claim such as "I like pistachio ice cream" must be made sincerely, but usually, that is all that is required of it. We are allowed to like what we like even if we cannot think of good reasons for doing so. This is not so for beliefs. If I say "Pistachio ice cream is tasty" I invite you to ask for reasons for this claim and to disagree with me by challenging the adequacy of my reasons. According to Robert Brandom (1994), the essential feature of the assertion is that it is a move in the game of giving and asking for reasons. Suppose I say "Pistachio ice cream is tasty" and you respond by saying "No, it isn't; it's not sweet enough for ice cream; it has a weird off-putting green color; and it usually doesn't taste remotely like real pistachios". I might be at a loss for words, unable to back my claim up with reasons of my own, but instead of saying "You're right; it is disgusting", I would be more likely to say "Well, I still like it". I thereby use the escape hatch not to cancel the normativity of truth but to cancel the normativity of objective assertibility. That is, I mitigate my justificatory burdens.

Price argues in the above quote that we must have the bull's-eye in place so that truth can serve as a common goal and thereby play its coordinative role. Yet, in the case of taste, this idea of a common goal makes no sense, either conceptually or practically. I have already discussed the practical side. Conceptually, there is also a difference between, for instance, scientific inquiry and inquiry about taste. In the case of science, when we are confronted with disagreement, the proper response seems to be to realize that our beliefs can be improved upon and that we must continue our inquiry and keep testing our hypotheses. We do so because we believe that we must coordinate our beliefs if we are to get to the truth. This is precisely the use that the truth norm has, as Price notes:

> Unless individual speakers recognize such a norm, the idea that they might improve their views by consultation with the wider community is simply incoherent to them. (It would be as if we

> gave a student full marks in an exam, and then told him that he would have done better if his answers had agreed with those of other students.) (Price 2003, 174)

Yet, the parallel story for disagreement about taste sounds quite ridiculous: learning that you disagree with me about whether pistachio ice cream is tasty should not make me think that something is amiss and that my belief can be improved upon. The idea that our opinions would ultimately come together as long as we just carry on with our inquiry by eating more and more pistachio ice cream, thereby testing our hypotheses, seems entirely misguided. To exploit Price's own words, once we have sufficient justification for our basic tastes, "the idea that [we] might improve [our] views by consultation with the wider community is [indeed] simply incoherent". (Price 2003, 174)

5. The Practical Functions of Disagreement

Price assumes that the practical function of disagreement is always to coordinate opinions and that, therefore, the proper functioning of disagreement will always ensure that there is a bull's-eye to aim for. In the above section, I have argued against this assumption. I have shown that disagreements on gustatory taste generally lack the bull's-eye. Price argues that the lack of a bull's-eye would mean that a disagreement would never get off the ground because we would have no reason to convince each other of anything. However, I think that a pragmatist account of assertions would reveal that disagreement still serves a purpose.

As mentioned above, Robert Brandom conceives of our assertoric practices as games of giving and asking for reasons. (Brandom 1994) Disagreement is an integral part of such games. Therefore, when asking about the function of disagreement, we must ask what the practical advantages are of engaging in these games of giving and asking for reasons. Price would presumably say that the coordination of our opinions is the only purpose of this

engagement. Yet, I think there is a more fundamental practical advantage to entering into games of giving and asking for reasons.

Without the external pressure to defend our opinions, we are at risk of opining uncritically, thereby allowing biases and prejudices to cloud our judgment. This has detrimental effects even in the case of taste judgments. If we are not prompted to seek out the reasons behind such judgments, we will not be able to draw out the patterns that link our dispositions, which in turn leads to the inability to effectively develop our taste in areas we are unacquainted with.

Given the practical advantages of finding out the reasons behind our taste judgments, one might wonder why we wouldn't internalize the game of giving and asking for reasons and why we must be prompted by others to make such evaluations. To a certain extent, we do indeed internalize this process. However, there is the danger of prejudice. This danger is that one, more often than not, simply assumes that there are good reasons for endorsing whichever opinions they happen to endorse. It is only when we encounter individuals who challenge our opinions and demand reasons for them that we start to engage with them critically.

An interesting question is what prompts these individuals to demand reasons from us in the first place. If there is no desire for coordination, then our opinions should be of little interest to them. Of course, a disagreement goes both ways, so what one gets in return is the opportunity to critically examine one's own opinions as well. Is it simply a matter of trading favors, of scratching each other's backs?

No, it isn't. There are more important reasons for others to take an interest in our opinions and for them to demand reasons for them. To use reason to subjugate one's opinions is to find or create a structure in them. Once we know this structure, we can use it to predict one's opinions and, thereby, one's actions.

In the case of taste, our opinions are rarely as structured as we would want them to be, and the reasons we cite to defend our opinions might have very little to do with why we actually have these opinions. We often have no idea why we enjoy certain kinds of food, music, or literary genres while we detest others. However, our actual capricious opinions or whimsical reasons for them might not be of much importance to our interlocutor. When we express our opinions in dialogue and draw up reasons for them, we subjugate ourselves to a set of rules and allow our interlocutor to hold us accountable if we violate these rules.

Disagreement is, therefore, a useful activity even without the aim for truth because it forces us to critically examine our own opinions, and it allows our interlocutor to make sense of our behavior and hold us accountable. If our discourse on taste lacked the assertoric mode and was instead purely expressive, then we would not have this kind of accountability. I could say 'Yum' or 'Yuk' whenever I am disposed to do so, without thinking of these expressions as standing in need of reasons and without worrying about being held accountable for any of my claims. Yet, if we were to engage in a game of giving and asking for reasons, e.g., "Why is beer yuk?" "Because beer is bitter, and bitter is yuk", we would simply seem to incorporate terms like 'Yuk' into assertoric discourse, and they would cease to be purely expressive.

As mentioned above, Price believes that even disagreements on taste must involve the aim for truth because without the need to coordinate our opinions we lose interest in the disagreement altogether. He then uses his "escape hatches" to explain away the appearance that disagreement on taste is different from scientific disagreement. In the previous section, I have argued, however, that the former disagreement is indeed different from the latter because the former is faultless in a way in which the latter is not. This faultlessness is explained by the lack of an alethic fault. I have argued in the present section that Price is wrong in thinking that this lack of the truth norm would make us lose interest in the disagreement altogether since there are practical advantages to engaging with each other's opinions that have nothing to do with

coordination. We should accept that, when it comes to taste, we are all differently inclined. Yet, this does not mean that we cannot enter into disagreement, thereby employing the mechanism of giving and asking for reasons to try to critically examine our different inclinations and understand how these differences affect the ways in which we are disposed to act.

6. Shapiro's Account of the Practical Role of Truth

Lionel Shapiro believes that once we think of assertoric practices as games of giving and asking for reasons, Price's account of the role of truth can be explained without having to think of truth as normative at all. (Shapiro 2021) According to Shapiro, we can use the term 'true' to share the justificatory burdens of our beliefs. I could say "*S* said *p*, and what *S* says is true", thereby deferring to *S*'s communicative authority or, in another dialectical context, allowing *S* to defer to my authority. It is convenient to pass the buck sometimes, even in the case of taste. If I happen to know that our friend has roughly the same taste as I do, then instead of going through the troubles of giving you all my reasons for my taste judgments, I could respond to our disagreement by saying that our friend agrees with me and that whatever they believe is true. In this case, I defer to our friend's authority and thus to whatever reason they have for their taste judgment.

While I agree that truth could play this role and that it need not be normative to play this role, I also believe that this is entirely consistent with the idea that the practical purpose of disagreement need not be to coordinate our opinions. Shapiro's explanation of truth shows how it enables us to engage with others with whom we agree, and to share the justificatory burden of our assertions. Yet, this does not imply that we should actively revise our opinions in an attempt to agree with others so that we can gain justification more easily. The reason for me to defer to our friend is that we already happen to be coordinated in our opinions.

As stated before, Price identifies two norms in addition to the truth norm: sincerity and justification. To simply change one's opinions so as to meet the norm of justification more easily would be to betray the first norm to satisfy the second. If one's opinion proves to be unjustifiable, revision might be in order, but not if one's opinion is merely unpopular, unless we aim to coordinate our opinions, which I have argued we do not do for matters of taste.

Shapiro aims, in his words, to give "an explanation of how ascriptions of truth and falsity motivate participating in reasoned dialogue." (Shapiro 2021, 11) I think he has succeeded in doing so. Still, this explanation does not require that we think of truth and falsity as coordination devices, so Shapiro's account is perfectly consistent with the idea that reasoned dialogue need not have convergence as its aim.

As I have mentioned above, Price explains why assertions are apt for truth, whereas questions or commands are not, by arguing that it is only in the case of assertions that we aim for convergence. Truth is a device for expressing agreement, and there would be no need for this device unless we aim for agreement. If this is true, then discourse on basic taste would fail to be apt for truth, and it would fail to be assertoric. However, Shapiro's account shows that there still is a use for this device of expressing agreement even if agreement is not our aim. This is why we could think of discourse on basic taste as apt for truth even though it does not aim for truth. What kind of truth such discourse is apt for and how it differs from truth for shared inquiry is the topic of Chapter 4.

7. Conclusion

In this chapter, I have used the work of Huw Price to show that truth is used as a kind of coordination device and that aiming for truth is useful insofar as aiming for convergence is practically beneficial. However, I part ways with Price by rejecting his presupposition that this aim is practically beneficial for all assertoric discourse. I have introduced the distinction between shared inquiry, which aims for truth, and personal inquiry, which lacks this aim. I then focused on the inquiry into matters of basic taste and argued, both on a conceptual and practical level, that this inquiry should count as personal inquiry. The next chapters spell out the different theoretical implications of this distinction. Chapter 3 links the distinction between shared and personal inquiry to different accounts of peer disagreement. Chapter 4 asks what difference the distinction makes to what truth consists in, and it develops a kind of pluralist, or rather dualist, theory of truth at the level of use. Chapter 5 discusses the effects that this distinction has on the way we must conceive of rationality and knowledge. Finally, Chapter 6 applies this framework to the case of philosophical inquiry and argues that philosophy should count as a form of personal inquiry, from which it follows that philosophy does not aim for truth.

3. Peer Disagreement and the Aim for Truth

1. Introduction

In the previous chapter, I have introduced the distinction between shared and personal inquiry; the former aims for truth, whereas the latter does not. The present chapter connects this distinction to the literature on peer disagreement. I argue that it is generally rational to conciliate in the face of peer disagreement in the case of shared inquiry, whereas we can be rationally entitled to remain steadfast in the face of peer disagreement in the case of personal inquiry.

Section 2 introduces the debate between the conciliatory and steadfast views. Section 3 argues that this debate turns on whether it is rational to assume that other individuals experience and process things in the same way. Section 4 makes the connection with the distinction between shared and personal inquiry, and with the aim for truth. It also argues against contextualism and relativism. In section 5, I develop the arguments for why the conciliatory view is useful in the case of scientific peer disagreement and why scientific inquiry is a paragon of a shared inquiry. Then, in section 6, I discuss several "unscientific" forms of inquiry, starting with the inquiry on matters of basic taste and then extending the discussion to aesthetic, moral, and spiritual inquiry. I argue why they count as personal rather than shared forms of inquiry. I conclude by stating my intention to extend this framework to philosophical inquiry as well, which I do in Chapter 6.

2. Peer Disagreement

Peer disagreement is, obviously enough, disagreement between peers. Yet what are peers? A commonly accepted definition has been offered by Thomas Kelly:

> Let us say that two individuals are *epistemic peers* with respect to some question if and only if they satisfy the following two conditions: (i) they are equals with respect to their familiarity with the evidence and arguments which bear on the question, and (ii) they are equals with respect to general epistemic virtues such as intelligence, thoughtfulness, and freedom from bias. (Kelly 2005, 174-175)

These demands for peerhood are rather strong, and one could wonder if peer disagreement actually ever happens in real life or whether it is a purely hypothetical notion. Nathan King (2012) argues that it happens only very infrequently that two people are actually peers. In real cases of disagreement, the interlocutors almost never actually share the same epistemic context or the same epistemic virtues. Yet, in deciding which response to peer disagreement is rational, it should not matter whether the disagreement is actually between epistemic peers. All that matters is whether the participant in the disagreement has antecedent reasons for discounting the opinion of their interlocutor; that is, reasons that are independent of the disagreement in question. If they do not have such reasons, then they rationally ought to treat their interlocutor as their peer.

One might argue that, as long as the participant is aware of the fact that it is very rare that two people share the same epistemic contexts and virtues, they should give a high credence to the belief that their interlocutor is not their peer. However, in those cases, one still would not know whether their interlocutor is epistemically superior or inferior. Therefore, if we are well enough acquainted with our interlocutor and know of no reason for why they or their evidential state might be epistemically superior or inferior concerning the topic in question, we should treat our disagreement as an instance of disagreement between peers.

There are unusual cases in which our supposed peer disagrees with us about something that we are so certain of that the fact of our disagreement is itself enough to discount the idea that our interlocutor is our peer, e.g., the case that our interlocutor disagrees with us about

whether 2 + 2 = 4. (Lackey 2008) This phenomenon has been called "peer-demotion". (Comesaña 2012) However, it should be clear that these cases are rare and that it would be irrational to employ peer-demotion as a strategy in ordinary cases of disagreement.

A question that dominates the literature on the topic of peer disagreement is whether it is rational to stick to your guns when faced with peer disagreement. Intuitively, if we take our interlocutor to be our peer, we should give as much weight to their view as to our own. However, this line of thinking has been argued against by, for instance, Kelly (2005), who thinks that we should be allowed to remain steadfast in our beliefs.

This view has appropriately been called the "steadfast" view, whereas the opposing view is the "conciliatory" view. (Elga 2010) The strictest version of the conciliatory view tells us that inquirers must always give equal weight to the views of their peers, meaning that they must split the difference; if I give *p* a credence of 0.3 and you give *p* a credence of 0.6, then, once we learn of this disagreement and our respective views concerning *p*, we both ought to give *p* a credence of 0.45. Yet, there are weaker versions of the view as well that are perhaps more appealing. I call any view which is not steadfast a conciliatory view — all views that say that the mere presence of disagreement should have *some* impact on our doxastic attitudes count as conciliatory views. I shall assume that the conciliatory strategy to disagreement would generally lead to a decrease of confidence with which we hold our beliefs, and in the most clearcut case, it will lead to a suspense of judgment concerning the topic of disagreement.

Kelly shows that the conciliatory view appeals to the idea of symmetry: if there is no reason to believe that my epistemic context or virtues are better or worse than yours, then we must conclude that your opinion is as good as mine because to suppose otherwise is to introduce an unwarranted asymmetry. However, Kelly argues that the situation would never really be symmetrical to begin with because by disagreeing, one does not merely take one's opponent to have reached a wrong conclusion, but one takes one's opponent to have reached a wrong

conclusion because of having misevaluated the available evidence. That is, one believes that the opponent has given too much weight to the evidence that supports their conclusion and not enough weight to the evidence that undermines it.

Now, it might seem that, in this case, the opponent is not an epistemic peer because one judges that the opponent is worse at evaluating the available evidence. However, Kelly argues that the opponent could still be seen as a peer because one deems them equally competent at evaluating evidence *in general*. He uses the analogy of chess players of equal skill, arguing that we should not expect them to draw all their matches. Similarly, two epistemic peers do not always draw their arguments. Sometimes there is a winner. Being peers merely means you expect to win about as often as you lose. According to Kelly, it is rational to assume, whenever we have evaluated our peer's arguments and have found ourselves unmoved by them, that we have emerged victorious from the particular disagreement. (Kelly 2005)

I think this argument is unappealing because if we are committed to thinking that we disagree *with a peer*, then we are likewise committed to thinking that there is no reason that our evaluation of the evidence is better than theirs. Hence, the situation remains symmetrical, and as a result, one must suspend judgment. Of course, in practice, this rarely happens because it is very difficult to step back from one's own reasoning and one's own evaluation of the evidence. In practice, the opponent's evaluation of the evidence will strike us as determinably worse, and we will feel the perhaps irresistible urge to stick to our guns. However, in order to determine whether one is justified in believing what one believes, we need to determine what one *ought* to believe given the presence of peer disagreement, and it seems to me that Kelly's argument fails to show why we are justified in privileging our own evaluation of the evidence and discounting the opponent's evaluation.

Nevertheless, it is not the case that we should always conciliate. As I argue in the next section, whether conciliation or steadfastness is rational depends on the context. There is room

for the steadfast view in those cases in which it is not clear that there is only one right evaluation of the evidence, i.e., those cases in which multiple different evaluations of the evidence might be permissible. Yet, in those cases, the situation is still symmetrical: both sides of the disagreement are allowed to remain steadfast.

There is another kind of steadfastness that is asymmetrical and that follows from a kind of externalism. According to this view, there is only one side of the disagreement that ought to stick to their guns: the side that, in fact, reasons correctly. (De Ridder 2014) However, in practice, we do not know which side actually reasons correctly. We shall always think that it is our side, but it would not be practically rational to act on this belief by always remaining steadfast because we would often turn out to be wrong. Therefore, this kind of external notion of rationality cannot and should not inform our practices. The question I aim to answer is which response to peer disagreement is rational from a first-person point of view. In what follows, I shall argue that the answer to this question depends on whether our inquiry is shared or personal and that it thus depends on whether our inquiry aims for truth.

3. Sameness and Uniqueness

Kelly has argued that the plausibility of the conciliatory approach depends on a controversial principle, namely the principle of "uniqueness". (Kelly 2010) This is how Roger White formulates the principle:

> **Uniqueness:** Given one's total evidence, there is a unique rational doxastic attitude that one can take to any proposition. (White 2005, 445)

The principle denies epistemic permissiveness, i.e., it denies that a single evidential state allows for multiple rational attitudes. Thus, if uniqueness holds, then, given any proposition *p*, one's evidential state either favors *p* — in which case it would only be rational to believe *p* — opposes *p* — in which case it would only be rational not to believe *p* — or has no bearing on *p* — in which case it would only be rational to suspend judgment.

The principle would be too strict if it said that it would only be rational, for instance, to give exactly a 0.86 credence to one's belief that *p*. In that case, our doxastic attitudes would almost never be rational. Instead, there must be some level of permissibility concerning one's doxastic attitude. For instance, if the evidence points in favor of *p*, any doxastic attitude that counts as believing *p* would be rational to a certain extent.

However, even on that interpretation, the uniqueness principle is too strict to be plausible. For instance, it seems perfectly reasonable that two inquirers evaluate the evidence differently because of collateral assumptions or background beliefs. These assumptions or beliefs could be said to be a part of the evidence, but if they are it would never really be the case that two different inquirers evaluate the *same* evidence. Also, even if they were, it could still be the case, for instance, that though they evaluate the evidence in the same way, one of them is slightly more cautious than the other, such that Inquirer A ends up believing *p* whereas Inquirer B is more cautious and suspends judgment, even though they both give *p* a credence of, say, 0.7. Unless B is excessively cautious or A is excessively careless, it would not be intuitive to say that either of them is irrational. This is why the uniqueness principle is generally rejected. Take, for instance, Gideon Rosen, who says:

> It should be obvious that reasonable people can disagree, even when confronted with the same body of evidence. When a jury or a court is divided in a difficult case, the mere fact of disagreement does not mean that someone is being unreasonable. (Rosen 2001, 71)

Given the controversial nature of uniqueness, the opponents of the conciliatory view, like Kelly or Rosen, argue against this view by arguing that it implies uniqueness. The mere fact that someone of equal competence could look at the same evidence and form an opposing belief about the proposition in question should not diminish our confidence in our beliefs *unless* there is only one right belief to be had.

I shall put the matter of whether the conciliatory view implies uniqueness aside. However, I do want to point out that the inverse relation holds: uniqueness leads to the conciliatory view. Peer disagreement is indicative of the fact that the evidence does not unequivocally point in one direction; parts of the evidential state might point in favor of *p*, while other parts might conflict with *p*. A supporter of the steadfast view could readily accept this as long as they reject the uniqueness principle. They could say that the evidential state leaves room for different rational evaluations and that we are allowed to stick to our own evaluation. However, if one accepts the uniqueness principle, then one believes that once we realize that the evidence does not unequivocally support a positive or negative doxastic attitude, the only rational attitude to be had is a suspension of judgment. This is what the conciliatory view prescribes.

The above explanation shows that the conciliatory view seems appealing insofar as we do not allow different evaluations of the evidence to coexist. Whenever we are presented with a case of disagreement, it casts doubt on our presupposition that we have evaluated the evidence correctly because either (1) the evidence is inconclusive or (2) one of us has made a mistake. Concerning (1), why would inconclusive evidence not allow us to both maintain our respective beliefs? Why should it lead to a suspension of judgment? This is because we believe that once all the evidence is in, there is only one correct evaluation to be had. Concerning (2), why would disagreement be indicative of a fault? Well, again, this is because we believe there is only one correct evaluation to be had.

Uniqueness is needlessly strong because it says that there is only one *rational* doxastic attitude, whereas it merely needs to say that there is only one *correct* doxastic attitude. Once we believe this, it would be rational for us to generally suspend our judgment when faced with disagreement (unless, for instance, it is abundantly clear that the fault is committed by our interlocutor, as it is when they say that 2 + 2 = 5). If we disagree, and we believe that there is only one correct doxastic attitude to be had, then either you are mistaken, or I am, or our evidence is inconclusive. If I have no good reason to believe that you are any worse of an inquirer than I am, then I must believe that the mistake might be mine or that our evidence is inconclusive. In either case, I would suspend judgment and continue carrying out my inquiry, e.g., by collecting more evidence until the relevant shortcoming has been exposed.

This seems much less controversial than uniqueness. However, it does depend on the assumption of *sameness*:

Sameness: We generally experience and process things in the same way.

I say that we *generally* experience and process things in the same way because this would fail to be the case if one of us has impaired cognitive or observational capacities, such as dementia or near-sightedness. Such cases do not challenge the assumption of sameness since they can be treated as anomalies.

The idea behind this assumption is that, for instance, if I can see that my couch is red, then so could you if you were in my position. If I can figure out that the sum of 25 and 13 is 38, then so could you. This should be fairly obvious and uncontroversial, though it by no means holds for everything. For instance, I could see my couch as ugly, but it would not be wise to assume that you would also see it that way.

This is not to say that when it comes to simple perceptual beliefs, such as whether my couch is red, there would never be any instances of disagreement. Still, they will be sufficiently rare for it to be *rational* for us to assume that we generally see things in the same way, and to treat these cases of disagreement as anomalies that could easily be resolved under closer scrutiny. It is rational to do so because this assumption will generally prove useful. However, when it comes to our aesthetic experiences, this assumption of sameness would pay no dividends since disagreements about such experiences are not sufficiently rare.

Without the assumption of sameness, the idea that there would be one correct evaluation of the evidence would not make sense. If we experience or process our experiences differently, then there would not be one unique way in which evidence is to be evaluated or interpreted. Suppose we inquire into the tastiness of a cake by sharing a slice. Our familiarity with the evidence is, in some sense, the same because we have eaten from the same slice of cake. However, I have no reason to assume that the way you process this evidence will be the same as the way I do so, and, therefore, I have no reason to believe that the evidence shall conclusively point in favor or against the hypothesis that the cake is tasty. Nor do I have a reason to believe that this is due to the incompleteness of the evidence. Instead, it is reasonable for me to expect that once we disagree, there is no promise of us settling this disagreement in the future. A more productive attitude would be to exercise tolerance toward disagreement and thereby agree to disagree.

As I argued in the previous chapter, if the assumption of sameness is not useful, then suspending judgment in the face of disagreement would be a rather perilous, and thus irrational, exercise. Doubt could be a useful tool that prompts us to find beliefs that serve us better in the long run, but too much doubt would be detrimental to our practices. The general idea of suspending judgment is that one trades in an adequate though plausibly less-than-perfect belief or set of beliefs for the promise of a better one. In the case of science, for instance, one gives

up their contested theory hoping to gain a theory that everyone finds agreeable. However, in the case of taste, it is highly improbable that our inquiry leads us to beliefs that we can all agree on. Therefore, if we suspend judgment in the face of disagreement, all we are doing is trading in adequate beliefs for an excess of doubt.

Many philosophical or metaphysical distinctions have been made to try to explain why sameness is a rational assumption for some matters, such as for the question of whether Marmite is savory, while it is not rational to make this assumption when it comes to other matters, such as for the question of whether Marmite is tasty. For instance, the property of being savory has been said to be more real than the property of tastiness; savory is said to be descriptive, whereas tasty is evaluative; the former represents reality, whereas the latter expresses a sentiment, etc. However, I want to treat these differences as, in some sense, primitive. The rationality of the assumption of sameness depends on the utility of this assumption. It is a brute empirical fact that we are disposed to react similarly in the case of basic observations, computations, episodic memory, etc., while we all tend to be differently inclined when it comes to our aesthetic, comical, or moral reactions. What I want to focus on is how we must respond to this difference and, in particular, how this difference warrants a distinction in how we should act in response to peer disagreement.

4. The Aim for Truth

At this point, the issue of how to respond to peer disagreement can be connected to the themes of the last chapter. In the previous chapter, I already briefly mentioned the assumption of sameness. I argued that it is only useful to pool our resources and turn our inquiry into a shared inquiry insofar as it is useful to assume sameness. The aim for truth facilitates this pooling of our resources since it turns the question of which beliefs to endorse into a shared concern.

4.1 Shared and Personal Inquiry

Other authors have noted contrasts that are in some ways similar to the one that I have noted between those disagreements that are guided by the assumption of sameness and those that are not. For instance, there is a contrast between ordinary everyday disagreements and disagreements that seem more intractable or fundamental. This is the distinction between "trivial" and "substantial" disagreements (in Robert Mark Simpson's (2013) terms) or between "ordinary" and "deep" disagreements (in Neil Levy's (2021) terms). I shall now argue that the assumption of sameness can explain these conceptual contrasts between disagreements.

Trivial or ordinary disagreements are disagreements about everyday basic observations or basic calculations, episodic memory, etc., whereas substantial or deep disagreements are disagreements about certain moral or political matters. The intuition that both Simpson and Levy build on is that a conciliatory approach works for the former cases but not for the latter cases. Levy uses the work of Katia Vavova (2014) to argue that this has to do with the surprisingness of the disagreements. If I see a horse in the distance, I would not expect you to say it is a cow. However, if I make a claim about a hotly contested political issue, I can expect the disagreement that ensues. Therefore, this disagreement would not offer me new information, i.e., it would not give me any new higher-order evidence in light of which I should reconsider my views. Since I already know people will disagree when it comes to politics or morality, I should not have to reduce my confidence in my beliefs when I indeed encounter such disagreements (unless, of course, I am persuaded by my opponent's arguments).

There are two things to note about this point. The first is that this can easily be extended to, for instance, cases of basic taste, aesthetics, comedy, etc., although "substantive" or "deep" would then seem like misnomers. The second thing to note is that this distinction of whether a disagreement is surprising is not very clear-cut. For instance, when we have to perform some moderately complicated calculations, and neither of us has a higher degree in mathematics, we

should hardly find it surprising that we end up disagreeing on what the right answer should be. Conversely, it could be extremely surprising to find out that someone in one's social circle, who enjoyed a similar upbringing, and with whom you are very amicable, turns out to have some outlandish moral or political convictions.

I believe these distinctions between ordinary and deep disagreements do not have that much to do with surprise *per se*. What underlies these differences is the acceptance or rejection of the assumption of sameness of our experiences and the relevant cognitive faculties, which then gives rise to the conviction that there is a unique doxastic attitude that is the right one to have. It is true that we do not have to conciliate in the face of moral or political peer disagreement. It is also true that this is because we do not expect that the opinions of our peers will generally align (we do not assume sameness). However, this does not mean that there could not be particular cases that strike us as surprising, and it also does not mean that we should conciliate in those particular cases. It is perfectly consistent to reject the assumption of sameness while still anticipating agreement with particular likeminded individuals. Yet, by rejecting sameness, we reject that there is one correct doxastic attitude to be had. We thereby adopt a permissive attitude, according to which people are allowed to evaluate or interpret matters in their own way.[3] This explains why there are many cases of disagreement about aesthetic matters, moral matters, matters of comedy or basic taste, etc., for which it intuitively seems that we are entitled to stick to our guns.

Alex Worsnip (2019) argues that we can conceive of disagreement as interpersonal incoherence. Disagreement bothers us, and we strive to resolve it because we strive for *inter*personal coherence, not just *intra*personal coherence. However, why would we pursue this?

[3] There are of course limits to this permissiveness. In Chapter 5, I argue that people are only rationally entitled to their own evaluations insofar as they can be interpreted as adhering to some structure or as belonging to some coherent set of beliefs.

Sure, the belief that pistachio ice cream is tasty does not cohere with the belief that pistachio ice cream is disgusting, and an individual who holds both beliefs is incoherent and thus irrational. Yet, it would not be irrational for two distinct individuals to endorse one of these beliefs each. Why should we be bothered by this?

Our pursuit for interpersonal coherence is explained by, and depends on, the assumption of sameness. If you generally experience the world in the same way as I would, and you process these experiences in the same way as I would, then we expect to have shared beliefs about the world. Thus, it should be possible to pool our beliefs and create a coherent system. Any incoherence that arises would be indicative of the fact that at least one of us has gone astray. Yet, without this assumption of sameness, there would be no expectation that we could pool our beliefs into a coherent system. In that case, the mere fact of peer disagreement should not bother us.

This is a practical difference since the adoption of the assumption of sameness turns our inquiry into a *shared* inquiry, in which we aim to converge toward the same doxastic attitudes. In this kind of inquiry, it would generally be rational to conciliate to aid this converge. The rejection of the assumption of sameness turns our inquiry into a *personal* inquiry, in which we are generally allowed to remain steadfast in the face of peer disagreement, provided that we find the opposing side's arguments unconvincing. This is because we do not aim to converge in our opinions. In fact, in the case of taste and morality, we often use our opinions to distinguish ourselves from other individuals or groups and to thereby articulate our (social) identity. As Levy argues, in the case of morality (but, again, also, though maybe to a lesser extent, in the case of taste), our opinions are central to our identity, which is why we are not inclined to give up on them and suspend judgment as easily as we would in the case of ordinary everyday beliefs. (Levy 2021, 137) And nor should we, for the assumption of sameness is

impractical and thus irrational for such cases, which is why the conciliatory approach would not be rational either.

The contrast between shared and personal inquiry accounts for why the intuitive appeal of the conciliatory response seems to evaporate once we shift the focus from ordinary everyday disagreement or scientific disagreement to disagreement about, for instance, matters of taste. Conciliation is a useful tool for reaching convergence; insofar as we are bothered by differences of belief, it would be rational to second-guess one's beliefs in the face of disagreement. However, it is not always practically advantageous to aim for convergence, and it is, therefore, not always rational to conciliate. In the case of matters of taste, there is no hope or expectation that our beliefs will converge in the long run, which is why it would be rather pointless to be bothered by differences of belief, and why it seems more rational to simply agree to disagree instead.

Since the aim for convergence is encapsulated by the aim for truth, it is the aim for truth that engenders the conciliatory approach to peer disagreement. It is this aim that provides the fault to our disagreements that prompts us to conciliate. Furthermore, for the aim for truth to have this effect on our disagreements, it must attribute an *a priori* fault to every instance of disagreement: we suspend judgment because, even in the cases in which we do not know of any particular mistakes or shortcomings, we still think that one of us must be at fault. This fault could only be interpreted as an *alethic fault*, i.e., the fault of asserting a falsehood. The mark of personal inquiries for which steadfastness could be appropriate would therefore be the lack of this alethic fault.

4.2 Relativism and Contextualism

I have argued that the difference between disagreements belonging to shared inquiry and those that belong to personal inquiry is best explained by the presence or lack of the aim for truth. Personal inquiry does not aim for truth, which is why we could be rationally entitled to remain steadfast in the face of peer disagreement. However, there are rival explanations, such as contextualism and relativism.

Contextualists argue that when we make a claim like "Pistachio ice cream is tasty", the content of the claim depends on the context of the utterance, and specifically on the person who makes the claim. The proposition expressed by the utterance is therefore bound to the speaker in a way that is explicitly stated by a claim like "For me, pistachio ice cream is tasty". This way, we could still aim for truth, but there would generally be no reason to reduce the credence of our beliefs in the face of peer disagreement since we could both be right.

Contextualism has been argued against because it turns disagreements about, for instance, matters of taste into a kind of pseudo-disagreement. (Kölbel 2004) On the contextualist picture, we aren't really disagreeing at all. This is because the propositions expressed might *look* incompatible — "Pistachio ice cream is tasty", "No, it isn't" — but if the contextualist is right, they aren't actually incompatible at all. So, the disagreement would be something like the situation in which someone in New York is skyping someone in Scotland, and the one says to the other, "It is a very hot day", to which the other replies "It is frigid outside". (Wright 2021)

Relativism about truth has a better chance of accounting for disagreements belonging to personal inquiry. The most popular and defensible form of relativism is that of John MacFarlane (2014). This kind of relativism focuses not on the speaker but on the assessor of a claim; the truth of a proposition is relative to whoever assesses it. This relativity explains why

disagreement can be faultless while ensuring that the disagreement still consists of incompatible contents.

One of the problems of this approach is that it might be difficult to extend it to moral disagreements. I haven't yet argued why moral inquiry should count as a personal inquiry, but I shall do so in section 6. For now, please allow me to assume that moral inquiry is personal. While we can agree to disagree in the case of taste, given that we recognize that we employ different standards of assessment, in the case of moral disagreement, we would not be inclined to agree to disagree since we would not simply accept that our interlocutor employs their own standard. We would want to say that there is only one truth in the case of morality. For instance, if we believe that slavery is wrong, we believe it to be wrong for everyone anywhere. Sure enough, there are also relativistic tendencies in the case of morality; we understand that different cultures might have different values. However, at the level of a disagreement between individuals who endorse fundamentally different moral values, these individuals involved would reject that their interlocutor is entitled to their own truth. The matter of which standards are the right ones for moral assessment is itself a part of the debate.

Furthermore, while the previous chapter has shown that the aim for a shared truth can be practically efficacious, the aim for a kind of personal truth, like the one relativists argue for, seems to lack practical bearings. What would distinguish it from merely aiming for beliefs that work best for me?

Instead of this relativist picture, which has all of us aiming for our own truths, it would be simpler to state that we do not aim for truth at all. John Davis suggests precisely this: get rid of the idea that there must be an alethic standard in addition to the standards of assertion. (Davis 2015) The only standards that guide personal inquiry are those that make up the rules of the language game in question. This does not mean that we reject that moral or aesthetic discourse is apt for truth. There is no need to hold such a radical and counterintuitive view. As I have

argued in the previous section, truth, as a device for registering agreement, can still be useful even if we do not aim for agreement. What I reject is that aesthetic or moral inquiry *aims* for truth. Wright (2021) endorses this solution for the inquiry into matters of basic taste. I shall argue in section 6 that the solution also holds for other personal inquiries, such as moral inquiry. Rejecting that personal inquiry aims at truth is all that is needed to make sense of the mechanics of peer disagreement in personal inquiry.

5. Shared Inquiry

I have argued that there is a distinction to be made between inquiries that are shared and those that are personal. The former kind of inquiry has *sameness* as its regulative assumption, whereas the latter does not. If *sameness* is a rational assumption to make, then it would also be rational to try to pool our resources, work together, and try to converge in our beliefs. This convergence would be aided by some degree of conciliation in the face of disagreement, and, as such, the assumption of *sameness* can turn conciliation into a (generally) rational response to peer disagreement. However, if our inquiry is personal, then disagreement should not bother us as much; why should I care that you believe pistachio ice cream is disgusting while I believe that it is tasty? Thus, if our inquiry is personal, then we could be rationally entitled to remain steadfast in the case of peer disagreement.

The clearest example of a shared inquiry is scientific inquiry. The scientific method centers around shared standards of evaluation. To ensure this, it tries to break down complex issues into simple observations and computations, the kind of observations and computations for which the assumption of sameness is reasonable.

There are clear practical advantages to treating scientific inquiry as a shared enterprise. Given the assumption of sameness, we can trust a peer's evaluation of the evidence. This allows

us to pool our resources. We could be stubborn and stick with our own idiosyncratic personal scientific theories, just like we stick with our own personal tastes and moral convictions. There are, unfortunately, many people out there nowadays who treat scientific theories as a kind of personal moral convictions. Yet, constructive collaboration requires that we exercise the virtue of intellectual modesty, which tells us that our peers are just as likely to be right, even if we cannot see it. Combined with the guiding assumption of the scientific method, the assumption of sameness, this makes conciliation a generally rational response to scientific disagreement.

The effectiveness of conciliation has been tested using agent-based models. I shall not pretend to understand the science behind these models, but I do think an interpretation of their results will be useful. It is my understanding that these models represent various agents who try to discover the value of a parameter, which is set to be somewhere between 0 and 1. Each agent starts with a random value, which represents their initial belief, and they update this value based on their research. Conciliatory populations split the difference between their value and the value of the peers they encounter. (Šešelja 2019) Igor Douven (2010) has shown that concerning easy inquiries, such as simple computations or observations, populations that split the difference when they encounter peer disagreement converge to the truth more quickly than steadfast populations. However, in the case of complex scientific inquiry, steadfast populations get within a moderate distance from the truth more quickly than conciliatory populations. Yet, while the progress of conciliatory populations is slower, they do end up closer to the value of the parameter. Thus, the conciliatory view trades speed for accuracy. This same effect holds if the value of the parameter is interpreted not as the true value but if, say, it represents the theory that is most practically advantageous.

While this shows the utility of the conciliatory response, it alone cannot determine that conciliation is the rational response. For conciliation to be generally rational, it must not only *actually* be useful, but it must also *seem* useful, i.e., this utility should be evident to the inquirer.

Luckily, Douven's results align with what we intuitively might expect. If scientific inquiry is personal, like aesthetic or moral inquiry, and we give a lot of weight to our own evaluations of scientific theories, just like in the case of our own aesthetic or moral evaluations, then these theories would have to be more rudimentary. It would be quicker to develop a rudimentary scientific theory that is moderately practically adequate if we are not led to a suspension of judgment whenever we are faced with disagreement.

However, this would not give us the most practically advantageous theory. It is the pressure to resolve our disagreements that carries our inquiry further. In contemporary science, we send robots to Mars to try to confirm our theories. The immediate practical payoff of doing so is negligible, and the costs are high. However, the expectation is that, more often than not, the long-term advantages do justify these costs. It is the intellectual curiosity, that is conveyed by us suspending our judgment and continuing our inquiry, that has given us our modern technology and medical sciences. It has given us theories that take longer to develop but that prove more practically beneficial in the long run. This is a reasonable expectation, and as such, it would be rational, in general, to suspend judgment in the case of scientific disagreement.

Rogier De Langhe (2013) shows that Douven's results only hold true for those populations for which the best beliefs — those that inquiry aims for — are the same for everyone. Again, this seems to make a lot of sense. Suppose that the best beliefs are those that work best in practice. If what works best for me, e.g., the belief that pistachio ice cream is tasty, is not what works best for you, then it would be better for me to remain steadfast in the cases in which we disagree. Thus, the conciliatory response would only be a generally rational response if the assumption of sameness is rational as well. This is exactly what I have argued for.

Douven's result can be interpreted as showing that, in the case of science, it is beneficial, in the long run, to conciliate in the face of peer disagreement. I have argued that this utility

does not seem to be due to some obfuscated accidental factors but that instead, the practical payoff of conciliation in the case of science is something that can reasonably be expected. Thus, it is practically rational to conciliate in the face of scientific peer disagreement. I have argued that the conciliatory view on peer disagreement presupposes that we aim for truth. Given the Jamesian assumption that we should aim for truth insofar as it pays to do so, it can be argued that we *ought* to aim for truth in the case of science and, thus, that we ought to treat science as a shared inquiry.

I believe this result holds not just for the empirical sciences but also for non-empirical sciences like mathematics, and it holds not just for the "hard" sciences but also for the social sciences. This is perhaps controversial, given, for example, the replication crisis in psychology: the results of many psychological studies are difficult, if not impossible, to replicate. This crisis challenges the assumption of sameness. Nevertheless, psychology does continue to operate under this assumption, and if it is shown that the crisis is so severe that it would simply be irrational to assume sameness, then it would cease to be *scientific* altogether. Thus, every inquiry that can justifiably bear the label 'scientific' is such that it is rational to carry it out as a shared inquiry.

6. Personal Inquiries

Given the above, it might seem that we could simply define personal inquiry as "unscientific" inquiry. This distinction might seem similar to that of the logical empiricists or Wittgenstein of the *Tractatus* (1921). I do not claim, like Wittgenstein does, that inquiries into taste, aesthetics, morality, religion, or philosophy are senseless, but I do argue that it would be a mistake to try to fit them into the scientific mold. They should not be treated as shared inquiries

but as personal inquiries. To conclude this chapter, I shall discuss these cases in turn to add detail to my thesis that they are all instances of personal inquiry.

6.1 Disagreement about Basic Taste

If conciliation is driven by the presence of an alethic fault, then the clearest example of a disagreement for which steadfastness is generally the rational response would be an instance of disagreement that contains no fault at all. A lot has been written on the relation between faultless disagreement and the issues of peer disagreement. Faultless disagreements are, as the term suggests, instances of genuine disagreement in which no fault has been committed by either side. (Kölbel 2004) According to Wright, a steadfast response is the right response to faultless disagreement:

> [F]aultless disagreements are to be examples where, in the terminology of the peer disagreement literature, *steadfastness* can be an appropriate response. (Wright 2021, 435)

The idea is as follows: if the disagreement contains no fault at all, why should we give up our opinions in light of it? In that case, the disagreement would not indicate that our opinions are somehow at fault, and if it gives us no reason to believe our opinions are at fault, we will not feel the pressure to revise or suspend them.

The most typical case of faultless disagreement, used in the previous chapter as well, concerns disagreements about basic taste. However, this is not to say that all disagreements about taste are faultless. For instance, if we take someone to have a superior taste, and our taste judgments do not align, then this does give us a reason to second-guess our judgments. It is, of course, an interesting question what would make someone superior in their taste judgments, or

what would make someone a peer, for that matter. Kelly mentions epistemic virtues such as intelligence to define peerhood (Kelly 2005), but that doesn't seem to matter much in the case of taste. The answer is likely to be strongly context-dependent. For instance, if the judgments in question are about what is "classy", we might treat those who belong to the same social class as peers. What this shows is that, as I have argued in section 2, it does not matter who actually is our peer, but only who we take to be our peer.

6.2 Aesthetic Disagreement

The instances of faultless disagreement discussed in the previous chapter are basic or primitive; "Pistachio ice cream is tasty", "No, it isn't". Yet, as we transition into more intricate aesthetic matters, there seems to be more of a possibility for a fault. For instance, you could fashion yourself a connoisseur when it comes to whisky, and you might say that those who dislike whisky altogether just haven't sufficiently developed their palate, or they just don't *understand* what makes a whisky good or enjoyable. These do seem to be attributions of genuine faults. In other aesthetic matters, such as literature, cinema, painting, music, fashion, design, etc., such attributions of fault occur very frequently. In various aesthetic matters, we seem to allow for the existence of a kind of experts, whom we call critics. This might suggest that there are some shared standards and that perhaps aesthetic inquiry can be shared.

However, aesthetic expertise is nothing like scientific expertise. It is clear that the books we like, the movies we watch, the music we listen to, and the clothes we wear are all expressions of our identity. And whom we count as experts is merely an extension of this expression; it is a personal decision with whom we want our tastes to align, which people or groups of people we take to have good taste in music, cinema, literature, or fashion.

A clear sign that aesthetic inquiry is personal is that, in order to make a genuine aesthetic judgment, it is required that we are personally acquainted with the object of our judgment. (Ferrari 2022, 89) It would be a bit off, for instance, if I were to say that whisky is tasty or that *Death of a Salesman* is a masterful play if I have never had a sip of whisky or been to the theater. In that case, I would be faking my aesthetic judgments, and it would be more honest to say: "I've *heard* that *Death of a Salesman* is a masterful play." Yet, it seems perfectly in order to judge that the speed of sound is 343 m/s even if I have never broken through the sound barrier myself or have conducted experiments to determine the speed of sound. In that case, it is allowed for me to rely on the observations and calculations of others since the inquiry is shared.

Thus, we may be swayed by experts in our aesthetic judgments, but our aesthetic inquiry is still personal because it is a personal choice whom we count as experts and because we can never fully rely on the expertise of others in our aesthetic judgments. We must have experienced and understood for ourselves why the play, painting, book, etc., is good or bad. This is consistent with the steadfast view. The conciliatory view states that the mere presence of disagreement should have an impact on what we believe and on how confident we are in this belief. The obverse of this is that the presence of overwhelming agreement, and thus the complete absence of disagreement, could also direct our beliefs. However, the steadfast view states that our attitudes should only be affected if we appreciate the force of the arguments and reasons behind our interlocutor's beliefs. This means that we should not be swayed by the mere presence of disagreement but also that the mere presence of consensus should not hold sway over us either. The ultimate arbiter is our own understanding. Even if the play received a standing ovation, if I didn't like it myself, and nobody could help me understand why it was a good play, it would be wrong for me to assert that the play was a masterpiece.

6.3 Moral Disagreement

While faultless disagreement might be the clearest example of disagreement for which steadfastness is a permissible response, there are domains of discourse which, like discourse on basic taste, differ from scientific discourse yet that do not contain the possibility of faultless disagreement at all, such as moral discourse. In spite of the dissimilarities between science and morality, there has been some pushback against the assimilation of the moral domain with, for instance, the domain of taste. The reason is precisely that moral disagreement is never faultless.

John Eriksson argues that there are strong practical reasons to try to come together in the case of moral disagreement. Eriksson argues that while there might at most be contingent reasons for trying to resolve disagreements about taste, this pressure to coordinate is part and parcel of moral judgments:

> In moral matters, by contrast, it is much more plausible to hold that, if your judgment conflicts with mine, then I am necessarily disposed to challenge it, and that this disposition is part of the judgment itself. This idea can be supported by considerations of the following sort. Moral thought and talk evolved in order to help us reap the benefits of cooperation. This requires that we coordinate our attitudes. However, coordination doesn't happen by itself. We therefore need some mechanism by which it is achieved. Simply having conflicting attitudes is consistent with indifference to other parties' attitudes (as often is the case in matters of taste). Having a disposition to challenge conflicting judgments seems to be a way around this problem. Of course, challenging a conflicting moral judgment doesn't by itself bring about agreement, but it is what triggers normative discussion. Nevertheless, by building this disposition into moral judgments evolution found a way for us to move towards consensus and thereby helped us disentangle practical problems, i.e., questions concerning how to live. Hence, to endorse a moral judgment is, in part, to be disposed to challenge conflicting judgments. This is why we cannot really turn the coordinating function off. (Eriksson 2016, 789)

Eriksson does not make mention of conciliation or steadfastness in this paper. Instead, he focuses on the normative force of moral judgments, which seems lacking in the case of taste judgments: when we make a moral judgment, we are thereby also judging that others ought to agree with us. However, demanding that others agree with us is not a productive strategy for coordinating our moral judgments unless we are willing to meet them halfway, and that is not something we appear to be willing to do. Moral and political debates are often as entrenched as any debate could possibly be. There aren't many of us who apply the virtue of intellectual modesty to our moral practices. Rather, they are typical examples of an "I am right, they are wrong" kind of mindset.

Eriksson does make a good point that the coordination of our moral convictions *would be* incredibly valuable to any society. He then makes a scientific speculation by arguing that evolution has enabled us to engage in moral inquiry in a cooperative fashion and that it has thereby enabled us to settle our moral disputes and coordinate our moral beliefs. However, I would speculate that this is not at all how we have evolved. Evolution has sadly disappointed us. Instead of employing the mechanism of conciliation, it has instilled a kind of tribalism within us: when it comes to our own tribe — our own narrowly defined society — our moral values are largely homogeneous. They are the marks we use to conceive of the very identity of our tribe or society. When faced with opposition, we welcome agreement, but only on our terms. When I disagree with someone on whether abortion is wrong, I do passionately want us to agree on the matter, but I am not willing to give up my belief to facilitate this agreement.

The question at issue here is, of course, not whether evolution has given us a steadfast or a conciliatory mindset with respect to moral disagreement but whether we *should* have a steadfast or a conciliatory mindset. Which approach would be more useful; which would be more rational? At first, it seems that it would clearly be more useful if we could coordinate our opinions. However, this kind of hypothetical reasoning is impractical. It would have been great

if our moral sentiments aligned just as well as our observation of colors or shapes do. Yet, the truth of the matter is that while it would be practically rational to assume that everyone hears the same bangs, smells the same smokey scent of the gunpowder, and sees the same redness of the blood, it would be precarious to assume that everyone sees the same wrongdoing. Not in the least because for most perceived wrongdoings, there is at least one person who sees things differently: the person holding the gun.

There are, of course, platitudes that are ingrained in all of us, like the platitude that murder is morally impermissible, but murders nevertheless do happen every day. Sometimes they are even sanctioned by a country or a state. In addition to this, we murder millions of animals every day. In the case of the slaughter of animals, there is even a conceptual debate about whether this counts as an act of murder. Similarly, we all agree that stealing is wrong. Yet, capitalism is thriving, and socialists or communists would classify many of the capitalist practices as acts of stealing (acts of stealing the fruits of another person's labor). Thus, saying that these platitudes show that we do have similar moral sentiments is a bit like saying that we have the same aesthetic experiences because we all enjoy things that are beautiful and detest things that are ugly (but the question is, of course, what is to count as beautiful and what is to count as ugly).

I, therefore, maintain that it would not be rational to assume sameness in the case of morality. In some sense, this counts as an attitude of tolerance toward other ways of life. However, we do not exercise this tolerance when faced with disagreement concerning any of our particular moral beliefs, such as our belief about the moral wrongness of slavery. In such a case, we are right, they are wrong, and that's the end of it. When asked why they are wrong, we shall simply say: "because they believe that slavery is morally permissible, and it isn't." This is exactly the way in which moral disagreement differs from scientific disagreement. The difference does not consist in the latter containing a fault while the former is fault*less*, but it

consists in the latter containing a distinctively *alethic* fault while the former contains a *moral* fault. It is part of a moral judgment that we judge those who disagree with us as suffering from a *moral* shortcoming because they endorse beliefs that are not just incorrect but also immoral.

Filippo Ferrari argues that this moral fault is also an alethic fault. He believes that there are many ways in which truth is normative. In the moral case, truth is deontically normative: if we judge *p* to be true, we judge that *p* ought to be believed and that it is thus *wrong* not to believe *p*. Ferrari claims that the difference between disagreement about basic taste and moral disagreement is that this deontic norm is present in the latter but not the former case. He says that the alethic and the moral fault are interrelated:

> [T]he alethic criticism of a contrary judgement about fundamental moral matters merges with a distinctively moral criticism targeting the holder of such a judgement. (Ferrari 2022, 87)

Yet. why must there be an alethic fault at all; why does the moral fault not suffice? It would be simpler to say that deontic normativity is not a property of moral truths but of moral *judgments*, like Eriksson does.

The role of the alethic fault in practice is to put pressure on us to try to resolve our disagreements. The moral fault that we attribute to those who disagree with us does not work this way. If anything, it stimulates the entrenchment of our moral disagreements. It could be put as follows: whereas the lack of the alethic fault explains why we are generally rationally allowed to stick to our guns, the presence of the moral fault explains why we generally feel *obliged* to stick to our guns. In the case of aesthetics, for instance, we allow ourselves some leeway to move with the tide in order to fit in. However, in the case of morality, we think this would be a testament to a weak moral character and that doing so would thus be immoral (unless, of course, we are genuinely convinced by the opposing side's arguments).

6.4 Religious Disagreement

Another inquiry that seems deeply personal is spiritual or religious inquiry. I do not mean the academic anthropological inquiry into different aspects of religion, but rather the inquiry that aims to answer questions such as "Is there a God?" and "Does the universe have a purpose?" The religious doctrines that people endorse are often a significant part of their social identity. Phrases such as "I am a Muslim", or "As a Christian, ..." are very common. Religion expresses not just a theory about the universe but also a way of life. As such, it seems highly unlikely that people adopt conciliatory strategies concerning such beliefs. But should they?

Well, it certainly seems that religious or spiritual experiences are not shared, so insofar as there are genuine religious or spiritual experiences, it would be irrational to assume *sameness* for them. However, what should the atheist believe? There might be an interesting methodological asymmetry between the proponents and opponents of religion. Atheists like to say that they believe in science instead of religion and that science shows that there is no God. Yet, this only means that the thesis that a God or various Gods exist cannot be tested scientifically, which is something that theists and atheists alike could agree on: to exercise faith is to believe without evidence. If spiritual experiences are indeed personal, it is no wonder that the scientific method fails; science also fails to tell us if pistachio ice cream is tasty or if *Death of a Salesman* is a masterful play. What lies behind the atheist's conviction is the conviction that the scientific world order leaves no room for a spiritual world order. Sometimes this conviction is overtly normative: the scientific world order is simply *better*; it sets us free of those stuffy religious values, it is more egalitarian, and it allows us to take matters into our own hands instead of thinking of our fortunes as determined by some deity. Yet, these are personal convictions instead of testable claims about reality.

In short, it would not be rational to assume *sameness* in the case of religious or spiritual inquiry because neither religious beliefs nor the rejection of religion or spirituality is the result

of shareable, repeatable, or testable observations or computations. An account of how people acquire their religious beliefs would probably be much closer to an account of how they come to have their moral beliefs. Therefore, just like in the moral case, we could be rationally entitled to be steadfast in response to religious disagreement.

6.5 Philosophical Disagreement

I argue in a later chapter that not only ethics but also theoretical philosophy should be seen as a personal inquiry. This goes against methodological naturalism, the idea that philosophical inquiry should be modeled after scientific inquiry. However, before I make my case for rejecting naturalism and for philosophical inquiry as a personal inquiry, I need to answer some important questions about what personal inquiry amounts to.

It is tempting to adopt a kind of non-cognitivism about personal inquiry and treat it as dealing merely with opinions and sentiments. Yet, this approach works better for matters of basic taste than for philosophical inquiry. By endorsing non-cognitivism about philosophical inquiry, we would also be committed to endorsing a non-cognitivist attitude toward the very position we are trying to defend; if we claim that philosophy is just about expressing sentiments, then we must take our very claim to be an expression of a mere sentiment as well. This looks like a potentially unstable picture. Furthermore, philosophers strongly believe that their inquiry is guided by reason and aims for knowledge. If non-cognitivism holds true for philosophical inquiry, then the whole enterprise might strike us as rather pointless and mistaken. In the next chapters, I argue that judgments belonging to personal inquiry can still be true, guided by reasons, and that this inquiry can still aim for knowledge, even if it does not aim for truth.

7. Conclusion

In Chapter 2, I have made the distinction between shared and personal inquiry. Shared inquiry aims for truth, which means that it has an *a priori* commitment to taking disagreements to imply alethic faults. Personal inquiry does not aim for truth, which is why disagreements belonging to this inquiry do not possess this alethic fault (though they might possess other faults). The present chapter has further explored the practical significance of this divide. Shared inquiries are characterized by a shared concern for the truth and by a pressure to converge, which gives the rational impetus for a conciliatory response to disagreement. Conversely, personal inquiries lack the aim of truth and the pressure to converge, which is why we can be rationally entitled to stick to our guns in the face of disagreement.

Now that the practical contrast between shared and personal inquiry has been explored in sufficient detail, I shall move on, in the chapters that follow, to adding some theoretical layers to this distinction. The next chapter asks what truth must be for shared inquiry and how truth for personal inquiry might differ from this. The purpose of Chapter 5 is to show that personal inquiries are still robust inquiries that are guided by reason and that aim for knowledge. This allows for the development of a theoretical framework for philosophical inquiry that rivals the scientific framework.

4. Truth and Objectivity

1. Introduction

In previous chapters, I have argued that some of our inquiries are aimed at truth, while others are not. I have furthermore argued that those inquiries that do not aim for truth are nevertheless apt for truth. Truth still serves a valuable expressive role for such inquiries as a device for expressing agreement. A question that presents itself quite naturally is: Whatever must truth be so that it could serve these uses? And does this variety of uses lead to pluralism about truth itself? Yet, by asking such metaphysical questions, we enter treacherous lands for pragmatists. After all, the pragmatist seeks to expose the practical bearings of truth, not to explore its metaphysical underpinnings.

However, I argue in section 2 that a deflationary attitude toward the nature of truth won't suffice since it cannot account for the normativity of truth in the case of shared inquiry. Section 3 then contemplates what truth might be instead, and it tentatively suggests that truth is correspondence to reality. Section 4 discusses the possibility of an epistemic account of truth, and it shows that this won't do. Section 5 then puts forth a pragmatist-friendly interpretation of the idea that truth is correspondence, and section 6 extends this interpretation to the notion of objectivity and the doctrine of realism. Finally, section 7 develops a kind of moderate pluralism about truth *at the level of use*. I argue that the truth property could be correspondence across the board and that the concept of truth could be described simply as a device for expressing agreement, but that there are two distinct ways in which we can use this device and that by using it to aim for agreement, as we do in shared inquiry, we commit ourselves to realism.

2. Deflationism and the Normativity of Truth

There are different versions of deflationism, such as minimalism (Horwich 1998) or prosententialism (e.g., Grover, Camp & Belnap 1975, Brandom 1994). Such views often start with the observation that adding "It is true that …" to a sentence does not add to its content. Take, for instance, Frank Ramsey's example: "'It is true that Caesar was murdered' means no more than that Caesar was murdered." (Ramsey 1927, 38) Deflationists contend that there is nothing more to the nature of truth than what is expressed by the equivalence schema: ⌜ *p* ⌝ *is true* if and only if *p*. Truth does have an expressive function in our language, but this is merely to express agreement ("What she said is true") or to endorse multiple claims at once ("Everything Wittgenstein says in *PI* is true").

However, it is important to note that, though arguments have been made to the contrary (e.g., Wright 1992, Chapter 1), deflationism as such does not entail the rejection of the normativity of truth. At least, it is not obvious that it would. Before showing why this is so, I must describe what it is to take truth to be normative in the first place. The normative profile of truth contains four elements: (1) a deontic, (2) criterial, (3) teleological, and (4) an axiological element. (Ferrari 2021, 4) What it is to take truth to be normative is to believe that (1) truth is what ought to be believed, (2) a belief is correct if and only if it is true, (3) truth is the aim of inquiry and (4) truth is valuable. I have ordered these elements in the way I take it that they should be explained. Truth is what ought to be believed because a belief is correct iff it is true. A belief is correct iff it is true because truth is the aim of inquiry. Truth is the aim of inquiry because the pursuit of truth is valuable. To show why truth should be taken to be normative is to show why the pursuit of truth is valuable. Within the pragmatist framework, this means showing that the pursuit of truth is *practically* valuable. Some deflationists, like Rorty (1995), believe that this is not the case. Other deflationists, like Huw Price and Paul Horwich, argue that it is.

Horwich argues that deflationism about truth has the resources to account for the value of truth. He uses the following principle as an example: "It is desirable that: one believe the proposition *that e = mc²* just in case e = mc²." (Horwich 2006, 356)[4] This normative attitude cannot be generalized straightforwardly without it being grammatically awkward, but it can be generalized indirectly by making use of the equivalence schema. Given that *p* is equivalent to ⌜ *p* ⌝ is true, we can say the following:

> It is desirable that: one believe the proposition *that e = mc²* just in case the proposition *that e = mc²* is true

This can be generalized as follows:

> It is desirable that: one believe *x* just in case *x* is true

This is a statement of the value of truth, and as Horwich shows, it seems perfectly compatible with deflationism about truth.[5] The same move can be made with respect to the other forms of the normativity of truth. Take, for instance, the criterial norm, which I shall take to be the most neutral expression of the Truth Norm:

[4] While it might seem trivial that it is desirable that one believe only what is the case, it might be hard for pragmatists to effectively argue for this. Nevertheless, let us assume for the sake of Horwich's argument that this is the case.

[5] This claim could be challenged by arguing, as Michael Lynch (2004) does, that the generalized claim is the basic statement of the normative commitment and that the particular instances of this claim depend on and are motivated by this general statement, rather than vice versa. I do not want to take a stance on this issue, and given that I shall provide my own independent argument for why the truth norm is ultimately incompatible with deflationism, I shall for now interpret the deflationist's interpretation of the normativity of truth generously by assuming that Horwich's argument works.

(Truth Norm) For all *x*, it is correct to assert that *x* if and only if *x* is true

Deflationists need not reject this norm, but given that they believe that the nature of truth is captured in full by the equivalence schema, they are committed to the view that the Assertion Norm below expresses everything that is expressed by the Truth Norm:

(Assertion Norm) For all *x*, it is correct to assert that *x* if and only if *x*

However, in the previous chapters, I argued that an assertion like "Pistachio ice cream is tasty" is not governed by the truth norm. If the truth norm is a coordination device, then assertions that belong to personal inquiry should be exempt from it. Nevertheless, it is correct to assert that pistachio ice cream is tasty if and only if pistachio ice cream is tasty. This is platitudinous. The assertion that pistachio ice cream is tasty is a move in a language game. Whether it is correct depends on the standards of this language game, i.e., the norms that govern the assertion. These norms won't add up unless they corroborate the platitude that it is correct to assert that pistachio ice cream is tasty if and only if pistachio ice cream is tasty. Yet, the equivalence schema can be applied to this biconditional to get: it is correct to assert that pistachio ice cream is tasty if and only if ⌜ Pistachio ice cream is tasty ⌝ is true. The Assertion Norm applies to all assertions, also the ones that belong to personal inquiry, and therefore the Truth Norm, insofar as it is equivalent to the Assertion Norm, applies to all assertions as well.

One could, of course, try to deny that claims belonging to personal inquiry are assertoric. However, on what grounds would this denial be based? I have already argued in Chapter 2 that personal inquiry is apt for truth. What other criterion could there be that personal inquiry fails to meet? It would be inconsistent with the deflationist or pragmatist methodology to try to argue for a metaphysical requirement, such as factuality or representationality. In that case, one

would have to explain assertoric content in terms of its relation to worldly entities. This is exactly the kind of explanatory order that pragmatists eschew by adopting deflationism.

The most obvious way forward, and the one that Price endorses, is to say that all assertions aim for truth. Yet, I have argued that this view does not align with some of our assertoric practices. We ought to believe that certain forms of inquiry do not aim for truth. This means that the aim for truth must be restricted to shared inquiry and that, thus, the range of applicability of the Truth Norm must be restricted to shared inquiry. This could only be done by driving a wedge between the Assertion Norm and the Truth Norm. Deflationism simply does not have the tools for this kind of job. To do this, truth, for shared inquiry, must be interpreted as going beyond the equivalence schema.

3. What Could Truth Be?

As explained in Chapter 2, Price argues that the truth norm adds an *a priori* alethic fault to disagreements:

> [The truth norm] is a norm which speakers immediately assume to be breached by someone with whom they disagree, *independently of any diagnosis of the source of the disagreement.* Indeed, this is the very essence of the norm of truth, in my view. (Price 2003, 164)

Price's conception of the truth norm is curiously close to Crispin Wright's notion of Cognitive Command. According to Wright, a discourse displays the feature of Cognitive Command if and only if "it is a priori that differences in opinion […] will involve something which may properly be regarded as a cognitive shortcoming." (Wright 1992, 144) If a discourse displays Cognitive Command, we have an *a priori* commitment to taking disagreement within this discourse to be indicative of a fault.

Wright takes Cognitive Command to be a feature of *realism*. Realist domains of discourse are, according to Wright, domains of discourse that aim to correspond to reality. Wright takes this aim for correspondence to explain Cognitive Command: the fact that we aim to correspond to reality — that we attempt to accurately represent the world — explains why we take disagreement to be indicative of some kind of shortcoming. In the example mentioned above, it would be the idea that there is an objective fact about the tastiness of pistachio ice cream that our assertion "Pistachio ice cream is tasty" aims to represent, that would explain why we would take disagreement concerning this matter to imply an alethic fault. But, of course, we do not tend to think that there is such an objective fact about the tastiness of pistachio ice cream or that the purpose of asserting that pistachio ice cream is tasty is to represent reality. Therefore, in this case, it is the absence of this objectivity and of the aim to represent reality that would explain why our disagreement lacks the alethic fault. This alethic fault would thus be the fault of misrepresenting reality.

Matthew Shields agrees with Wright's analysis and argues that a commitment to an evidence-transcendent notion of truth follows directly from Price's description of the truth norm. Shields argues that the alethic fault in disagreement is reflected in the "assumption that there is something that outstrips any inquirer's language or thought to which inquirers are all accountable." (Shields 2022, 18)

However, I think Price would argue for the inverse of this explanatory relation: this intuition that the correctness of our assertions depends on how things stand independently of us is *explained by* the practice of taking disagreements to contain an alethic fault. This practice has more to do with the nature of inquiry than with the nature of truth. It is useful to assume, for certain inquiries, that we experience and process things in the same manner, and it is, therefore, useful to assume that, whenever we disagree, something must have gone wrong, even if we cannot locate this fault. These assumptions and their utility are what explains the truth

norm. The truth norm must not and cannot be explained by looking into the metaphysical nature of truth. Our aim to coordinate our beliefs is not explained by our supposed aim to correspond to reality. Rather, it is because we aim to coordinate our beliefs that we get this sense that our beliefs must respond to a standard that is external to us. If the lack of coordination is by itself a deficiency, then it is conceivable that one might have carried out one's inquiry perfectly but that one's beliefs might nevertheless still be deficient, solely because they are not shared by one's peers.

This does not take away from the fact that truth is normative in a way that goes beyond what a deflationary account can account for and that truth is an evidence-transcendent notion of correctness. Yet, this does not directly lead to correspondence. Sure, truth might transcend an individual's ability to detect it, but what about the community's abilities as a whole? Is it required that we believe that the truth could outstrip even our collected efforts? This shall be the topic of discussion in the next section.

Yet, before I move to this section, I must briefly consider the option of primitivism. Could we not maintain that truth is a substantive notion — given its normative role and the deflationist's inability to account for this normativity — but deny that it could be analyzed? It may be that once we have described the role of truth in our inquiry, we have gone as far as we can go.

Primitivism is a foundationalist theory of truth, which treats truth as a conceptual atom. According to this theory, inquiry is to be explained in terms of truth, and the role of truth in inquiry is fundamental, meaning that it cannot be explained, non-circularly, in terms of any other concept. (Asay 2013, 81) Yet these thoughts conflict with the theory proposed here. First of all, this primitivist theory would not hold for personal inquiry since, for this kind of inquiry, truth is merely a device for expressing agreement. It seems that primitivism about truth presupposes that the conceptual role of truth is uniform across all assertoric domains of

discourse. Given that it isn't uniform, and truth performs different roles for different forms of inquiry, it becomes very hard to maintain that "that's just what truth is like", so to speak, and that we shouldn't look to explain these differences.

Even if we were to maintain a kind of primitivism about truth only for the case of shared inquiry, this would be in conflict with the kind of theory I am hoping to develop. The difference between truth for personal inquiry and truth for shared inquiry is the normative character that it has in the latter case due to the fact that we aim for agreement. It is this normative role that we would be primitivists about. However, I have adhered to the Jamesian principle that we aim for truth only insofar as we aim to do what pays. Thus, the normativity of truth is not primitive; it is explained by our obligation to do what pays.

4. Friction as a Regulative Assumption

Price maintains that we need not look into the nature of truth to explain the role of truth. To a certain extent, I agree with him. The truth norm is merely an expression of something already there: the pressure to coordinate our opinions. Yet, as argued above, in failing to distinguish the truth norm from the assertion norm, deflationism lacks the resources for the truth norm to be an expression of the pressure to coordinate our opinions. For truth to be used as a norm that provides this kind of friction to our disagreements, it needs to be something more than what the deflationist could account for.

This does not yet show that truth has to be correspondence, though. Could it not be an epistemic notion, like superassertibility? The definition of superassertibility is as follows:

> A statement is superassertible […] if and only if it is, or can be, warranted and some warrant for it would survive arbitrarily close scrutiny of its pedigree and arbitrarily extensive increments to or other forms of improvement of our information. (Wright 1992, 48)

Wright's notion of superassertibility is similar to Hilary Putnam's conception of truth. (Putnam 1981, 49) Putnam makes a distinction between justification, which is potentially unstable, and truth, which is stable. He aims to bridge these two notions by defining truth as what one would be justified in believing under ideal epistemic circumstances. I prefer Wright's account of truth because it does not rely on an idealization of our epistemic abilities. On both accounts, however, truth just is a kind of sustainable justification, i.e., justification that would survive closer scrutiny. Superassertibility is also closely related to Peirce's conception of truth, as we shall see later on in this section.

The question is, however, what kind of sustainable justification we are talking about. There is a difference to be made between justification that is merely personal and a kind of *interpersonal* justification. This has everything to do with the difference between personal and shared inquiry. It seems that, *for me*, the claim that pistachio ice cream is tasty is superassertible: I am confident that I have examined the matter closely enough and that my current judgment will probably be final. However, if we interpret the notion of superassertibility as interpersonal, then I would not have much faith in my claim being superassertible because I do not believe that the assertion that pistachio ice cream is tasty — or, for that matter, the assertion that it is not — is such that *everyone* could endorse it as their final opinion. Nor is this what I am aiming for. This is why Wright, at one point, suggested that in the case of basic taste, superassertibility should be interpreted as *personal* superassertibility, which would render truth, once understood as consisting in superassertibility, relativistic. (Wright 2006)

For the aim for superassertibility to serve as a coordination device, it must be the aim for interpersonal superassertibility. This is still an epistemic notion of truth, though. It is

perfectly imaginable that there are multiple theories that are incompatible, but that could both be warrantably believed without fail by everyone, no matter how closely we scrutinize them. In that case, to say that only one of these theories is correct and the others are wrong would be to endorse an evidence-transcendent standard of correctness. Therefore, to stick with the superassertibility account of truth, we must either say that all these theories are interpersonally superassertible, and thus true, or that none of them are. In either case, we could not attribute an alethic fault to any disagreement that arises from the endorsement of such theories.

From a practical point of view, this does not seem to matter much. We attribute an alethic fault to our disagreements because doing so is useful, but it is only useful insofar as it helps us to gain beliefs that work better. If both theories are superassertible, then they are both maximally practically beneficial — since any practical failing would count as evidence against the theory. This means that the attribution of fault ceases to be useful in this scenario. Once we have reached the end of inquiry, and our beliefs work as well as they possibly could, we could treat disagreements as mere differences of preference, as faultless matters of taste.

This is essentially what Putnam argues for by saying that whenever two incompatible theories are equally adequate in every way, their incompatibility is merely apparent and that this appearance is caused by the endorsement of two different but equally adequate conceptual structures. (Putnam 1987) This is not to say that every proponent of an epistemic notion of truth has to accept Putnam's conceptual relativity. All they need to do is find some way of rejecting evidence-transcendent alethic faults. Putnam sees this as the rejection of, what he calls, metaphysical realism, which is, roughly, the thesis that an inquirer who has carried out their inquiry impeccably and has taken all possible evidence into account, could still get things wrong. (Putnam 1981)

Instead of saying that disagreement *always* implies a fault, we can say that it implies a fault unless inquiry has been carried out impeccably. Thus, the proposal is that if two inquirers

carry out their inquiry impeccably, taking all the relevant evidence into consideration, and ending up with optimally beneficial beliefs that are nevertheless (interpersonally) incompatible, then their disagreement is faultless. In this case, the alethic fault would be a kind of stand-in for a yet-to-be-discovered procedural fault, e.g., the fault of having based one's belief on inconclusive evidence. The alethic fault would then only apply to ordinary cases of disagreement and not to those rare cases in which the opposing views are both optimally conducive to one's practices. As argued, this fits well into our pragmatist framework: coordination isn't a goal *an sich*, the aim for coordination is in service of the aim to find beliefs that work best. If the beliefs involved in the disagreement are both maximally beneficial to our practices, then it no longer matters which of them we choose to endorse, and it no longer matters that we side with the same beliefs, so the disagreement becomes faultless, since it would merely be a matter of personal preference.

Sadly, though, the rule "conciliate unless you've reached the end of inquiry" is not a rule we could meaningfully follow since we can't possibly know when the end of inquiry has been reached; we cannot possibly know if there aren't further beliefs out there that might serve us even better, or if our current beliefs that work so well for us today might fail to work for us tomorrow. There might be many instances in which we find that our peers endorse theories that are incompatible with our own theories, but in which case, all theories on offer seem internally coherent. In such instances, the temptation might arise to think that the end of our inquiry has been reached, thereby treating the disagreement as faultless and sticking with one's preferred system. However, this kind of conduct would rarely be useful. We normally think of such conduct as evincing the vice of intellectual immodesty. It would be more virtuous to believe that there will always be things left to discover and that there will always be more that is unknown to us than known. It is generally better to inquire too far than not far enough.

Therefore, the rule couldn't be "conciliate unless you've reached the end of inquiry", and the alethic fault must indeed be *a priori* and possibly evidence transcendent.

There is an alternative, which is to say that there couldn't be any disagreements at the end of inquiry. My argument for evidence transcendence presupposes that it is possible that two inquirers who carry out their inquiry impeccably end up with incompatible beliefs. This premise could be rejected by assuming that inquiry is such that all disagreements will be resolved as we go along. What would substantiate this assumption?

Peirce argues that it is a regulative assumption of *scientific* inquiry (as opposed to inquiry into matters of taste) that it will converge toward a single final opinion and that, thus, all disagreements will ultimately be resolved one way or another. (Misak 2004, 126) According to Peirce, the purpose of inquiry is the cessation of doubt, and persistent disagreement gets in the way of this purpose. This is why it would only be worth it to engage in inquiry if we assume that all disagreements will ultimately be resolved and that we will slowly converge toward a single final opinion.

However, this argument would only work once we assume that disagreement bothers us, i.e., that it should lead to a suspension of judgment. This would not be the case for our opinions on ice cream, so why would scientific inquiry be any different? I have argued that the difference lies in the former being guided by the truth norm, which the latter lacks. However, Diana Heney shows in her article "Reality as Necessary Friction" that, for Peirce, it is not the concept of truth but the concept of reality that does the heavy lifting. (Heney 2015)

Disagreement need not cast doubt on my opinion unless I would believe that our opinions cannot both be correct. If we take the norm of correctness to be epistemic, then this means that I must be convinced that at most one of our beliefs can be justifiable in the long run. Why would I believe this? Why must I preclude the possibility that both remain equally

justifiable? Well, if Peirce is right, then this belief is warranted by a conception of reality that shows that we are destined to converge toward a single opinion. He says:

> [A]ll the followers of science are fully persuaded that the processes of investigation, if only pushed far enough, will give *one* certain solution to every question to which they can be applied […] Different minds may set out with the most antagonistic views, but the progress of investigation carries them by *a force outside of themselves* to one and the same conclusion. This activity of thought by which we are carried, not where we wish, but to a *foreordained goal*, is like the operation of *destiny*. No modification of the point of view taken, no selection of other facts for study, no natural bent of mind even, can enable a man to escape the *predestinate opinion*. This great law is embodied in the conception of *truth* and *reality*. The opinion which is *fated* to be ultimately agreed to by all who investigate, is what we mean by the truth, and the object represented in this opinion is the real. That is the way I would explain reality. (Peirce 1878, 299-300 *emphasis added*)

This quote shows that Peirce's regulative assumption that disagreements will be resolved is a regulative assumption about reality: we must assume that reality is such that it carries our inquiry "to a foreordained goal" and to "one and the same conclusion".

In short, according to Peirce, reality must be such that whichever opinion will be believed by all in the long run (at the hypothetical end of inquiry) will correspond to it. It is not because an opinion is indefeasible that we think it is the one true opinion to be had, but because we assume that our indefeasibly justified opinions *correspond to reality*. In a roundabout way, then, we have ended up with a version of the correspondence theory. In fact, there are some contemporary pragmatists, such as Robert Lane (2018) and Andrew Howat (2020), who argue that Peirce can be read as a correspondence theorist.

However, this kind of correspondence theory would not take truth to be evidence transcendent. Therefore, one could perhaps argue that this epistemically constrained version of correspondence "isn't the real thing". In any case, I think this epistemic version of the correspondence theory is decidedly worse than a conception of correspondence that takes it to be an evidence-transcendent notion of correctness. In trying to keep truth from being evidence transcendent, we would have to make the very controversial move of denying the existence of lost facts; we would have to commit ourselves to the assumption that inquiry could ultimately resolve disagreements about the most minute facts of the distant past, even when it seems intuitively possible that these facts are no longer discoverable.

Perhaps it would be useful for some of our inquiries to believe that reality will convey itself to us and lead us to some final conclusion. Yet, to think this of all shared inquiries is to risk going off on wild goose chases, trying to solve puzzles that have already been shown to be unsolvable. This conception of reality either shrinks the world down to a size small enough for it to be consumed wholly by the human intellect, or it blows up this intellect to astronomical proportions just to ensure that no stone will be left unturned. I think neither of these moves is particularly helpful to our practices.

We could also opt for the exact opposite of this picture. We could think of the attribution of an alethic fault to disagreements not as a commitment to a yet-to-be-discovered procedural fault but as an act of intellectual humility, an expression of the conviction that no matter how much we know there will always be much more we can never know. Given this ignorance, your guess is as good as mine, so to speak, and holding on to our own beliefs in the face of disagreement with our peers is mere stubbornness, a sin of intellectual immodesty.

Both strategies express the conviction that disagreements imply a fault. Yet, the Peircean strategy takes disagreement as a sign that we haven't yet reached the end of inquiry and that we must dig deeper, whereas this strategy takes disagreement as a sign that the correct

answer might be out of reach altogether. I think that it has fewer drawbacks. On a theoretical level, it avoids the controversial move of denying the existence of lost facts, while on the practical level, it safeguards against unproductive inquiries into intractable issues.

On both strategies, what it is for a belief to be true is for it to correspond to reality in some sense. On the Peircean strategy, it is the concept of reality that does the heavy lifting, whereas on my preferred strategy, it is the concept of correspondence itself, interpreted as an evidence-transcendent notion of correctness, that tells us that at most one of us can be right and that thereby adds the alethic fault needed for truth to be used as a coordination device.

I understand that pragmatists might instinctively resist an evidence-transcendent notion of correctness at all costs because they have been told time and again that such a notion of correctness has no bearing on our practices. However, I have now shown that it does have such bearings. After all, aiming for an evidence-transcendent notion of correctness does alter our practice; doing so adds friction to our disagreements, and it thereby makes us more inclined to coordinate our opinions.

5. A Pragmatist Interpretation of Correspondence

I have argued in the previous section that it would not be advantageous to endorse an epistemic account of truth, which analyzes truth in terms of a kind of indefeasible justification. First of all, this analysis isn't particularly informative on its own since we aim for indefeasible justification also in the case of personal inquiry. What it is to aim for truth is to aim not just for indefeasible justification but also for the coordination of our beliefs. We must thus aim for beliefs that are indefeasibly warranted for everyone. This could be built into the epistemic account of truth by taking truth to consist in warrant that is not just indefeasible but also interpersonal. Yet, it is conceptually possible that there are multiple beliefs of this status that

are incompatible. Given our inability to recognize such a situation — that is, our inability to reliably tell when our beliefs are indefeasibly warranted as opposed to merely warranted here and now — we can only follow the general rule that an alethic fault be attributed to *all* disagreements. This makes the alethic fault an *a priori* fault that is possibly evidence transcendent. We could also assume that it is simply not possible to have two incompatible beliefs that are both interpersonally indefeasible and that, instead, our inquiry will lead us to a unique set of beliefs. I have shown that Peirce believes this and that he takes this to be a regulative assumption about reality. Yet, I have argued that this approach has some drawbacks and that it would be better to think of truth as evidence transcendent.

The pragmatist might raise their familiar worry: how could we aim for a truth that possibly outstrips our collective ability to detect it? Well, I have already shown how this is to be done: by adding an alethic fault to all disagreements. Granted, this is still a bit vague, but in the previous chapter I have explained that this amounts to the rule that one ought to conciliate in the face of peer disagreement, irrespective of the details of the disagreement. Therefore, my conception of the aim for truth, though it may commit one to an evidence-transcendent norm of correctness, comes with a fairly detailed set of practical instructions as to how to implement it.

As mentioned above, this interpretation of the aim for truth also commits us to metaphysical realism. In some sense, then, Wright (1992) is correct to link this *a priori* alethic fault to realism. Yet, what are we to make of his talk about correspondence? Wright believes that the idea that we aim to correspond to reality explains why we add this *a priori* fault to our disagreements. We must believe that whenever we disagree, at least one of us must be mistaken because only one of our beliefs can correspond to the way things really are.

As argued above, the pragmatist must think that Wright gets things the wrong way around. Luckily though, we could inverse the explanatory relation. I have shown in the previous

section that Peirce introduces his metaphysical beliefs not as an explanation of why inquiry must lead to convergence but as a *regulative assumption*: in order to engage in our inquiry in the way we think we should, we must believe that reality is such that it will guide us to a "foreordained goal" and a "predestinate opinion". (Peirce 1878, 300) The metaphysical realism implied by the aim for truth could be interpreted in the same fashion: to engage in our inquiry in the way we think we should, we must take it that we aim for our beliefs to correspond to reality.

This move is also inspired by Price, who, in *Facts and the Function of Truth* (1988), takes the notion of 'factuality' to be explained by, and thus downstream from, our commitment to taking disagreements to imply a fault. Likewise, we can think of our aim for correspondence as explained by this practical commitment: what it is to take the correctness of a belief to outrun our, or everyone else's, ability to judge its correctness, is for us to be intolerant to disagreement, irrespective of its particulars.

It might almost sound like correspondence has become nothing more than a useful fiction: for us to carry out our inquiry in the way in which we think we ought to, we should just *pretend* that our assertions aim to correspond to reality. However, as I have shown in "Alethic Pluralism for Pragmatists" (Kaspers 2022, 10), pragmatists cannot actually endorse this kind of response when it comes to truth.[6]

There are two reasons for calling a belief a useful fiction: 1) it is useful to act on the belief, but it is actually nonsense, or 2) it is useful to act on the belief even though it is actually false. Pragmatists would want to go for the first option by arguing that the correspondence theory is "empty metaphysics". Their argument for this view is that aiming for correspondence, as opposed to justification or an epistemic notion of truth — i.e., having an external instead of

[6] What follows is a quick summary of the argument given in this paper. Please see my paper for the extended version of the argument.

an internal standard — has no bearing on our practices. However, I have shown in this chapter that this is wrong and that the aim for correspondence *is* practically significant. Therefore, the pragmatist argument to the effect that the correspondence conception of truth is nonsensical or empty does not work.

The only option left is to say that, though it would be useful to pretend that truth is correspondence, this is not actually the case. Yet, this response would be self-undermining. To say that the idea that truth is correspondence helps us to properly conduct our shared inquiry is to concede that this idea works great for our shared practices. We are justified in believing that truth is correspondence for shared inquiry because it aligns with our intuitions about how shared inquiry should be conducted, and this belief won't let us down because this way of conducting shared inquiry is practically beneficial. In what way would this belief be false, then? In order to insist that it is false, we would have to rely on a notion of truth that goes beyond our practices, that goes beyond our internal standards. This means that we would end up presupposing what we are trying to reject: an external standard of correctness.

Therefore, pragmatists cannot treat correspondence as a useful fiction, and nor should they. As long as they interpret the aim for correspondence as the practical commitment to taking disagreements to contain *a priori* alethic faults, and as long as they maintain that its metaphysical implications are *explained by* this practical commitment, the idea that truth is correspondence, even in the inflated sense, has clear practical bearings, and it thereby satisfies the Pragmatic Maxim.

6. Objectivity

In Chapter 2, I mentioned two examples: the first example was the belief that it pays to bet on the third leftmost horse, and the second example was the belief that the food at the Yellow

Pigtail is excellent. These examples were meant to show that in some cases, it makes sense to pool our resources, whereas in other cases, it doesn't. One should cease to maintain the former belief if one finds out that this belief does not work for others, whereas one is entitled to stick with the latter belief even if one finds out that the belief does not work for others. Intuitively, this difference is explained by the subject matter: matters of gustatory taste are *subjective*, whereas there is an *objective* fact as to which betting strategy is most successful.

However, I have argued that, as pragmatists, we should take this practical difference to be primary. Sometimes it is practically beneficial to assume sameness, while other times, it is not. Sometimes it is useful to assume that the beliefs that work best for me are those that would work best for everyone else as well, while other times, it is more useful to assume that different beliefs work best for different people.

This does not mean, however, that this distinction between objectivity and subjectivity couldn't follow from this practical distinction. We could say that we treat a subject matter as objective if and only if everyone ought to have the same beliefs concerning this subject matter. It follows from this that we treat a subject matter as objective whenever we commit to the idea that one is only allowed to believe *p* if one also believes that everyone else ought to believe *p* as well. In the case of personal inquiry, we may believe that our doxastic commitments apply only to ourselves, but in the case of shared inquiry, to be doxastically committed to a proposition is to take everyone else to share in this commitment.

This might seem too strong. What about those who only tentatively believe a proposition? What about those who are aware of their lack of expertise in the subject matter and who might think that, though their current epistemic context warrants their belief, this context might be flawed? What about those who are mindful of the possibility that others with superior epistemic contexts would not be committed to the same belief? I think this is not a problem for my view. First of all, those who are aware of their lack of expertise might have a

doxastic attitude that stops short of belief. In one possible situation, they might know that the experts do not believe the proposition, so they might reason in the following way: if *p* is true, then everyone ought to believe *p*, but nobody does, so *p* is probably not true. In another possible situation, they might not know what the experts think, and they might only know that their limited epistemic context warrants their belief in *p*. In that case, they might believe *p*, thereby believing that everyone else ought to believe *p* too, but they would cease to believe *p* after finding out that this doxastic commitment is not shared by others.

There might be problems going in the other direction as well. Suppose that one takes their epistemic context to be *better* than most others. In that case, they might believe *p* while realizing that others lack warrant for this belief due to their more limited epistemic context. Would this be an instance of believing *p* without believing that everyone else ought to believe *p* as well?

I do not think that it is. In some sense, we might excuse our fellow inquirer from not having the right doxastic commitments, but we do not absolve them completely. This is evidenced by the fact that we take their epistemic context to not merely be *different* from ours (like we might in the case of personal inquiry), but we take it to be *worse* than ours. Even if those in worse epistemic contexts are warranted in believing what they do, they are still at fault simply in virtue of their commitment to false beliefs and their failure to commit to true beliefs. This is exactly the role of the truth norm; it adds this kind of alethic fault to our disagreements.

I have argued that we treat a proposition as objective insofar as we believe that our doxastic attitude toward this proposition ought to be shared by everyone. This simply is to believe that there is ultimately only one correct doxastic attitude to be had. Yet, the ordinary non-pragmatist explanatory order goes as follows: the belief that the subject matter is objective explains this belief that there is only one correct doxastic attitude to be had, which in turn explains the belief that our doxastic attitudes toward the subject matter ought to be shared by

everyone. I take it that the pragmatist explanatory order goes in the opposite direction. The assumption of sameness is useful and thus rational for propositions belonging to certain subject matters, which is why we should believe that our doxastic attitude toward these propositions ought to be shared, and why we thus ought to believe that there is only one correct doxastic attitude toward such propositions to be had, which is *ipso facto* why we ought to believe that this doxastic attitude is *objectively true*.

This requirement of objectivity stems not from ontological but from practical considerations. I take it that it is in this practical sense that shared inquiry could be said to be *realist*. Its realism does not entail metaphysical presumptions about the existence or nature of the entities it refers to, but it does entail the presumption that one's beliefs are *objectively* true. This kind of realism is in the spirit of Wright's *Truth and Objectivity* (1992) in the sense that Wright's account of realism also has less to do with the mode of existence of the subject matter and more with the role that the subject matter plays in our inquiry. The matter of objectivity is to be separated from that of mind-independent existence.

7. Alethic Pluralism

I have argued in this chapter that truth for shared inquiry is to be interpreted as correspondence. I have also argued in previous chapters that truth for personal inquiry is merely a device for expressing agreement. Where does this leave us? *Prima facie*, this seems to commit us to a kind of pluralist theory of truth, namely a dualism about truth. In that case, truth would be deflationary for personal inquiry and correspondence for shared inquiry. Yet, what would bind these two notions of truth? In what way do they denote the same concept? And is this dualist view even stable? I hope to answer some of these questions in the present section.

7.1 Deflationism and Correspondence

The basic idea of correspondence, that a proposition is true if and only if it corresponds to the way things are, does not need to be read as implying a commitment to objectivity or realism. It could be seen as a mere platitude. This is what deflationists usually argue. For instance, Paul Horwich believes that it is a strength of deflationism, when compared to epistemic or use-based theories of truth, that it can accept the correspondence platitude. Furthermore, he argues that though the general idea that truth is correspondence seems to be intuitive, particular correspondence *theories* have never gained much traction:

> The common-sense notion that truth is a kind of 'correspondence with the facts' has never been worked out to anyone's satisfaction. Even its advocates would concede that it remains little more than a vague, guiding intuition. But the traditional alternatives — equations of truth with 'membership in a coherent system of beliefs', or 'what would be verified in ideal conditions', or 'suitability as a basis for action' — have always looked unlikely to work, precisely because they don't accommodate the 'correspondence' intuition[.] (Horwich 1998, 1)

Indeed, I also have said nothing about the metaphysics of the correspondence relation. For instance, I have not endorsed isomorphic or causality-based accounts of this relation. In some sense, then, I also take the idea that truth is correspondence to be nothing more than a "vague, guiding intuition." (Horwich 1998, 1)

At the same time, though, I have shown that the aim for truth commits to metaphysical realism, which flies in the face of deflationism. Yet, this does not have to do with the correspondence relation itself, but it is a regulative assumption that we make to facilitate convergence. It is thus the *aim* for truth, not the nature of truth itself, that entails this commitment to realism.

If we don't aim for truth, and we only employ truth as a linguistic device for expressing agreement, then truth is not a potentially evidence-transcendent standard, only because it wouldn't really be a standard at all. If we accept the laws of classical logic, then we may say that all propositions are either true or false, even if it is not within our capacity to decide for all propositions whether they are true or false. Again, if truth is not a standard at all, then this does not commit us to an evidence-transcendent standard. We wouldn't say that, if two inquirers have carried out their inquiry impeccably and end up disagreeing, one of them must still be *at fault*. Yet, once we *aim* for truth, thereby aiming for agreement, truth is used as a standard, and we do incur this possibly evidence-transcendent alethic fault. It could be argued, however, that this is simply the result of (1) us aiming for truth and (2) the acceptance of classical logic. It need not be substantialized by exploring the metaphysics of the correspondence relation.

7.2 Truth as One or as Many?

Where does this leave us with respect to the question of whether truth is plural? I have shown above that there is little need for two different truth properties. Truth could be seen as correspondence across the board. For personal inquiry, this would be correspondence in a fully deflated sense, normatively as well as metaphysically. For shared inquiry, truth is both normative and metaphysically committing, but it is neither required nor, I think, conducive to inquire into the metaphysics of the correspondence relation to explore these commitments. These commitments are a consequence of using truth to aim for convergence. They are thus a consequence not of endorsing a different kind of truth but of using the same kind of truth for a different purpose.

What about the concept of truth? The pragmatist argues that to analyze the concept of truth, we must look at its role in practice. I have shown that truth serves different roles in

practice. For personal inquiry, we use truth to mitigate our (or another's) justificatory burdens. As Lionel Shapiro (2021) argues, we can use truth to express our agreement with another's claim or claims, and thereby we could either defer to them or, in another dialectical context, allow them to defer to us. In shared inquiry, agreement is what we aim for. Truth can still play the role described by Shapiro in the case of shared inquiry, but in addition to this, we also use truth as something to aim for, which turns it into a genuine standard.

However, as Amy Thomasson notes, there is a difference to be made between the function of a concept and its uses. (Thomasson 2020, 53) This distinction is subtle. For instance, a hammer can serve many purposes, such as cracking nuts, but its function is to drive nails. In the case of a hammer, we may say that it was designed to drive nails, not to crack nuts. This wouldn't work in the case of concepts unless we mean it in some convoluted sense as "designed by the forces of evolution". However, while I shall not make a principled distinction between functions and uses, I do believe that there is an intuitive distinction to be made in the case of truth. The function of "is true" has to do with the role it plays in our linguistic system. We could say, like P. F. Strawson does, that the primary function of truth is to express agreement; truth just is a linguistic device for expressing agreement. (Strawson 1949, 94) Yet, just like we can use a hammer in different ways, we can use this linguistic tool in different ways as well. In the case of personal inquiry, we merely employ it to take on someone's justificatory burdens or, in a different dialectical context, to transfer our burdens onto someone else. This use can easily be made sense of while also thinking of truth as essentially a linguistic device for expressing agreement. In the case of shared inquiry, we aim for agreement, and we thus use truth as a standard, but as argued above, this can also be made sense of while thinking of truth as essentially a linguistic device for expressing agreement. If a truth attribution is an expression of agreement, and we aim to agree, then we could be said to aim for truth. We may pick up some collateral commitments by using truth in these different ways, such as the commitment

to realism, but we could ultimately think of it as having the same function across the board and, therefore, as the same concept across the board.

Does this count as a pluralist theory of truth? It could do since many pluralist theories are "moderate" instead of "strong", which means that truth is unified at some level. (Pedersen 2012) Moderate pluralists usually take the *concept* of truth to be unified and believe that there are multiple properties that instantiate the concept for different domains of discourse, such as correspondence, superassertibility, coherence, etc. (e.g., Wright 1992 and Lynch 2009)

Michael Lynch gives a functional account of the concept of truth, but he doesn't distinguish between function and use in the way I do above. He argues that the function of truth can be analyzed in terms of a network of platitudes, one of which is the platitude that truth is the goal of inquiry. (Lynch 2009, 12) Lynch seems to presuppose that truth always serves the same use, namely as a norm of belief and as the goal of inquiry. He then argues that there are multiple properties that play this role for different domains of discourse. In some sense, my account is a reversal of Lynch's account. Lynch's pluralism is formulated at the metaphysical level, whereas mine is formulated *at the level of use.*

7.3 Meta-Alethic Pragmatism

Given this pluralism at the level of use, we could ask which use truth has when it comes to discourse *about truth*. This presents a puzzle for my view. If we do not aim for truth in our discourse about truth, then we are not realists about truth, and then we have to maintain that there are no objective truths about truth itself. How could this be?

On the other hand, if we do aim for truth in alethic discourse, then it seems that we must say that *all* truths are objective. This is because the statement "⌜ *p* ⌝ *is true*" is a statement about truth, regardless of what *p* is about, and therefore, if we disagree about this claim, and

we take our discourse about truth to be aimed at truth, then we must presuppose that our disagreement contains an alethic fault. But then we must also presuppose this same alethic fault for our disagreement about *p* itself, regardless of what it is about, given that this proposition is equivalent to the truth attribution, so we must presuppose that all disagreements contain alethic faults.

I believe that this first problem is merely apparent. If we reject the truth norm for alethic discourse, then we aren't saying that *all* disagreements about truth are faultless. We are only saying that we shouldn't presuppose that all our disagreements about truth contain alethic faults. The second problem makes it clear why we shouldn't have this presupposition: if we presuppose that all disagreements about truth contain alethic faults, this commitment would spread to all disagreements whatsoever, which would turn every inquiry into a shared inquiry. Insofar as a proposition *p* belongs to a shared inquiry, we must presuppose that disagreement about *p* contains an alethic fault. Given that *p* is equivalent to ⌜ *p* ⌝ *is true*, this alethic fault carries over. This is not only the case for alethic faults but it applies to other faults as well. For example, if *p* belongs to the moral domain, then disagreement about *p* implies a moral fault, and this fault also carries over to our disagreement about ⌜ *p* ⌝ *is true*; if I think you are immoral for believing *p*, I also take you to be immoral for believing ⌜ *p* ⌝ *is true*. Yet, moral inquiry does not aim for truth, so in that case, disagreement about *p* does not contain an alethic fault, and neither should the disagreement about ⌜ *p* ⌝ *is true*.

If alethic discourse is not guided by the truth norm, it isn't realist, which means we would be meta-alethic anti-realists. I have tried to defend a kind of meta-alethic anti-realism, specifically a meta-alethic *pragmatism*, in my article "Alethic Pluralism for Pragmatists". (Kaspers 2022) In that article, I argued for a pragmatist notion of truth for alethic discourse instead of a deflationary notion. However, as Chase Wrenn shows, this option does not work.

Wrenn starts his argument with an example of a proposition that he assumes, for the sake of the argument, to be true but unknowable:

Expand The universe expands and contracts eternally, without beginning or end.

Consider this truth-attribution:

T-Expand Expand is true.

Because **T-Expand** entails **Expand**, any warrant for **T-Expand** would supply warrant for **Expand**. If you could be indefeasibly warranted in believing or asserting **T-Expand**, then you could be indefeasibly warranted in believing or asserting **Expand**. If **T-Expand** were knowable, **Expand** would be too. So, since **Expand** is *unknowable*, **T-Expand** must be unknowable too.

Assume for *reductio* that aletheiological truth is epistemic: <<p> is true> is true only if <<p> is true> is knowable. Since **T-Expand** is unknowable, it follows that **T-Expand** is not true. (That doesn't mean **T-Expand** is false; epistemic truth may tolerate truth-value gaps.) By hypothesis, **Expand** is unknowably true, so:

Expand is true, but **T-Expand** is not true.

Substitution yields:

(*) **Expand** is true, but <**Expand** is true> is not true.

which has the form:

p, but <*p*> is not true.

> This is inconsistent with ES [the equivalence schema] and the non-negotiable logical principle of Capture: $p \vdash \langle p \rangle$ is true. So, by *reductio*, if truth is non-epistemic in any discourse, it must be non-epistemic for truth-attributions. (Wrenn 2020, 313)

I think this argument is correct and that we must, therefore, accept its conclusion: truth for truth-attributions cannot be epistemic (at least not for all truth-attributions). I thereby reject that the realism/anti-realism distinction turns on the issue of mind-(in)dependence. The distinction I have introduced above revolves around objectivity and the aim for truth. I have argued that we do not aim for truth in the case of, for instance, matters of basic taste because it is not practical to aim for convergence for such matters. This is what drives the sense that such matters are *subjective* since we deem it possible that two people have contradictory beliefs, yet neither is (alethically) at fault. On the other hand, it is the presupposition that we experience and process things in the same way that, when rational, justifies the aim for convergence, and thereby the *a priori* attitude that, whenever people disagree, at least one of them must be at fault. It is, in turn, this attitude that drives the sense that the subject matter is *objective*. This is not the same as the mind-dependence/mind-independence distinction, but it is nevertheless a distinction in the vicinity, and it could serve as the grounds for some kind of realism/anti-realism distinction.

I think it does a better job of explaining the paradigm cases than the mind-(in)dependence distinction. Discourse on basic taste is one such paradigm cases since it seems very far removed from realist discourse. It is clear that matters of taste are subjective in the way described above. For instance, I could believe that pistachio ice cream is tasty, and you could believe that it is not, and neither of us need be at fault. However, it is far less clear that such matters cannot be mind-independent. It might be unknowable, for instance, whether the meat of the dodo is tasty. If we accept an epistemic notion of truth, we must also accept that it

is neither true nor false that the meat of the dodo is tasty. This has ramifications for our logical system since it entails the rejection of the bifurcation thesis and, thus, of classical logic. Yet, we do not need to take this theoretical baggage on board if we accept a deflationary notion of truth.[7]

8. Conclusion

I have argued in this chapter that deflationism about truth is incompatible with the distinction between shared and personal inquiry. Furthermore, if the aim for truth is to be used as a coordination device, it must commit to an evidence-transcendent notion of correctness. This means that it must commit to realism. However, I have argued for a pragmatist interpretation of this metaphysical commitment that takes it to be explanatorily downstream from our practical commitment to the aim for truth. I have proceeded to argue that this metaphysical commitment has less to do with the property of truth *per se* and more with how we use truth, i.e., whether we aim for truth. It is, therefore, possible to think of truth as consisting in correspondence for all domains of discourse, as long as we keep in mind that, in the case of shared but not personal inquiry, this property is used as an evidence-transcendent standard (and is hence incompatible with a deflationist interpretation of correspondence). The concept of truth could also be interpreted as uniform as long as we distinguish between the function of truth and its uses. Truth could be seen as simply a linguistic device for expressing agreement. Yet, it is the different ways in which we *use* this device that explains why shared inquiry is realist, whereas personal inquiry is not. Therefore, my account of truth is plural, but only *at the level of use.*

[7] This is not to say that there may not be other reasons for rejecting classical logic.

5. Rationality and Knowledge

1. Introduction

The previous chapters have introduced and developed a distinction between personal and shared inquiry. Shared inquiry aims for truth, whereas personal inquiry does not. It is rational to conciliate in the face of peer disagreement in the case of shared inquiry, while we could be rationally entitled to remain steadfast in the case of personal inquiry. Shared inquiry is objective and realist, whereas personal inquiry is subjective and anti-realist.

It may seem in line with the abovementioned features of the distinction to say that shared inquiry is a rational inquiry that aims to accumulate knowledge, whereas personal inquiry has to do merely with sentiments. However, this would be a mistake. In this chapter, I argue that personal inquiry is still robustly rational and that it aims for knowledge, even though it doesn't aim for truth. Nevertheless, there are still distinctions to be drawn when it comes to rationality and knowledge. In section 2, I argue for the distinction between personal and shared rationality. In section 3, I introduce a similar distinction with respect to knowledge. In section 4, I show how these different senses of knowledge explain the difference in the mechanics of testimony between personal and shared inquiry.

2. Rationality

Would it not be best to say that shared inquiry deals in beliefs, whereas personal inquiry deals in sentiments? Beliefs are shareable mental states, whereas sentiments, though they could be shared, are also deeply personal. Sentiments need not be irrational, but they are often *arational*. While this contrast between beliefs and sentiments may be intuitive, it is not so easily made. My finding Van Gogh's paintings beautiful is certainly a sentiment, but may I not also believe

that Van Gogh's paintings are beautiful? Is this something I could believe, or could I only opine it? Relatedly, we could ask whether we could *know* that Van Gogh's paintings are beautiful. I will answer this question in the next section. For now, I am primarily concerned with the question of whether personal inquiry deals in beliefs.

Robert Brandom argues that there are two ways of understanding what a belief is: either it is the attitude of taking the matter at hand to be true or it is something that can "serve as and stand in need of reasons". (Brandom 1994, 5) As I argued in Chapter 2, personal inquiry is apt for truth, so if believing is the attitude of taking to be true, then there is no reason why personal inquiry couldn't accommodate beliefs. However, given that the distinctive feature of personal inquiries is that they do not aim for truth, and given that their truth is deflated, it seems that this definition is bound to be uninformative.

The second definition is more promising. If personal inquiry is guided by reasons, then it surely deals in cognitive contents and thus in beliefs. I have argued in Chapter 2 that disagreement is a useful activity even if we do not try to coordinate our opinions. By engaging in the activity of giving and asking for reasons, we can make sense of both our own opinions and those of our interlocutor, which enables us to keep each other in check by making our actions more predictable and by allowing us to hold each other accountable. This story presupposes that personal inquiry is still a rational inquiry, unlike what a noncognitivist might think; our assertions are not mere expressions of sentiment, but they are moves in the game of giving and asking for reasons. However, more needs to be said about what reasons are to see how personal inquiry can still be guided by reasons, though it is not guided by a truth norm.

Ferrari and Pedersen argue that there are two types of rationality, which they label external and internal rationality. (Ferrari & Pedersen 2019) According to the externalist view, a reason is a good one if it is "the result of properly functioning cognitive capacities or reliable belief-forming processes." (Ferrari & Pedersen 2019, 13) In order to understand which

functioning of cognitive capacities is "proper" and which belief-forming processes are "reliable", we need to think of them as serving a goal outside of merely being rational. If it is rationality itself that we are after, then the "proper" functioning of cognitive capacities would be one that adheres to reason, and our characterization of reasoning as being the result of proper functioning cognitive capacities would be plainly circular. This is why this account is external: what it is for our cognitive capacities to function properly is for them to be in touch with reality in the right kind of way, and what it is to be in touch with reality in the right kind of way is to be in touch with reality in such a way that our cognitive capacities produce truths about reality. On the externalist view, rationality is a tool for getting to the truth. This makes it incompatible with personal inquiry. If personal inquiry is to be guided by reasons but not by truth, rationality must be defined in different terms.

The internalist view differs from the externalist view in that the standards of reason are internally accessible. How are we to make sense of this? What serves as the standard of correctness for reasoning if not truth? Ferrari and Pedersen give being "a member of a coherent belief system" as an example of an internal standard for rationality. (Ferrari & Pedersen 2019, 13) This would not be a very precise criterion for rationality unless we can say something more about what coherence consists in. Yet, as it turns out, the more we say about coherence, the more likely it is to refer to truth. Take, for instance, C. I. Lewis's influential analysis of coherence. Lewis argues that proposition *P* coheres with a set of propositions if and only if the assumption that this set of propositions is true increases the probability that *P* is true as well. (Lewis 1946) If P1 is a good reason for P2 because P2 coheres with P1, then, given Lewis's definition of coherence, what it is for this reason to be a good reason is for P2 to probably be true whenever P1 is true. This is not something that is internal to us; there is an external fact of the matter whether the truth of P1 makes P2 likely to be true as well, and it would be this external fact that would determine whether the reason is a good one.

An alternative characterization of coherence is in terms of inference. Laurence BonJour says the following about the connection between coherence and inference:

> As the classical proponents of coherence have always insisted, coherence must involve some sort of positive connection among the beliefs in question, not merely the absence of conflict. […] But what sort of positive connection is required and how strong must it be? The obvious answer to the first question is that the connections in question are *inference relations* […] This much would be accepted by most, if not all, proponents of coherence theories. (BonJour 1985, 96)

P2 coheres with P1 if and only if there are good inferential patterns that go from P1 to P2. Yet, the question of what makes an inference a good inference is virtually the same question as the question of what makes a reason a good reason. BonJour also shows that the most straightforward analysis of inference is in terms of truth:

> The basic requirement for such an inference relation, as suggested in the earlier discussion of epistemic justification, is that it be to some degree truth-preserving; any sort of relation which meets this requirement will serve as an appropriate positive connection between beliefs, and no other sort of connection seems relevant here. (ibid.)

The question "What is inference" is very close to the question "What is reason?", so little to no progress is made by moving from reason to inference. The most straightforward answer to these questions brings in truth and thereby blurs the distinction between external and internal reasons.

The difficulty is that rationality is normative; it is a standard. Where does this normativity come from? Well, ordinarily, we would say that it comes from our aim for truth,

yet this answer would not work for personal inquiry. In fact, I argue, in what follows, that this does not work for shared inquiry either. We must reject this easy answer and try to dig deeper into our practices of giving and asking for reasons.

2.1 Reasons and Rules

If the standard for our reasoning is not provided by truth, then it would make sense to say that it is provided instead by the norms of assertion: to reason well is to adhere to the rules of the language game. Just like how the rules of chess determine which chess moves are allowed, the rules of our language games determine which patterns of reasoning are allowed. Yet, if we take this route, we must say something about what it is to grasp a rule and what it is to follow a rule, which is a notoriously tricky topic. Wittgenstein has posed a famous challenge to the very idea that our linguistic practices are guided by rules. (Wittgenstein 1953) This challenge has been brought into prominence in part because of Saul Kripke's interpretation of the problem in his book *Wittgenstein on Rules and Private Language* (1982). Below I shall give a quick exposition of Kripke's arguments, and I shall argue that we should accept his skeptical solution. I then show how we can think of Ferrari & Pedersen's divide between internal and external reasons, given these considerations.

Wittgenstein has shown that our application of a concept can be made to accord not just with the rule that *seems* right to us but with an infinite set of alternative rules as well. Using Kripke's example, suppose I have made many additions in my life, but I have never added 57 to 68. It seems to me that to act in accordance with the rule implied by the concept of addition, this should add up to 125. Yet, there is no way of ruling out that the correct application of the concept dictates that the answer is 5. (Kripke 1982) I have acted in accordance with not just my own envisaged rule of addition — provided to me by my grasp of the concept of addition

— but also with a rule that states that all my previous applications of the concept of addition were correct but that adding 57 to 68 should make 5. My past application of the concept of addition can be seen as in accordance with countless alternative rules, and therefore it would be hard to understand this application as rule-governed in the first place. Wittgenstein thus concludes: "no course of action could be determined by a rule". (Wittgenstein 1953, §201)

Wright characterizes two responses to the problem of rule-following, which he argues Wittgenstein ultimately finds lacking: the communitarian response and the Platonist response. (Wright 2001) The Platonist response is to say that there is an objective rule that guides the application of the concept of addition. To appeal to truth is essentially to embrace a Platonist response. We can say, for instance, that the correct application of the rule for addition demands of us that we say that 57 and 68 add up to 125, not 5, because it is true that 57 + 68 = 125 and false that 57 + 68 = 5. The problem with this response is that it renders our grasp of rules mysterious. It turns the correctness of application into something that is settled objectively, and that is thus external to us. It says that the solution of the addition ought to be 125, but it does nothing to show how we could possibly come to know this and exclude the possibility that the correct solution is 5. It gives us a rule, but no way of following this rule; it is a rule that cannot guide us. Therefore, it does not give us a rule at all. What is needed for our responses to count as rule-governed is the possibility that they satisfy the rule not due to a mere instance of good fortune but as the result of having grasped the rule, and, as such, it must not be wholly external. This means that the rule-following problems are not solved by claiming, as the Platonist does, that the rules are set from a mind-independent third-person perspective.

However, we cannot do away with a third-person perspective entirely since if the access to such standards is wholly internal, we could not distinguish between the correct application of a rule and the application that we merely take to be correct. This would lead to the inverse problem: whereas the Platonist response gives us rules that we cannot meaningfully follow

because they are not internally accessible — and that therefore aren't rules at all — an internalist response gives us rules that we could not fail to follow — and that, thus, likewise aren't rules at all.

The communitarian account offers a solution by taking the correct application of a rule to be determined by the community so that individuals who belong to this community can err in their application of the rule. A problem with this response is that it would still not be possible for the community to err, to have a collective failing, so that the correct application of a rule is merely the application that happens to be in accordance with our community, which would not be a very robust sense of correctness. As Wright says: "the evident awkwardness with this idea is that it seems to reduce the correctness of an assessment to a kind of marching in step." (Wright 2007, 485)

However, this should not be of much concern in the case of personal inquiry. I have argued that there are no objective truths for personal inquiry. There is no evidence-transcendent standard of correctness for personal inquiry, and therefore there is no one right view to be had about what is *really* beautiful, moral, funny, tasty, etc., independently of what anyone might think. We could say that the subject matter of personal inquiry resides in the social realm and has no significance outside of this realm anyways.

Even for the realist, on my construal, this would not be very problematic since what they are after, practically speaking, is to coordinate opinions. In fact, this curious compatibility with realism is exactly why the communitarian response won't work for anti-realism. Suppose I do not like apples because I think they are too tart. You disagree with me. On the view we are currently exploring, my reason for not liking apples would be a good reason if it is a correct application of the concepts in question. It would thus be a good reason if the concept of unpleasant tartness (or something like that) is such that it applies to apples. However, whether it does or not is exactly what we disagree about. The Platonist response has it that there is an

objective fact as to whether it does, which basically leads to the external view on reasons, according to which a reason is good insofar as it is truth-preserving. The communitarian response says instead that it depends on what the community agrees on, so that the legitimacy of our opinions is entirely a matter of whether they align with our peers, which would motivate us to join the fold. This would require us to bring the realist apparatus of shared inquiry into play by adhering to the truth norm and by conciliating in the face of peer disagreement.

2.2 Kripke's View on Rules

Kripke himself offers the interesting and much-discussed "skeptical solution" to the problem. He offers this not as a straightforward solution because it does not show how facts about rules can be discovered or constructed, but he instead accepts the conclusion of the paradox by endorsing a kind of non-factualism about rules. He shows that, in the absence of actual correctness conditions for rules, there are still interpretations about which performances *ought* to count as the correct application of a rule. The matter is thus normative all the way down, so to speak.

Kripke says:

> The solution turns on the idea that each person who claims to be following a rule can be checked by others. Others in the community can check whether the putative rule follower is or is not giving particular responses that they endorse, that agree with their own. (Kripke 1982, 101)

Put this way, the solution is closely related to the communitarian response in that the "correct" application of a rule is one that aligns with the application of the community. The only difference is that the communitarians attempt to create a genuine sense of correctness from this alignment, whereas for Kripke, there is no such thing as a correct application of a rule as such;

it is only that the community *interprets* an application as correct when it aligns with its own application of the rule. There are no facts about which application is correct, only interpretations.

This difference seems of little help to our current project. However, the idea that there is only interpretation — only attitudes about which application of a rule *ought to be* correct, and no facts about which application *actually is* correct — can be exploited for our current purposes.

The focus on the social element of rule-following, evident in both the communitarian and Kripke's response, is meant to create a normative wedge between one's disposition to apply a rule and the "correct" application of a rule. Yet, both responses do so by relying on a kind of collective, shared interpretation of a rule which the individual violates.

Kripke gives the example of a grocer. Since we all agree that addition is such that when two is added to three, we get five, I am entitled to keep my grocer in check. When I ask for two apples and three oranges and I find only four pieces of fruit in the bag, I won't exercise tolerance toward the grocer's individual application of the rule of addition, but I force them to play by my rules. A disagreement about what the *real* rule for addition dictates is not in order. It is no use for the grocer to appeal to such philosophical contemplations since the harsh reality is that such philosophical stubbornness will not be appreciated by their clientele, which is not concerned with finding out the *real* sum of two and three, but which will simply go to another grocer who conforms to their arithmetic proprieties. (Kripke 1982)

This example shows that the push to convergence is built into Kripke's account of rule-following. Kripke endorses, what Brandom calls, an *I-we* conception of rules. Brandom himself endorses an *I-thou* conception of rules instead. (Brandom 1994) He shows that the proper application of a rule almost always depends on one's collateral beliefs or opinions. For instance, if I find tart food disgusting and you do not, then, while the circumstances in which we may

take food to be tart might be the same, the consequences of doing so would be vastly different. Hence, while apples being tart might be a good reason for me for finding them disgusting, it would not be a good reason for you. Kripke would argue that, in this case, we appeal to the perspective of the community. Yet, since everyone's taste is unique, everyone has different collateral beliefs and, therefore, everyone has a different way of applying concepts such as tartness. The idea of appealing to the community would break down here.

Brandom avoids the appeal to the community altogether. The mere presence of multiple different interpretations of a concept is enough to create the normativity needed to speak of a rule. On my interpretation, I am right and you are wrong, and vice versa. Brandom thinks that we attribute a fault to applications of concepts that we interpret as incorrect and that the normative force of the rule stems from attributing such a fault (which entails, on the practical level, some kind of sanctioning). On this view, whether a person's actions are guided by a rule is a matter of whether they can be interpreted as such, not by this person themselves, but by another person. Whether they can be interpreted as following a rule by another person depends, in turn, on whether this other person can take them to be more or less aligned with the rule they take themselves to be following (though the situation is, of course, symmetrical and whether they are themselves following this rule can only be determined from the perspective of another).

Attributing fault to (what one takes to be) divergent applications of a concept is like attributing fault to a disagreement. It stems from the same kind of intolerance toward incompatible opinions. Therefore, even Brandom's *I-thou* model implies that our inquiry is shared. It implies that we believe that there is only one correct interpretation of the rule in question. However, normativity could be generated even if we allow people to follow their own rules. This may seem like handing over the authority to judge the correctness of performances to the performer. In that case, it would fall prey to Wittgenstein's and Kripke's criticisms since the performer cannot distinguish between correct performances and performances that they

merely take to be correct. However, this need not be the case. An example is in order to show the crucial difference between allowing people to follow their own rules and allowing them to determine for themselves whether their performances satisfy their rules.

Suppose that I have only played tennis once or twice in my life and that the way I have learned to play tennis deviates from the way in which you play tennis in that I think the ball can bounce not once but twice. You might first try to have me play by your rules, which you, of course, take to be correct, but you might quickly discover that this is not much fun because you are a much more experienced player. You then suggest that we instead play by our own rules: I can let the ball bounce twice, whereas you can only let it bounce once. Allowing me to play by my own rules does not prevent you from keeping me in check, although it might at first seem that it does. Suppose, for instance, that we are at match point, you play a magnificent drop-shot, and it bounces thrice before I get there. You claim the victory, but I say that one of my rules, which I forgot to mention, is that on match point, the ball can bounce thrice. You, of course, see me for the cheater that I am, and you say that this is not what we agreed upon. I then say that we agreed to play by our own rules and that what *my* rules are is up to me, not you.

However, in allowing me to play by my rules, you are not thereby granting me all of the authority regarding these rules, and nor could you be: if my rules are solely up to me, then, Wittgenstein shows, there are no rules to begin with. Rules must be socially articulated for them to have their distinctive normative force. Thus, what you are doing instead is allowing me to play by *your* interpretation of my rules, not by *my* interpretation of my rules.

What is required to keep each other in check is not that we all try to play by the same rules, but only that we base our judgment of a performance of an individual on the rules that *we* think *they* play by. Therefore, the only thing that is required is that one's performance can be interpreted as satisfying *some* rule. Of course, all performances whatsoever could

theoretically be interpreted as satisfying some rule, but it is up to the interpreter to decide whether someone's actions can be seen as rule-governed or not. They do so chauvinistically, of course, by putting themselves in another's shoes and determining what rules should bind them. Absurd skeptical interpretations of rules are thereby excluded because they defy our imagination. This chauvinism — this willingness to interpret someone as following a rule only insofar as we can imagine ourselves to be bound by such a rule had we been in their shoes — is of great practical value because the utility of taking others to be guided by rules would be defeated by allowing all possible actions to accord to some rule.

2.3 Personal and Shared Rationality

For personal inquiry to function as it seems it does, we must interpret the opinions of others (and of ourselves) as structured. This practice of connecting our opinions is what makes it appropriate to speak of something like "coherence" as an aim for this kind of inquiry (and as a standard of internal reasons).

This is where Kripke's skeptic tries to intervene by suggesting that one's opinions can be interpreted as structured in some deviant way. The case of arithmetic actually displays this structure, or lack of structure, very well. We think that there is some neat structure to addition, to the pattern it creates (1, 2, 3, 4, 5 …), but the skeptic tells us that, for all we know, our rule for addition breaks up this pattern somewhere (123, 124, **5**, 126 …). There is no way of ruling out whether we or anyone else are not *actually* following such a deviant rule. However, there is no practical point in keeping this option open because this kind of structure is beyond our comprehension, and therefore we cannot see any structure at all anymore. In interpreting others, we must restrict ourselves to structures we can make sense of. Yet, this does not mean that we must suppose that everyone structures their opinions in exactly the same way as we take

ourselves to structure our opinions. If the purpose is only to keep each other in check, we can allow for different structures to be in play as long as we can make sense of these structures.

I find the labels "internal" and "external" misleading because rationality can neither be wholly internal nor wholly external. Yet, there is still a contrast between the kind of rationality that guides personal inquiry and the rationality that shared inquiry aims for. I shall use the labels "personal rationality" and "shared rationality" to highlight this contrast. Please keep in mind though, that personal rationality is not meant to convey that people can make up and follow rules without any social constraints. Rationality must still be socially articulated. The difference between personal and shared rationality is perhaps most easily understood by keeping the tennis example in mind. If rationality is *shared*, we force each other to play by the same rules, whereas if it is *personal*, we allow others to play by our interpretation of their own personal rules.

It might be asked whether structure is really all that important when it comes, for instance, to matters of basic taste or comedy. Could we not just find things tasty or funny without worrying about coherence? We could, but not if we are participating in a game of giving and asking for reasons. To have disagreements about such matters, we must think of our opinions as structured. Perhaps we only start to look into the patterns between our opinions in response to these games of giving and asking for reasons. Nevertheless, asking why we hold certain views, and finding some structure between our views to respond to this question, is practically useful since it allows us to keep each other in check and because it puts us in a better position to predict how our tastes will respond in unexplored territory. For example, understanding *why* you thought this comedian was funny better enables you to seek out comedy acts that you will find funny and avoid those that you would find a bore, even though you are expected to find countless of exceptions to whichever rules you may come up with.

I argue in the next chapter that philosophical inquiry is personal as well. However, one might struggle to understand how this could possibly be the case. It seems, *prima facie*, that philosophical inquiry is on the other end of the spectrum from, for example, matters of taste or comedy. Structure is incredibly important to our philosophical views. Internal coherence is probably the principal criterion by which we evaluate philosophical theories. Furthermore, it seems that this structure must be shared. How do we make sense of the idea that you can follow your own rules and I can follow mine?

The next chapter argues that the common perception of philosophical inquiry as a scientific inquiry has made it seem as if our standards are, or ought to be, wholly shared. However, as I mentioned in the introduction to this thesis, one of the ways in which philosophy differs from science is that its methodology is by no means common ground. How to carry out philosophical inquiry is itself a philosophical issue. Methodological principles are the consequence of object-level philosophical commitments nearly as often as they are their premise. Given the academic nature of the discipline, however, we tend to group together in philosophical schools or isms. I argue in the next chapter that this is less like scientific practice and more like artistic movements, which also clump together in schools and isms. A pragmatist could read a work in speculative metaphysics, and they will most likely find it mistaken and false, but they may appreciate that it is internally coherent, just like we can appreciate a work of art as a skillful execution of the cubist style of painting while also having a general distaste for cubism.

Just because a philosophical work is rational does not mean that we must think it is correct. The case of philosophy really is more similar to the case of basic taste than it might appear. I mentioned the example of whether apples being tart would be a good reason for apples being disgusting. This depends, obviously, on whether we find tart food disgusting. If we do, then this is a good reason for us. If we know that our interlocutor happens to love tart food,

then we would object to their using apples being tart as a reason for them being disgusting. In philosophy, the question of whether a reason is a good one also depends on one's background commitments, and it is these background commitments that *may* ultimately turn on a mere matter of taste. As such, we can call an opponent's view false without attributing a fault to this view, just like we can say that it is false that tart food is disgusting without attributing a fault.

It is, of course, also the case that, even in scientific inquiry, our reasons depend on our background commitments. However, in that case, we are never just allowed to have different background commitments as if they are mere matters of personal taste. This would get in the way of convergence. Personal rationality is mostly a tool for keeping each other in check, but if our goal is to cooperate with one another, then we need a kind of shared rationality. If our purpose is to converge in our opinions, as it is for shared inquiry, then we must try to play by the *same* rules. When we come up with the same views, our confidence in the premise that we are indeed following the same rule grows stronger (even though it would theoretically be possible that we are all following different rules that only happen to lead us to the same opinions).

Thus, one is *personally* rational whenever one can be interpreted as following their own rules, whereas one is rational in a *shared* sense whenever one can be interpreted as following the same rules as the ones we take ourselves to be following. This extra requirement on rationality is explained by the cooperative nature of shared inquiry. In shared inquiry, we pool our resources so that we can help each other gain better beliefs. Yet, we can only do so if we take each other to be guided by the same rules.

Once we think that the rules we take ourselves to be following are the same as the rules most others are following, we can think of our interpretation of a rule as the *correct* interpretation. In the case of basic arithmetic, we interpret ourselves as following the same rule as pretty much everyone else, and we think of this rule as the "correct" one. Therefore, as far

as we are concerned, we are judging whether the grocer adds in a way that is *objectively true*. We are not concerned with finding out whether the grocer acts in accordance with their own rules for addition but whether they act in accordance with *our* rule for addition. This is the sense in which we can take the standard for the reasoning of the grocer to be *external* to the grocer, and it is the sense in which we take the standard to be objective. Yet, while this standard might be external to the grocer, it is not wholly external since it is internal to the interpreter. This is why it would be misleading to call this kind of rationality "external rationality".

2.4 Steadfastness and Conciliation

Ferrari and Pedersen (2019) draw a connection between the two different kinds of rationality and what would be the rational response to peer disagreement. Using their terminology, Ferrari and Pedersen argue that the conciliatory response depends on internal rationality. I take it that their argument is as follows: if the rationality of one's response is determined externally, then the person who defends the view that turns out to be true is rationally entitled to stick to their view, whereas the person who defends the opposing view is rationally required to change their belief. Yet, if there is no external notion of rationality, then there is no external perspective from which one side can be said to be correct and the other incorrect, so in that case, the situation is symmetrical, and both parties must conciliate.

I think that this line of reasoning is wrong and that the mistake is caused by Ferrari and Pedersen using the labels "external" and "internal". They think of external rationality as wholly external to anyone's point of view; what is externally rational depends only on the objective truth. Yet, such a notion of rationality does not survive the rule-following problems mentioned above. A rational constraint on our doxastic attitudes must be a normative constraint. Therefore, what it is for our doxastic attitudes to be rationally constrained is for them to be the result of

our acting on a rule. However, the rule "Stick to your beliefs if they are true and change them if they are false" is not something we can meaningfully act on, which is why it cannot be a genuine rule at all.

Since there cannot be such a notion of "external rationality", symmetry holds for all peer disagreements. Yet, it is not the case that conciliation would always be the rational response. As argued in Chapter 3, steadfastness is a generally rational response to disagreement in the case of personal inquiry. My account of personal rationality shows why this is the case. If our inquiry is shared, we must think of ourselves as playing by the same rules, and disagreement would be an indication that we might have misapplied these rules. Yet, if we allow others to play by their own rules, disagreement should not come as a surprise. It would not be indicative of the possibility that we have misapplied our own rules. Instead, it is indicative of the possibility that our interlocutor plays by a different set of rules.

For instance, I can allow my interlocuter to use the belief that apples are tart as a reason for finding apples disgusting, though I do not believe that this rational constraint applies to myself; I can both maintain that apples are tart and that they are not disgusting. Therefore, once we find ourselves disagreeing on whether apples are disgusting, we *could* believe that we might have made a mistake and reduce our confidence in our belief, but we are *rationally entitled* to believe instead that the disagreement is the result of diverging sets of rules, and we are thus rationally allowed to remain steadfast in our belief.

The picture is the exact opposite of what Ferrari and Pedersen argue: for those inquiries that are constrained by shared rationality, conciliation would generally be the rational response to peer disagreement, whereas, for those inquiries that are constrained by personal rationality, we can be rationally entitled to remain steadfast in response to disagreement. This aligns with what I have argued for in Chapter 3.

3. Knowledge

I have defined personal inquiry as inquiry that lacks the truth norm and that is instead guided by non-alethic norms, such as the norm of justification. The previous section has shown how this norm of justification can be divorced from the norm of truth, by showing how we can aim to be rational without thereby implicitly aiming for truth. The question that I shall ask in this current section is whether personal inquiry can successfully aim for knowledge, and what personal knowledge would amount to. In other words, can I really know that apples are tasty or that sunsets are beautiful? And, if so, could personal inquiry aim for such knowledge?

3.1 Attributions of Knowledge

The difference between merely being justified in believing that *p* and *knowing p* has traditionally been thought to turn on the truth of *p*. This is the truth condition of the JTB (justified true belief) account of knowledge. If this is the difference between knowledge and mere justified belief, what it is for our inquiry to aim for knowledge rather than for mere justification, is for it to aim for truth in addition to justification. This would preclude us from saying that personal inquiry aims for knowledge since what makes this inquiry personal is exactly that it does not aim for truth. If this is right, then instead of saying that the aim of personal inquiry is the accumulation of knowledge, we should say that the aim is merely to gain justified opinions.

One important thing to note is that personal inquiry is still truth-apt, which is why we can unproblematically say that there is knowledge in those cases in which one's belief is justified and true (and not a Gettier case (1963)). The question is whether we can *aim* for knowledge without thereby aiming for truth. If there were no truth at all, then we for sure

couldn't aim for knowledge because there would likely not be any knowledge at all; there would only be opinions. But again, this is not the case for personal inquiry.

There is a curious difference between two ways of speaking of knowledge. I shall argue below that there is a personal and a shared aim for knowledge. This distinction is based on the distinction between the attribution of knowledge to another person and the self-attribution of knowledge. When attributing knowledge to someone else, we can say (ignoring Gettier cases) that *S* knows *p* if and only if *S* justifiably believes that *p and p is true.* This is the perspective I assumed above. Yet, when we assess whether we know *p* ourselves, we try to do so by assessing how sure we are of our belief in *p*.

This difference is striking. For us to be able to say of ourselves that we know that *p*, we must be certain of this belief. Absolute certainty is, of course, impossible to achieve, and, as pragmatists, we should not take absolute certainty to be our standard. Still, we must at least be *practically certain* of this belief, meaning that we do not have any actual doubts about *p*, nor do we anticipate that we shall come to doubt *p* in the future. (Misak 2004) Yet this requirement of practical certainty is not placed on others when we attribute knowledge to them. The only thing that matters from this perspective, is that *S* does in fact justifiably believe *p* — even if *S* is only somewhat confident in their belief that *p*, e.g., if they give this belief a credence of 0.7 — and that *p* is true.

3.2 Does Personal Inquiry Aim for Knowledge?

The above has shown that I can reasonably take myself to know that *p* if and only if I am practically certain that *p.* Insofar as we aim for practical certainty in personal inquiry, it could be said that we aim for knowledge. This might sound confusing. If knowledge is something like justified true belief, does it not follow that if we aim for knowledge, we thereby aim for

truth in addition to justification? Not quite, because, from the first-person perspective, we cannot aim for knowledge by asking ourselves: "Do I believe *p*? Check. Am I justified in believing *p*? Check. Is *p* true? Check." If we want to figure out the truth of our belief, we have to figure out whether our belief is justified, and that is all we can do. Therefore, what we actually do when we try to gain knowledge is to try to gain certainty, which amounts roughly to trying to gain justification that is so strong that we can reasonably expect it not to be defeated, come what may. The aim for knowledge, from the first-person perspective, is the aim for certainty.

One might ask whether there is any difference at all, from a first-person perspective, between aiming for justification and aiming for knowledge. This question is similar to Richard Rorty's question, discussed in Chapter 2, about the difference between aiming for justification and aiming for truth. (Rorty 1995) We saw then that the main difference cannot be appreciated when focusing on a single inquirer, but it has to do with the interactions between inquirers; the aim for truth puts pressure on us to coordinate our beliefs, whereas the aim for justification does not. Below I shall argue that, *from the third-person perspective*, the difference between the aim for knowledge and the aim for justification is just like this difference between the aim for truth and the aim for justification; the *shared* aim for knowledge puts pressure on us to coordinate our doxastic attitudes.

However, in the case of personal inquiry, this shared aim is lacking, and there is only the personal aim for knowledge, interpreted as the aim for practical certainty. Thus, the question still stands: is there any difference, from a first-person perspective, between aiming for justification and aiming for knowledge? I argue that there is and that, therefore, even this personal aim for knowledge is distinct from the aim for mere justification. There is an important *practical* difference between these two aims. The difference is that when we take ourselves to know *p*, we have a reason to take contexts of inquiry that count against *p* to be misleading. That

is, when we take ourselves to know *p*, we have *p* factor into the evaluation of the epistemic credentials of contexts of inquiry. Take, for example, the situation in which we take ourselves to know that apples are tasty versus the situation in which we take ourselves to merely be justified in believing that apples are tasty. In the latter situation, eating an apple that we do not find tasty would count against our belief that apples are tasty. It would make us doubt this belief, and it would prompt us to inquire further into the matter. However, in the former case, eating an apple that we do not find tasty would not count against our belief that apples are tasty, but it would count against the epistemic credentials of this particular data point. Perhaps we suffer from dysgeusia, have covid, or are pregnant. More likely, we would take something to be wrong with this particular apple. We might suspect, for instance, that it is slightly rotten. In any case, this situation would not make us doubt our belief that apples are tasty because we already *know* they are.

A possible worry is that this approach to knowledge leads to the dogmatism paradox. According to this paradox, once you take yourself to know that *p*, you take yourself to know that *p* is true and that, therefore, any evidence that counts against *p* must be misleading, for *p* is true. If knowledge closes inquiry, then we must discount all evidence that counts against *p*. Yet, this would be highly dogmatic. (Kripke 2011) There are countless of counterexamples imaginable in which we take ourselves to know something and then find out that our supposed knowledge is defeated. For instance, we could take ourselves to know that the sun is out because it is the middle of the day, and we were outside two minutes ago and saw that there is not a single cloud in the sky. Yet, someone might tell us that there is a solar eclipse happening at the moment. The paradox tells us that we must treat this testimony as misleading. We might then go outside and see that the sun is eclipsed with our own eyes, but again the paradox tells us that we must treat this evidence as misleading as well.

The solution, I think, is to say that knowledge closes inquiry only insofar as it removes the doubts that prompted the inquiry in the first place. Yet, practical certainty is not absolute certainty, so we must be fallibilists. Knowledge closes inquiry, but it doesn't seal it shut never for it to be opened again. It can be reopened, for instance, when we find that the different beliefs we hold with practical certainty conflict. For instance, we can be practically certain that our eyesight doesn't betray us. This is also not an infallible belief, but it is stable enough for practical certainty. So, when it conflicts with another belief that we hold with practical certainty, then this certainty is defeated, and something has to give. Yet, this does not undercut the difference between aiming for justification and aiming for practical certainty. It is still the case that once we have gained practical certainty, we can close the inquiry and add the belief in question to our background beliefs moving forward.

This practical difference can be understood without bringing in the aim for truth. According to Peirce, the aim of inquiry is not the pursuit of truth but the cessation of doubt. He says:

> The irritation of doubt is the only immediate motive for the struggle to attain belief. It is certainly best for us that our beliefs should be such as may truly guide our actions so as to satisfy our desires; and this reflection will make us reject every belief which does not seem to have been so formed as to insure this result. But it will only do so by creating a doubt in the place of that belief. With the doubt, therefore, the struggle begins, and with the cessation of doubt it ends. Hence, the sole object of inquiry is the settlement of opinion. We may fancy that this is not enough for us, and that we seek, not merely an opinion, but a true opinion. But put this fancy to the test, and it proves groundless; for as soon as a firm belief is reached we are entirely satisfied, whether the belief be true or false. [...] The most that can be maintained is, that we seek for a belief that we shall *think* to be true. But we think each one of our beliefs to be true, and, indeed, it is mere tautology to say so. (Peirce 1877, 6)

Doubt impedes our acting, so we inquire insofar as we have reason to actively doubt our belief. Once we have dealt with any particular doubts we may have, the belief is settled, and we use it as a guide for our actions. We have no reason to further defend it from theoretical doubts, such as Cartesian skepticism, because these "paper doubts", as Peirce calls them, do not impede our acting. The goal is thus to have beliefs that we have no reason to doubt.

From a first-person perspective, the aim for knowledge is the aim for practical certainty. If we take ourselves to know *p*, we can add *p* to our background information and allow this belief to factor into the evaluation of epistemic contexts. This aim can be understood without it implying an aim for truth. What we are after instead, might merely be the cessation of doubt. From a first-person perspective, our aim for knowledge can be thought of as fulfilled once our beliefs enable us to effectively satisfy our desires and to go about our business in such a way that it does not prompt us to further doubt our beliefs.

3.3 Where Did the Third-Person Perspective Go?

Suppose we can indeed make sense of our aim for knowledge from a first-person perspective without invoking the aim for truth. There would still be the third-person perspective from which the aim for knowledge does seem to involve the aim for truth. However, Brandom has shown that this perspective allows for a more deflationary reading. Brandom asks what the difference is between attributing a justified belief to *S* and attributing knowledge to *S*. The difference is, he argues, that one who attributes knowledge does not only attribute a commitment to *p* to *S* but also undergoes this commitment themselves. This allows for a deflationary reading of the JTB account of knowledge:

> First, the scorekeeper must *attribute* an inferentially articulated, hence propositionally contentful, *commitment*.

This corresponds to the *belief* condition on knowledge.

> Second, the scorekeeper must *attribute* […] *entitlement* to that commitment.

This corresponds to the *justification* condition on knowledge. What is it that then corresponds to the *truth* condition on knowledge? For the scorekeeper to take the attributed claim to be true is just for the scorekeeper to endorse that claim. That is:

> Third, the scorekeeper must *undertake* the same propositional commitment attributed to the candidate knower. (Brandom 1994, 515)

A scorekeeper is someone who attributes doxastic commitments to someone else, such as beliefs or knowledge. According to Brandom, a scorekeeper can attribute a belief *p* to *S* by stating or showing merely that *S* is committed to believing *p*. Yet, for them to attribute *knowledge* of *p* to *S* they must additionally undertake a commitment to believing *p* themselves. Put more straightforwardly, the difference between me thinking that you believe *p* and me thinking that you *know p* is just that in the latter case, I believe *p* myself, whereas this is not necessarily so in the former case.

I argued before that from a first-person perspective, it makes no sense to first ask "Am I justified in believing *p*?" and then ask "Is *p* true?", since the latter question can only be answered by addressing the former. Yet, from a third-person perspective, this does make sense. Thus, if knowledge is something like justified true belief, we can ask whether *S* knows *p* by asking: "Does *S* believe *p*?", "Is *S* justified in believing *p*?" and "Is *p* true?" However, as

Brandom shows, what we do when we address the third question is address the first two questions when applied to ourselves: "Do *I* believe *p*?" and, if so, "Am *I* justified in believing *p*?"

A small correction of Brandom's view is in order, though. If I take both myself and *S* to justifiably believe *p*, then I would not automatically say that either I or *S* knows that *p* because I might not be practically certain of it. We would only attribute knowledge to others when we take them to justifiably believe *p* and we believe *p* as well *with practical certainty*.

The practical role of knowledge attributions is similar to what deflationists like Lionel Shapiro (2021) take the practical role of truth attributions to be: they are ways of showing our alliances, and of passing on the burden of justification to somebody else. Just like I can say "What *S* said is true", thereby offering to share the burden of justification for *S*'s claim — by either deferring to *S* or allowing *S* to defer to me — I can say "*S* knows that *p*", thereby doing virtually the same thing.

This analysis of knowledge attributions allows us to now make sense of both self-attributions of knowledge as well as attributions of knowledge to others without even having to mention truth. To say that "I know that apples are tasty" is to say that I am practically certain that apples are tasty, that I do not expect to ever doubt this belief. To say that "You know that apples are tasty" is to express that you believe that apples are tasty and that I am practically certain of this belief. It is conceptually possible to conceive of some of our practices as containing knowledge attributions, both to ourselves and to others, but as only aiming for knowledge in the first-personal sense. Such practices only aim for personal knowledge, i.e., practical certainty.

3.4 Two Ways of Aiming for Knowledge

As mentioned above, the aim for knowledge, from a personal perspective, is just the aim for practical certainty. This could be a rather individualistic enterprise. If I aim to know whether apples are tasty, I do not have to think of this aim in relation to others. Yet, in the case of shared inquiry, our aim for knowledge is interpersonal. For instance, we aim to know scientific facts not as individuals but as part of the scientific community. Brandom's analysis suffices to show that this aim amounts to the aim of coordinating our opinions. If knowledge attributions are expressions of shared doxastic commitments, then the aim for knowledge as something like justified *true* belief is the aim for beliefs that we are not just practically certain of ourselves, but the aim for us to have shared beliefs in the community. We aim to be on the same page, to be able to attribute the same doxastic commitments as we undertake ourselves.

This coheres with the story I have told about the aim for truth. By aiming for knowledge (in this shared sense), we aim for truth, and by aiming for truth, as Price (2003) has shown, we aim to coordinate our opinions. Brandom seems to agree with Price on this matter:

> Whenever two believers disagree, a diagnosis of error or ignorance is appropriate for at least one of them. Though agents with differing practical commitments can also be criticized on the grounds of error and ignorance, mere difference of desire or preference is not sufficient in general to make them liable to such criticism. We come with different bodies, and that by itself ensures that we will have different desires; what is good for my digestion may not be good for yours; my reason to avoid peppers need be no reason for you to avoid peppers. Our different bodies give us different perceptual perspectives on the world as well, but belief as taking-true incorporates an implicit norm of commonality — that we should pool our resources, attempt to overcome the error and ignorance that distinguish our different sets of doxastic commitments, and aim at a common set of beliefs that are equally good for all. Talk about belief as involving an implicit commitment to the Truth as One, the same for all believers, is a colorful way of

> talking about the role of testimony and challenge in the authority structure of doxastic commitment — about the way in which entitlements can be inherited by others and undercut by the incompatible commitments they become entitled to. (Brandom 1994, 240)

Brandom assumes that our practice of sharing justificatory burdens by making truth attributions and, indeed, also by making knowledge attributions, is driven by the aim to coordinate our opinions. However, in the first chapter, I argued that these mechanisms of sharing justificatory burdens do not actually presuppose this "norm of commonality". It is one thing to say that those who happen to agree in their opinions — those who happen to be similarly inclined — can exploit this agreement to defer to one another, but it is another thing to say that, for the sake of deference, we must *aim* to agree in our opinions, even if we are differently inclined. The utility of a mechanism for expressing that we endorse the same opinions is independent of the utility of the mechanism for converging toward the same opinions.

Knowledge attributions express shared doxastic commitments. We could take it as our aim that our doxastic commitments be shared, in which case we would aim for shared knowledge. Yet, this aim is optional, and, as argued in the previous chapters, it is not always practically advantageous to have such an aim. In the absence of this aim, we could still aim for personal knowledge by aiming for practical certainty. In that case, we would succeed once we have reached practical certainty, regardless of whether our community of inquirers agrees or disagrees with our beliefs. This is the way in which personal inquiry aims for knowledge. Shared inquiry, on the other hand, aims for knowledge in a more ambitious sense, by aiming not just for practical certainty but also for shared belief.

The latter requirement also affects the former requirement. If I am a dissenting scientist, I might hold uncommon beliefs that turn out to be correct, yet I cannot say that I *know* these things I believe since the mere fact that my community does not share my beliefs is generally enough to undercut my justification for these beliefs. This is what the conciliatory view on peer

disagreement tells us, and conciliation is the rational response to peer disagreement *because* we aim to coordinate our beliefs. We stand or fall together, so to speak. Yet, without the aim for coordination, I could say that I *know* things even if my community disagrees with me — as is shown by the steadfast view on peer disagreement — because the mere fact that my community does not share my beliefs does *not* undercut my justification. I can therefore be practically certain of the fact that, for example, salty licorice is tasty, and I can ascribe knowledge of this fact to myself, in spite of being fully aware that most people do not share this belief.

4. Knowledge and Testimony

The above makes it seem as if knowledge is more easily gained by personal inquiry than by shared inquiry. If we are constantly put in doubt by the beliefs of our peers, then we can only gain knowledge when disputes are settled. Personal inquiry does not have this constraint: if we are practically certain that salty licorice is tasty, then this certainty need not be undermined by the mere fact that many people disagree with us. However, there is a flip side to this coin: we cannot rely on others as much for gaining knowledge in personal inquiry as we can in shared inquiry. I can say that I know, for instance, that the gravitational constant is 6.6743×10^{-11} m^3 kg^{-1} s^{-2}, even though I do not understand why this is so. That is, I can know this even though I am not aware of any reasons that count in favor of this belief. I only need to be aware of the fact that the scientific community has settled this issue and that this is the consensus that has been reached. In the case of personal inquiry, matters are more complicated.

In the case of aesthetics, for example, it is widely assumed that though our beliefs can be informed by aesthetic testimony, we cannot really *know* that the play is masterful or that the book was disconsolate without understanding why this might be so, i.e., without evaluating the

matter at issue ourselves and appreciating the reasons behind these aesthetic judgments. (Robson 2012) Robert Hopkins has argued that this is because the mechanics of testimony are different in the case of aesthetics when compared to cases of, for example, mundane observational beliefs. (Hopkins 2000) In the case of mundane testimony, we are warranted to believe the testimony *unless* there is a particular defeater, such as when we know the source to be unreliable. Yet, in the case of aesthetic testimony, we should first establish that the source is reliable, i.e., that their aesthetic beliefs generally overlap with ours, before we could believe the testimony.

Hopkins suggests that this difference is due to the lack of Cognitive Command in the case of aesthetic inquiry. I want to make a similar claim that it is due to the fact that aesthetic inquiry is personal; it is not rational to assume that another's aesthetic experiences and judgments are generally the same, which is why it is only rational to pay heed to someone's testimony if we have evidence that our aesthetic experiences and judgments are generally the same. Yet, in the case of shared inquiry, it *is* rational to assume sameness which is why it is generally rational to believe someone's testimony unless there are concrete reasons not to.

The same pattern seems to hold for moral testimony. The literature on moral testimony generally focuses on whether we are *morally at fault* if we fully rely on testimony for our moral beliefs. For instance, Philip Nickel argues that an individual's actions are only morally good if they are based on moral understanding; the individual must not merely know which actions are morally right and act on this knowledge, but they must understand the reasons behind this knowledge. (Nickel 2001)

Karen Jones argues against this thesis that it would be morally wrong to rely fully on testimony. According to Jones, there are moral experts whose testimony we should trust even if we cannot understand the reasons behind it. For instance, we should generally think of women as being in a better epistemic position to know which actions are sexist, and, as such,

if men do not understand why certain actions are sexist, they should simply take women's word for it, instead of rejecting all moral testimony that they cannot understand. (Jones 1999)

Both Nickel and Jones try to answer the question of whether morally virtuous actions require moral understanding in addition to knowledge. The question that concerns me is different because it asks not whether it is morally virtuous to aim for moral understanding in addition to knowledge, but whether testimony without understanding could give us any moral knowledge at all. I do not deny that it is rational to believe the testimony of those we take to be experts and that this will give us justification for certain moral beliefs. Yet could testimony on its own give us moral knowledge? From a third-person perspective, it could because it could give us justified true belief, but from a first-person perspective, it does not seem that moral testimony alone could generally give us practical certainty. One of the reasons for this is that moral experts are hard to find, and it is difficult to believe with practical certainty that someone is a moral expert. It seems likely that, just like in the case of aesthetics, we must understand the reasons behind a moral judgment to be practically certain of it.

An exception to this rule is basic judgment, the kind of judgment for which no further reasons can be given. We can be practically certain of a basic moral judgment, such as the judgment that causing pain is wrong, even without understanding *why* this is so. The same holds for basic taste judgments. Yet, it seems unlikely that we can gain such basic judgments by means of testimony; the very thing that makes them basic is that they just strike us as correct independently of what others think. If we gain such beliefs through testimony, they are not basic for us.

This shows that in the case of personal inquiry, the opportunity to gain knowledge through testimony is significantly smaller than in the case of shared inquiry. In the ordinary cases in which someone tells us that a play is masterful or that an act is immoral, we have to either formulate our own understanding of why this would be the case or, if the belief expressed

is a basic one, we have to experience it for ourselves or form our own intuitions on the matter. Therefore, while on the one hand, we could gain knowledge more easily because we need not be bothered by the beliefs of others, on the other hand, it could be more difficult to gain knowledge in personal inquiry because we generally cannot rely on the beliefs of others either.

5. Conclusion

In this chapter I have argued that personal inquiry is rational and that it aims for knowledge, just like shared inquiry. However, there is a distinction to be made in the kind of rationality that personal inquiry is governed by and likewise in the kind of knowledge it aims for.

Personal rationality allows us to play by our own rules, or at least by another's interpretation of our own rules. The constraints are looser than in the case of shared rationality; our beliefs must still be structured and coherent in a way that allows others to make sense of them, but there is no expectation that by employing the tools of rationality, we shall come to adopt the same beliefs.

I have shown that personal inquiry aims for personal knowledge but not for shared knowledge. This personal knowledge consists in practical certainty. For personal inquiry, knowledge is possible even without consensus, whereas in the case of shared inquiry, the lack of consensus undercuts justification, and thereby keeps knowledge out of reach. That said, there is a much greater opportunity to know without understanding in the case of shared inquiry, by relying on testimony, than in the case of personal inquiry. Therefore, personal inquiry must be seen as aiming not just for knowledge but for knowledge through understanding, or in the more basic case, as aiming for knowledge through a kind of personal acquaintance with the relevant intuitions or experiences.

In the next chapter, I shall apply these thoughts to the case of philosophy. In philosophy, understanding also seems like a general requirement for knowledge. Something would not quite be right if a philosophy student were to claim that they know that moral realism is true just because their professor said so. Note that in the case of science, this would be more acceptable, e.g., a physics student could say that they know that the gravitational constant is 6.6743×10^{-11} m^3 kg^{-1} s^{-2} because their professor said so. Also, if we understand philosophical inquiry as aiming for personal knowledge, we can see how widespread persistent disagreement need not get in the way of the accumulation of knowledge and, thus, of philosophical progress.

6. Philosophy as a Personal Inquiry

1. Introduction

In previous chapters, I have made a distinction between shared inquiry and personal inquiry. I have argued that personal inquiry differs from shared inquiry by lacking the aim for truth. Nevertheless, personal inquiry is still apt for truth, guided by reasons, and it can still aim for knowledge. The way in which it aims for knowledge also differs from the way shared inquiry aims for knowledge. In shared inquiry, conciliation is generally the rational response to peer disagreement, which is why widespread disagreement undercuts the confidence which is needed for us to attribute knowledge to ourselves. In personal inquiry, widespread disagreement need not get in the way of knowledge because we can be rationally entitled to remain steadfast in the face of disagreement. Therefore, we could say, for instance, that we know that pistachio ice cream is tasty or that eating meat is wrong, even though we are aware that many people disagree with us on these matters. The flip side of this coin is that we cannot say that we know, for example, that *Death of a Salesman* is a masterful play without having seen it ourselves. That is, we cannot rely on the testimony of others in the same way as in shared inquiry because we cannot assume that everyone experiences or processes things in the same way. Knowledge without understanding is harder to come by in personal inquiry because we can generally only take another's testimony on board once we understand why they hold the beliefs that they do, and once we find ourselves persuaded by their reasons.

In this chapter, I argue that philosophical inquiry is personal, not shared. The mechanics of testimony and peer disagreement are similar to other personal inquiries and unlike shared inquiry. In section 2, I point out how many of our philosophical practices — including a philosopher's day-to-day activities, such as attending talks, teaching, writing, and peer reviewing — indicate that philosophers have a higher tolerance for disagreement than scientists.

In section 3, I make the case that this seems to be more in line with how we carry out personal inquiry than with how we carry out shared inquiry. Section 4 argues against methodological naturalism about philosophy, which is the view that philosophical inquiry should be conducted like scientific inquiry. Section 5 considers what the thesis that philosophical inquiry is personal tells us about the nature of philosophical inquiry. I argue that philosophy is generally not about describing reality, and that it is not objective in the way that science is. Instead, philosophical theories offer worldviews, which are ways in which we *ought* to think about the world. Finally, section 6 argues that there is progress in philosophy, in spite of the existence of widespread and persistent philosophical disagreements.

2. Philosophical Practices

There are countless differences between philosophical and scientific practices, some of which I have already mentioned. One of the most important differences is how philosophers deal with disagreement. This seems to follow the model of personal inquiry. Disagreement is expected in philosophy, and it is something that we allow to exist. I shall mention how this is evidenced by the philosopher's hiring, teaching, peer reviewing, and writing practices, as well as by the code of conduct of philosophy talks.

2.1 Hiring Practices

A tolerance for disagreement in philosophy is reflected in the make-up of a department of philosophy. In the sciences, research at a department is often structured into large research groups. Those who want to work in such research groups are expected to have ample overlap in their convictions concerning the project; they need to be on board with it, so to speak. Why

would anyone join the project if they think it will fail? This would make no sense, especially since one's academic fortunes in the sciences seem conditioned to a considerable extent on the success of the project.

Yet, in the case of philosophy, there is no expectation at all that a job candidate shares the philosophical convictions of the other members of the department. Often enough, there is no large overlap at all in a department. Sure, there are research clusters, but these only determine the topic of research of the members of the department, such as self-knowledge or conceptual engineering. They do not indicate that everyone endorses the same views about such topics, and more often than not, they don't.

2.2 Teaching

Philosophers do not exclusively teach theories that they take to be true or theories they take to be on the right track. In scientific disciplines, it is common to teach theories that are deemed wrong or outdated only to explain the development of the contemporary theories that are taken to be correct. This makes it appealing to think of the progress of the discipline as steady and linear, even though this is usually a mere myth. This is nothing like how we teach the history of philosophy. As Catherine Elgin notes:

> We require our students to study Plato, even if we are convinced that the forms do not exist. We require them to study Spinoza even if we are confident that there's more than one thing in the world. We require them to study Kant even if we consider it obvious that transcendental idealism is a non-starter. You can continue the list, enumerating plainly untenable views that you were required to study, and those that you in turn will require your students to study. (Elgin 2022, 4)

We do not teach these theories as mistakes from the past, but we teach them to give a pedigree of our current philosophical disagreements. In fact, we still take ourselves to be in a continued conversation with these philosophers of the past. Yet, this conversation is one of sustained disagreement. Additionally, when we teach contemporary philosophy, we feel obliged to give equal weight to all views and not be swayed by which theory we, personally, take to be true.

2.3 Peer Reviewing

When philosophers serve as peer reviewers, it is common, and it is considered intellectually virtuous, for them to accept papers for publication even though they disagree with their conclusions. A peer reviewer is expected to judge whether the paper is internally coherent and to overlook the fact that it does not cohere with their own idiosyncratic philosophical beliefs (unless, of course, there is a particular reason for this incoherence that is especially interesting). In scientific cases, it is allowed to doubt the results of the paper under review simply because they contradict earlier well-established research. Therefore, reviewers can press the authors to account for this contradiction and explain why their research holds up. In philosophy, the mere fact that the conclusion of the paper contradicts that of an earlier paper in a well-placed journal is not enough to press the author to account for this. At most, one could give a friendly suggestion that the author might want to look into this paper because they might find interesting reasons for why they ended up with a contradictory conclusion.

The explanation for the scientist's reviewing decisions is fairly straightforward. Well-established research supposedly has a decent chance at getting to the truth. It is generally deemed rational to believe the results of a well-researched scientific paper. Therefore, it is *prima facie* rational not to believe that the conclusion of the paper under review is true if it contradicts these results. The reviewer could, of course, decide to accept the paper anyways,

without revisions, perhaps because they find the evidence for its correctness overwhelming, but the matter is not resolved until the research of either one of these papers has been shown to be misleading. Scientific inquiry has gone astray whenever there are various publications in well-established places that contradict each other, and nobody is trying to account for these contradictions. Yet, in the case of philosophy, this is a relatively frequent occurrence, and it is seen at most as an opportunity, not an obligation, to try to sort out some of these contradictions. This displays a higher tolerance for disagreement than in ordinary scientific cases.

2.4 Attending Talks

When philosophers attend talks, they don't just plan to sit and absorb information, to hear about the results of new research, but they ready themselves for disagreement. So does the speaker. Catherine Elgin has an amusing anecdote of when she replied to someone who disagreed with her by agreeing with him. She said: "He's right. His alternative is lovely. It is much better than ours. I wish we had thought of it." (Elgin 2022, 2) This comment silenced the room, and nobody knew how to proceed. It is clear that the audience did not come to this talk to see the speaker concede the point; they came to see the speaker perform all kinds of mental gymnastics to try to get their views to remain coherent in light of the criticisms offered, even if it is just for show.

The ultimate purpose of the Q&A is, I think, to help the speaker make their theory better insulated from counterarguments. We are thereby helping them find creative ways of sticking with their theory, of remaining steadfast in the face of disagreement.

2.5 Writing

As Elgin notes, philosophical papers often start with a review of the literature. Sometimes they show how their work builds on the work of others, like one could expect to see in a scientific paper, but more frequently, the literature review is to show with which views the author disagrees. Elgin says:

> A literature review in philosophy typically consists of rehearsing extant positions and explaining why they are inadequate. That done, we go on to present our own (which, naturally enough, does not suffer from the flaws we highlighted). The literature review sets the stage; it frames our discussion by explaining how and why we disagree with others who have worked on the same topic. (Elgin 2022, 3)

This is so pervasive that it would actually be strange to see a philosophical paper that does not take on and criticize an extant view. We might think that a paper is insubstantial if it does not defend a contrary philosophical position. We might even think it rude for the author not to make mention of the views with which they disagree. To a philosopher, it might be less offensive to claim that their work is mistaken than to not engage with their work at all.

Another notable difference with scientific papers is that those papers are very often cowritten, while the practice of cowriting papers is relatively rare in philosophy. This shows that philosophy is indeed less collaborative and more of an individual enterprise. Two philosophers must be very likeminded to agree on enough to write a paper together. The fact that we usually do not seem to have enough overlap with our colleagues is again a sign of widespread disagreement. It is also a sign that one's philosophical theories are deemed to be very personal. Perhaps this is why we identify with them, e.g., by saying "I am a pragmatist" or "I am a pluralist". This resonates with other personal inquiries (e.g., "I am an atheist" or "I am a socialist").

3. The Aim for Truth

I have shown that our philosophical practices display a high tolerance for disagreement, but this alone does not show that philosophical inquiry lacks the aim for truth. To make the case for this latter thesis, I must rule out the possibility that our tolerance for disagreement is in service of the aim for truth. Furthermore, I must show that the interpretation of philosophy as a personal inquiry can make better sense of our philosophical activities, given this high tolerance for disagreement, than any of the alternative views that try to stick to the idea that philosophical inquiry is shared.

3.1 Cognitive Diversity to Aid the Aim for Truth?

While I have emphasized the importance of convergence for our scientific practices, some level of cognitive diversity does seem valuable as well if we are seeking truth. We should not want to prematurely block certain roads of inquiry. So, even in the case of science, we should value disagreement. It would be too risky, given the pursuit of truth, to go all-in on the most promising theory. Zach Barnett shows why this is the case with the example of a doctor, Holly, who does research into how a certain illness is to be treated. It seems to Holly that drug Y would be best for treating the illness, but then she learns that the large majority of the medical profession believes that drug X would be more effective. Barnett says the following:

> [S]uppose that Holly does not see herself as any more likely to make accurate assessments of drug efficacy than the other doctors. Indeed, in the past, when there has been an absence of consensus among the doctors about which of two drugs is most efficacious, the larger group has tended to be right nine times out of ten. Given this track record, and given that a great majority of the doctors judged X to be more effective on this occasion, it would seem irresponsible for Holly (or any other doctor possessing the same evidence) to do anything other

> than administer X to an ailing patient. At the same time, it may well make good sense for Holly to investigate the efficacy of Y in her research. As a member of the research community, Holly should do whatever will aid the group in its efforts to determine conclusively which drug is most effective. Toward this end, it may not be optimal to have all fifty doctors devoting their research efforts to the same drug—even if that drug is currently the most promising one. Instead, it may be more efficient to have a majority of the doctors researching the most promising drug, with the rest researching alternatives that still have a decent chance of turning out to be the best. If we add that the doctors, as a rule, tend to produce their best research when they are permitted to investigate whichever drug sincerely seems to them to be the best, then it will make sense for Holly to research Y. (Barnett 2019, 126-127)

Barnett believes that we value disagreement and cognitive diversity in philosophy for the same reason.

However, the way in which diversity is valuable in the example of the doctor is nothing like the case of philosophy. We could imagine scientific inquiry, when successful, as a road to truth. The most promising theory will be the main road walked by most inquirers. Yet it is imaginable that this theory lets us down, that we encounter a roadblock. Therefore, it is useful to keep some adjacent roads open. Holly walks on one of these adjacent roads. Note, however, that it is neither possible nor useful to keep all possible roads open. For instance, the theory that the earth is flat should not be treated as a live possibility in any scientific inquiry. Many theories must be left behind, and only the most promising fringe views should be allowed to be developed further. Our expectation is that most fringe views will die off pretty quickly, while some might survive to become mainstream. In the case of science, we can reasonably expect that (1) many roads will be blocked off and, as a result, many inquirers will leave their own paths and return to the fold and (2) the adjacent roads do not frequently lead us in a completely different direction (such as to the belief that the earth is flat).

In the case of philosophy, there should be no such expectations. First of all, there often is no main road to begin with. It seems that more or less everyone takes their own path, paths which frequently head into completely opposite directions. Secondly, it happens very infrequently that a philosopher encounters a roadblock (i.e., conclusive arguments for why their views are a no-go) and has to desert their path.

Therefore, in the case of science, one could allow for some diversity of opinion to aid the aim for truth, but one can only do so given the expectation that these alternative views shall ultimately be shown to be conclusively better or worse than the mainstream view. In philosophy, however, such an expectation would be irrational since it happens only very infrequently that we develop conclusive reasons for or against any theory. Thus, there would not be an opportunity to close the roads of inquiry once opened. As a result, we all come to develop our own idiosyncratic theories, and there ceases to be a mainstream view at all.

One could argue that this is not actually a bad thing if one's aim is truth since the more views we develop the likelier it is that one of them is true. However, if our aim is indeed truth, then widespread disagreement undercuts our justification for believing our own theory. Barnett admits as much, arguing that we cannot rationally believe any of our philosophical views. (Barnett 2019) If nobody can justifiably believe their own theories, then the truth is still lost on us, for we could not know it, even if one of our theories does turn out to be true. The goal of science, or any other shared inquiry, is ultimately to *converge* to the truth.

David Lewis tries to use an argument like Barnett's to argue that the hiring practices of philosophy departments make sense given the aim for truth: we are trying to hedge our bets to give ourselves the best chance at uncovering the truth. (Lewis 2000) The above should suffice to show that this argument is unappealing. Instead, we could think that cognitive diversity is non-instrumentally valuable; we can appreciate the cognitive achievement of building a pretty coherent system of beliefs that radically departs from our own intuitions or convictions. Or, if

the value of cognitive diversity is instrumental after all, it is so because disagreement is a fruitful exercise for increasing the level of coherence of one's own views, regardless of the aim for truth.

The examples of philosophical practices, mentioned in the previous section, show that we are much more expecting and tolerant of disagreement in philosophy than in science. This indicates that the pressure to converge in our philosophical beliefs is lacking. Furthermore, it is unlikely that we exercise this kind of tolerance to disagreement with the aim for truth in mind. A likelier interpretation is that, for instance, when we hire philosophers, the question of whether their views are true does not even enter our minds. It neither does so when we assess the work of a peer. What we want to know instead is whether their views are internally coherent. This disregard for truth is a clear sign that philosophical inquiry does not aim for truth at all.

3.2 Steadfastness, Belief and Knowledge

If philosophy were to aim for truth, then the presence of widespread disagreement would be very problematic. If philosophy is carried out as a shared inquiry, then it would be generally rational to conciliate in the face of disagreement. Yet, given that there is so much disagreement in philosophy, we would have to suspend judgment all the time. There are various philosophers (e.g., Bernáth and Tőzsér 2021, Barnett 2019, Goldberg 2013, Kornblith 2013, Elgin 2022) who have argued from conciliation to the conclusion that it would be irrational for philosophers to believe their own theories, and therefore that there is no philosophical knowledge.

If we construe our philosophical practices as a shared inquiry, like scientific inquiry, we thus face the problem that it would be irrational to believe our own philosophical theories, and that philosophical knowledge is very likely out of reach. Why would we carry on doing philosophy at all if our prospects are that bleak? Some of the authors referred to above give an

answer to this question by arguing that there are doxastic attitudes of support for or endorsement of a philosophical theory that stop short of belief, but which we can rationally maintain.

Sanford Goldberg motivates this thought as follows:

> Unless we want to condemn philosophers to widespread unreasonableness (!), we must allow that their doxastic attitude towards contested propositions is, or at any rate can be, something other than that of belief. (Goldberg 2013, 282)

What would these doxastic attitudes be? According to Goldberg, it is attitudinal speculation:

> [O]ne who attitudinally speculates that *p* regards *p* as more likely than *not-p*, though also regards the total evidence as stopping short of warranting belief in *p*. (Goldberg 2013, 283)

The idea is that we can remain steadfast in the face of disagreement as long as we believe that our theories are likelier to be true than false. Barnett argues that even this won't explain why we can hold on to our theories because there are many cases in which there are so many competing theories that it would be rational for us to believe that there is a higher chance that one of these other theories is true than that our own theory is true. Barnett suggests that we could still hold onto our own theories by tweaking the interpretation of attitudinal speculation:

> *Attitudinal Speculation**: One who attitudinally speculates* that *p* regards *p* as the likeliest option (given some set of options), though also regards the total evidence as stopping short of warranting belief in *p.* (Barnett 2019, 115)

This way, we could keep endorsing our theory as long as we find it more plausible than any other theory. Elgin opts for something similar and says that the relevant doxastic attitude is not belief but acceptance:

> A theory or network of commitments is acceptable when it is at least as good as any available rival. The goodness in question is a matter of reflective equilibrium: the commitments that comprise a tenable network are reasonable in light of one another, hence mutually supportive; and the network as a whole is at least as reasonable as any available alternative in light of the antecedent commitments we deem relevant. (Elgin 2022, 7)

Barnett argues that the main difference between deciding which theory to *accept* and deciding which theory to *believe* is that, in the case of the former, we must insulate our reasoning from the higher-order evidence that our awareness of disagreement provides us with; we should endorse the philosophical theories that seem right to us when we put the disagreements aside. (Barnett 2019) The result of this "insulated reasoning" is that we are generally allowed to remain steadfast in the face of philosophical disagreement. Yet, we are only allowed to remain steadfast when it comes to the matter of which theory to support or defend, not when it comes to what to *believe*.

All of these responses are meant to show that we can be rationally entitled to keep endorsing our theories in the face of disagreement. Yet, all of these philosophers also argue that this could only be achieved by giving up on the premise that we could rationally *believe* our philosophical theories, and by thereby giving up on the possibility that we could come to *know* any philosophical truths. (Kornblith 2013)

This skeptical conclusion is widely accepted, but I want to resist it. I want to argue that we could be rationally entitled to believe our theories in the face of disagreement and, therefore,

that it is possible to have philosophical knowledge (which is not to say that it is easy to come by, though). Widespread disagreement need not prevent us from knowing philosophical truths, just like widespread disagreement does not prevent me from knowing that salty licorice is delicious. If philosophical inquiry is personal, then steadfastness could be an appropriate response to disagreement. In that case, widespread disagreement need not undercut our justification for our philosophical beliefs. This offers a much more appealing account of our philosophical practices than the skeptical account that tells us that none of us are ever justified in believing any philosophical theories at all. As section 6 shall show, once we reject that philosophy aims for truth, we can conceive of our philosophical inquiry as making steady progress, in the form of the accumulation of knowledge, even in the absence of convergence.

4. Scientific vs. Philosophical Inquiry

Before developing my own positive view about the nature of philosophical inquiry and philosophical progress, I first want to take a step back and assess the contentious premise that has led so many philosophers to the skeptical conclusion that nobody is justified in believing any philosophical theories, namely the premise that philosophical inquiry is a shared inquiry, and that it aims for truth. The above has shown that this view seems to conflict with how we carry out our philosophical practices, so why do we not question it? I think this is largely due to the popularity of the naturalist belief that philosophical inquiry is, or should be, like scientific inquiry. Therefore, a productive way of arguing that philosophical inquiry is personal would be to argue that philosophy *should not be* scientific. This is the aim of the current section.

4.1 Naturalism

In spite of all the differences between philosophical and scientific inquiry mentioned above, one could think of science and philosophy as having a shared aim: to explain the nature of reality. The most prominent difference between the two disciplines, it could be argued, is their level of abstraction — though this would not be a very neat distinction since, for instance, mathematics and theoretical physics can be incredibly abstract. Perhaps there are those that would even be willing to say that theoretical physics at the highest level of abstraction is indistinguishable from philosophy. Given this high level of abstraction, philosophical inquiry is highly speculative, which goes some way to explain why there is so much disagreement in philosophy.

This view, that philosophical inquiry is like scientific inquiry, is called methodological naturalism. It has a normative dimension as well: philosophical practices *should* emulate scientific practices. Philosophical inquiry should be as rigorous and cooperative as scientific inquiry. This would give us the best shot at converging toward the truth. I take it that some version of this view is presupposed by many, if not most, analytic philosophers.

However, the idea that philosophy is about inquiring into the fundamental nature of reality completely overlooks practical philosophy. The primary purpose of practical philosophy is not to inquire into reality but to inquire into how we ought to live our lives. I have already argued that moral inquiry is unscientific, and that it is personal. It is no wonder, then, that, for instance, the logical positivists, who were proponents of the idea that philosophical inquiry should be scientific, tried to find various ways of disregarding ethics, e.g., by arguing that it is not truth-apt. (Ayer 1936) Naturalism provides a *prima facie* reason to try to push practical philosophy to the periphery of philosophical inquiry.

This is exemplified by the emphasis on logic. Logic serves as a kind of ideal for philosophical reasoning. Suppose one believes that there is One True logical system that

describes the most general rules that all of our reasoning has to abide by. In that case, we can think of the function of logic as breaking down intricate problems into simple deductions for which the assumption of *sameness* is rational, just like one does in the case of science, using simple observations and computations. Philosophy should then count as a shared inquiry. This is not to say that philosophical debates *about* logic — such as whether intuitionistic logic is better than classical logic — would count as exemplars of the scientific approach to philosophy as well. Nevertheless, there are philosophers, such as Timothy Williamson (2013), who pick their philosophical views based on whether they cohere with their preferred logical system, and who take this to be a successful implementation of scientific rigor into philosophical inquiry.

4.2 Waiting Around for Convergence

We could say that logical deduction is at the one end of the spectrum of philosophical inquiry, and moral inquiry is at the other. What they have in common is that we frequently engage with both outside of philosophical contexts. Yet, if logical deduction is meant to serve as an ideal for our reasoning, our moral reasoning falls far short of it. It is a very messy business that relies heavily on intuitions, which we try to bring into reflective equilibrium. I think this model of reasoning, rather than the logical reasoning model, is a truer reflection of our everyday philosophical practices.

Even in the case of Williamson's philosophy, this messy pick and choose is evident: he is willing to accept rather contentious theories, such that everything exists necessarily (Williamson 2013) or that there is a single grain of sand that turns the heap into a heap (Williamson 1996) because he is committed to standard logic. Yet not everyone feels so strongly about standard logic that they are willing to accept these counterintuitive results. There are various packages of theoretical commitments, none of which is optimally coherent or

intuitive. However, as Williamson would argue, scientific practices are just like this; they are often just as messy and require a kind of reflective equilibrium of theoretical commitments. They do not neatly break down into easy observations or computations. What is the difference?

In both cases, which of the rival views one supports is essentially a value judgment. There isn't one view that fits the evidence better, or that is much more elegant or simple. It rather has to do with which theoretical baggage one is willing to burden. Therefore, it does become like an instance of disagreement that is, in some ways, a faultless disagreement of taste: you like this theory better, I prefer that one. However, the difference is whether it is rational or not to believe that the continuation of the inquiry might resolve the disagreement. In the case of science, the answer is usually yes, whereas in philosophy, it is usually no. It is to be expected that even in the case of scientific inquiry there shall be many instances of disagreement that are messy and that boil down to a matter of taste. Yet, the question is whether we take such an instance of disagreement as ineliminable due to the nature of the inquiry, or whether we hold out hope that it is due to the inadequacy of the current evidence and that our future inquiry is likely to dissolve the disagreement.

It seems clear to me that we tend to believe the latter in the case of science. This is not because we expect that there would be one vital experiment that would prove the truth of one theory beyond a shadow of a doubt. We know that the interpretation of experiments can itself also be a messy business. However, we do expect that one theory will be the aesthetic favorite; it will generally be seen as offering the more elegant and simple account. The other theory will then slowly be pushed out. In the case of philosophical disagreement, however, we usually do not expect any significant changes in our evidence; there would be no new experimental discoveries that would have a bearing on the disagreement. Furthermore, we also lack the mechanisms for pushing out any theories. Therefore, even if a theory loses support over time, the situation is not unstable like it would be in the scientific case. An example of this might be

mind/body dualism, which, once upon a time, was very popular, then it suffered a heavy blow, but ever since, contemporary versions of it have received steady support from a large minority of the philosophical community. (Bourget & Chalmers ms.) Given the unlikeliness that our philosophical disagreements will be resolved by some new piece of evidence, it would be rational to presume that these disagreements are not due to a lack of evidence but that they are due to the nature of philosophical inquiry.

While philosophical inquiry is highly disciplined, like scientific inquiry, and unlike other personal inquiries, its standards stop far short of ensuring that there is only one rational belief or set of beliefs in the face of disagreement. This is also the case for scientific inquiry. However, I have shown that in the case of scientific inquiry, it is generally rational to assume that further evidence might significantly alter the state of the disagreement. This is why it is generally rational suspend judgment in the case of science, whereas this is not the case for philosophical inquiry. In the case of philosophy, if we had to wait around for new evidence that could lead to convergence, we would have to wait a rather long time. Therefore, while the mythology of there being at most one right view to be had is functional in the case of scientific disagreement, because it creates the impetus to conciliate, it breaks down in the case of philosophical disagreement. It is not useful to treat philosophy as a shared inquiry, which is why we ought to take philosophical inquiry to be personal.

This might explain why so often a philosopher's response to disagreement is not conciliation but rather entrenchment. Of course, if one thinks that the other theory is actually better than their own, they should change sides, but if they find their opponent's arguments unconvincing, they should not be bothered by the disagreement. Likewise, we should not cease to support our theory even if 70% of the community disagrees with us, as long as we find their arguments unmoving. In that case, we could appreciate their arguments and understand their point of view, but ultimately think of it as a matter of taste, and of which theoretical

commitments one is willing to burden. This could be seen as nothing more than a personal inclination.

4.3 Just a Matter of Taste?

Given that we could interpret *some* philosophical disagreements as merely a difference of personal taste, what is to stop us from treating all disagreements as such? We want to say that our theories must abide by standards that are independent of truth. Simon Blackburn has argued, for instance, that inquiry that does not aim for truth could still be guided by the following standards: "information, sensitivity, maturity, imagination, coherence." (Blackburn 1998, 318) I have shown in the previous chapter that we can understand the aim for rationality — which, I take it, also involves the aim for coherence — independently of truth. However, the question could still be asked why, if we do not aim for truth, we should pay any heed to such standards at all. Why couldn't I just endorse philosophical beliefs haphazardly, based only on my personal philosophical tastes?

This might be more common than we think. I think people often treat their philosophical views as mere expressions of personal taste or inclination. Everyone has their "own philosophy", so to speak. One's worldview does not necessarily *need* to be coherent. However, this concerns philosophy in the broadest sense, not only as an academic discipline but also as encompassing any beliefs people might have that are remotely philosophical. Everyone must have *some* philosophical views about life and reality. Yet not everyone is a philosopher. Philosophy, in the stricter sense, as an academic discipline, could not merely be about expressing one's own personal sentiments.

A common way to get into this thing we call "philosophy" is through philosophical dilemmas or paradoxes. We could imagine a high school teacher presenting their students with

the trolley problem or the lottery paradox. Many students might find these problems irrelevant. Some might even struggle to conceive of them as problems at all. Then there are, hopefully, also a few who find such problems fascinating. They might go on to become philosophers. They are the kind of folk that aims for coherence since they believe it problematic to have various independently appealing philosophical beliefs that do not add up — such as the belief that we should prioritize the wellbeing of the many over the few (in the trolley case), or that it is rational to accept beliefs that are very likely to be true (in the lottery case). The activity of a philosopher is not to spew philosophical beliefs but to argue for them, which would be the activity of trying to make them add up.

A comparison with art might be useful. It is fairly typical to think of, for example, the activity of painting or sculpting as a purely personal expression of aesthetic sentiments. However, art is also a conversation. The only way to adjudicate between good and bad art within this conversation is by appealing to aesthetic principles. It is no use to simply say: I like this, I don't like that. Contemporary academic philosophy is a bit like the *Académie des Beaux-Arts* in France. Back in the 19th century, the *Académie* had a monopoly on the high arts. It promoted realist and naturalist art that adhered to an extremely rigorous set of aesthetic principles. Perhaps unsurprisingly, the academy was also very hierarchal and cliquey. For the members of the *Académie*, the suggestion that art is merely an expression of personal sentiments must have seemed very far removed from their activities. Likewise, for those belonging to academic philosophy, the suggestion that philosophy is merely an expression of personal sentiments seems outlandish as well.

Philosophy may be a personal inquiry, but it is nevertheless a shared activity. It is a conversation. As this chapter has shown thus far, a central task of philosophers is to disagree. We can only meaningfully disagree as long as we agree on a set of rules that determine what counts as a good argument and what counts as a bad argument. Philosophy is, as an activity,

encapsulated by the game of giving and asking for reasons. To do philosophy is to join in on a conversation, and if we don't do that then nobody will listen. If we disregard the norms and principles that guide our conversation, and we think of our philosophical views as purely personal expressions of philosophical sentiment, then nobody could meaningfully engage with these views.

This doesn't mean that our inquiry should be shared and that there are universal rules as to how to do philosophy well. Just like in art, the principles of philosophy are constantly evolving and under scrutiny. There is no neat boundary to be drawn between the methodology and the activity itself. An artwork could be both a piece that adheres to aesthetic principles as well as a piece that reflects on these principles. A metaphilosophical text is also a philosophical text, so it could put forth principles of how to carry out philosophical inquiry, but it is itself also subject to such principles. A methodology for philosophy cannot be formulated in the abstract, without any commitments to substantial philosophical views. Just like art, then, in philosophy, the matters of style and method are themselves constantly under scrutiny. There are many different schools and isms in which philosophical views and metaphilosophical principles run together.

As mentioned in the previous chapter, this explains how we could think of some philosophical views as mistaken yet faultless because we take them to be internally coherent. This is just like disagreeing with cubist art in general but still being able to judge whether an artwork belonging to this movement is good or bad, according to the principles of the movement. The admission that multiple philosophical styles or methods can be rational does not mean that all philosophical views are admissible, that all our philosophical disagreements are faultless, or that there are no standards at all in philosophy.

It also does not mean that philosophy is just a game, and that its rules are entirely arbitrary. Even though our following of the rules of rationality may not be in service of the

pursuit of truth, it does have practical bearings. Academic philosophy can seem a bit detached from our practices, but there have been many worldviews that ultimately influence every aspect of our lives. Philosophy is a part of the battle of great ideas that have shaped societies. Philosophy is guided by the conviction that we can change things using *reason* rather than swords, guns, or killer drones. Even if this exercise of reason turns out to merely be a game with an arbitrary set of rules, it is still a game that requires less bloodshed than war. It would be better to be a pawn in a game of chess than in the game of war. That being said, while we could try to change another's worldview using philosophical arguments, by trying to poke holes in their views, this is, as I have argued, not a recipe for convergence. For the sake of convergence, it might be more effective to try to kill our opponents than to engage in a philosophical debate with them (but it would, of course, be much better to exercise tolerance toward different worldviews).

The standards mentioned above are also practically relevant in a more direct sense. We follow these standards because we aim to do what works. Incoherent beliefs generally lead to disappointment; uninformative beliefs fail to sufficiently guide us; insensitive beliefs might not enable us to cope with changing circumstances; unimaginative beliefs risk overlooking better alternatives, etc.

The worry that this subsection started with was roughly as follows: philosophical inquiry shares many theoretical virtues with scientific inquiry, yet in the case of scientific inquiry, these standards seem motivated by the aim for truth, so what motivates them in the case of philosophical inquiry? What legitimacy do these standards have without the aim for truth? Now we see that there are two parts to this answer. The first part is that *some* standards are needed to facilitate the game of giving and asking for reasons. The second part is that our standards are not just arbitrary rules because they are motivated by our aim to do what is best.

This does not yet explain why we have many of the *same* theoretical virtues for scientific and personal inquiry. Is this merely a coincidence or merely due to the historical connections between the two disciplines? I do not think that it is. We have assumed that the theoretical virtues of science are motivated by the pursuit of truth, but is this really the case? It seems that the only indication we can have of our scientific beliefs being true is whether they work in practice. Therefore, what motivates the scientific virtues is, first and foremost, the promise of practical efficacy, and perhaps only indirectly, or not at all, the promise of truth. Given that truth is not analyzed in terms of utility, it requires a leap of faith to believe that the theory that works best is the one that is true. As argued in this thesis, the practical role of the aim for truth is that it gets us to try to converge toward the same beliefs. The aim for truth might have nothing to do with motivating us to follow the scientific virtues, which we seem to follow instead for their practical efficacy, just like in the case of philosophy. Therefore, it is not due to coincidence or historical contingencies that science and philosophy share many theoretical virtues, but it is because both inquiries are ultimately guided by practical efficacy.

4.4 The Self-Refutation of Conciliation

One last thing worth mentioning is that, given the link between personal inquiry and steadfastness, and shared inquiry and conciliation, it would actually be rather difficult to consistently believe that, in spite of appearances, philosophical inquiry is shared. There is a problem of self-refutation that troubles any global conciliatory view: if we believe that we are always rationally obliged to conciliate in the face of peer disagreement, we must likewise believe that we are rationally obliged to conciliate when it comes to the conciliatory view itself (given that many skilled philosophers support steadfastness). As a result, the supporter of the conciliatory view is not rationally allowed to be confident in their own view, whereas the

supporter of the steadfast view is. (Weatherson 2013) Yet this issue is avoided altogether if one supports the conciliatory view only locally, for instance, for scientific disagreements. If philosophical inquiry is personal, then steadfastness could be a rationally permissible response to the disagreement about whether the steadfast view or the conciliatory view is correct. Therefore, in addition to there being good reasons for believing that philosophical inquiry is personal, it would also be difficult to confidently believe otherwise.

5. The Nature of Philosophical Inquiry

It is important to stress that philosophical inquiry is not bound by method or subject matter. To give a theory of philosophical inquiry is to presuppose that it is uniform enough to be captured by a single theory. It probably isn't. I must thus be more careful when I say that philosophical inquiry is personal, especially since the boundary between philosophical and scientific inquiry is a fuzzy one. Surely there are many philosophical projects that are scientific and, thus, shared.

As Philip Kitcher points out, science used to be a part of philosophical inquiry. Yet, this was before the modern scientific method. Kitcher argues that though many questions that preoccupied the minds of pre-Socratic philosophers were factual questions about the nature of reality, we have since come to realize that such questions are best answered not by *a priori* reflections but by extensive experimentation. (Kitcher 2011) We found the philosopher ill-equipped to engage with the project of describing nature, which is why the inquiry branched off into different scientific disciplines.

While we still have plenty of philosophers who venture into science and scientists who venture into philosophy, I have shown it to be mistaken to think of philosophical inquiry in general as modeled after scientific inquiry (or vice versa, for that matter). I argued above that if we think of philosophy as an inquiry into the nature of reality, we tend to push areas of

practical philosophy to the periphery of what we take the essence of philosophical inquiry to be. Perhaps we should argue for the inverse: the core of philosophical inquiry is practical philosophy — the chief concern of philosophical inquiry is to find out how we ought to live our lives.

This would help explain the personal character of the inquiry since we do not, or at any rate should not, presuppose that there is only one correct way of living one's life. However, this conception of philosophical inquiry seems at odds with most areas of theoretical philosophy. A more encompassing characterization of philosophy, which I already briefly made use of in the previous section, is that it is about developing one's "worldview".

Kant first used this term, *Weltanschauung*, to denote the reflective judgment of nature as a harmonious whole. Kant believes that this judgment "constructs" a world in which the laws of nature and our moral practices are united. The main standard of such a judgment is not predictability or truth (like in the case of scientific judgments) but *coherence*. And, importantly, Kant does not conceive of a *Weltanschauung* as a factual account of reality but as a normative account. It is, in Alexander Englert's terms, "a value-judgment about existence itself." (Englert 2022, 9)

The task of philosophy is thus, according to Kant, to develop an account of how we ought to think about reality. Englert puts it as follows: "As a technical term, worldview connotes an active process of orienting oneself relative to the world." (Englert 2022, 9)

I would be remiss if I did not point out that Kant believed that one's worldview is composed of *a priori* necessary judgments and that, therefore, everyone who reasons well shall adopt the same worldview. This is basically an expression of the assumption of sameness, which ensures that the inquiry is shared. However, without this controversial assumption, his idea that philosophical inquiry is not about reality *per se* but about how we ought to think about

reality, and that it aims not for truth but for coherence, can give us some idea as to what philosophical inquiry as personal inquiry amounts to.

This conception of the nature of philosophy works well with my overall framework. Suppose philosophical inquiry is indeed about how we *ought* to conceive of the world. This would vindicate my methodology of arguing for the distinction between personal and shared inquiry not on ontological but on practical grounds: some of our practices are better served by carrying out the relevant inquiry as a shared inquiry, which is why we *ought* to conceive of them as shared, whereas this utility is lacking in the case of other practices. In arguing for the philosophical thesis that certain inquiries are shared whereas others are personal, I argue that it pays to conceive of our inquiries as such — it is a useful distinction — and that is why we *ought* to take them to be this way.

Further important philosophical distinctions follow from this, such as distinctions in truth, rationality, knowledge, and a realism/anti-realism distinction, but none of these further distinctions are meant to give an epistemological or ontological justification for the distinction between shared and personal inquiry. On the contrary, they are themselves *explained by* this distinction, e.g., science aims for truth whereas philosophy does not *because* we ought to take science to be aimed for truth, whereas we shouldn't treat philosophy as having this aim.

Furthermore, it would be strange if I offered my thesis that philosophical inquiry does not aim for truth merely as a description of philosophical inquiry. Many philosophers would disagree with such a thesis. Perhaps if we asked philosophers whether they take themselves to aim for truth in their inquiry, most would say yes. Thus, if my thesis offered a mere description of what philosophical inquiry is like, it would not be appealing at all.

It is imaginable that a philosopher, persuaded by my arguments, would say: "I thought I was aiming for truth, but now I realize that this was never truly my aim". However, this is not to be interpreted as a simple mistake, such as the one made by someone who says: "I thought I

was heading for Scotland, but now I realize that I have boarded a flight to Spain." In this latter case, one thinks back on their activities — driving to the airport, getting to the terminal, looking at the screen to find their gate, etc. — and they are required to believe that some of these activities have been executed incorrectly. However, the philosopher is not required to believe that they have been going about their practice in the wrong way. Quite the opposite, in fact. They might realize that how they have been going about their practice, especially on the social level, is best interpreted by taking themselves to aim not for truth, but toward something other than truth, such as coherence, certainty, knowledge, or understanding. It is not how our inquiry is carried out that is at issue, but only how best to conceive of how it is carried out; we *ought* to conceive of it as lacking the aim for truth. Other philosophers are free to disagree with this, but they have misunderstood the nature of the thesis if they think they can refute it by taking a poll on how many philosophers are "actually" aiming for truth in their philosophical activities.

5.1 Norms and Facts

If the character of philosophical inquiry is normative rather than descriptive, something needs to be said about the nature of this normativity and about this distinction in general. Does the distinction work for other personal inquiries? In the case of basic taste or aesthetics, it would be more proper to speak of "evaluative" rather than "normative". Nevertheless, the idea that personal inquiry deals in value judgments rather than descriptions is a familiar one that holds up pretty well. However, I shall argue that the distinction is too crude.

Suppose that philosophical talk is normative rather than descriptive. What kind of norms would it express? I am partial to a pragmatist interpretation: when we express that the world *ought* to be one way rather than another, we are implying that it would be more *useful* to

take the world to be one way rather than another. This usefulness would concern a general utility: our practices are best served, in general, by thinking of the world in this way.

Yet, while I am partial to a pragmatist interpretation of the normativity of our philosophical talk, other interpretations are possible as well. We need not assume, for instance, that the normativity of ethics can directly be reduced to a kind of practical utility. Also, one might worry that the pragmatist interpretation would be hostile toward metaphysical discourse, which is about matters that are notoriously insulated from any practical effects. It is worth focusing on metaphysical discourse in particular because this is the kind of discourse that seems most straightforwardly to be descriptive of reality. Therefore, if the thesis that most of philosophical inquiry is not descriptive of reality has any merit, it must be possible to show how it could incorporate at least some of our metaphysical inquiries.

The interpretation of metaphysics as an expression of a practical attitude rather than a description of states of affairs has been defended, for instance, by Rudolf Carnap:

> [M]etaphysics does indeed have a content; only it is not theoretical content. The (pseudo)statements of metaphysics do not serve for the *description of states of affairs* […] They serve for the *expression of the general attitude of a person towards life*. (Carnap 1959, 78)

A caveat here is that Carnap is a methodological naturalist, and he does think that philosophy should describe reality, which is why he, at times, treats most of metaphysics as an inquiry we should not engage with.

A more constructive, yet still undeniably Carnapian, interpretation of a part of metaphysical discourse has recently been developed by Amie Thomasson in *Norms and Necessity* (2020). In this book, Thomasson takes on claims concerning *metaphysical modality* and argues for a "normativist" as opposed to a descriptivist approach.

Not all metaphysical claims are about possibility or necessity, especially given the recent turn toward grounding and fundamentality. (Schaffer 2009) However, it still holds true for many metaphysical inquiries that they are not about what reality is like but about what it *could* or *must* be like. Thomasson interprets these metaphysical *coulds* and *musts* as normative.

She argues that a descriptivist account of metaphysical modality leads to significant ontological and epistemological problems; the idea that our modal claims are somehow tracking and being made true by modal states of affairs is very obscure. Thomasson argues instead that statements about what is metaphysically necessary, possible, or impossible are concerned with *semantic* normativity. According to Thomasson, when we make such claims, we give an object-level expression of a semantic rule, specifically a rule concerning the application of a concept.

Take the issue of whether the classical vocal rendition of *Autumn in New York* by Billie Holiday and the contemporary instrumental interpretation by Makaya McCraven are renditions of the same song. One side might argue that it is the same song because it has the same melodic line, whereas the other side could argue that McCraven's version has strayed so far from the composition that it is no longer recognizable as the same song. This is a metaphysical inquiry: has McCraven destroyed the song by altering the key, tempo, removing the lyrics, and adding extensive improvisations and new melodies, or could the song survive these alterations? Yet, in a different guise, it could also be a meta-semantic debate: how should we apply the concept *song*?

Some might say that the meta-semantic level reveals that this is merely a verbal dispute instead of a substantive metaphysical disagreement. However, there could still be a substantive disagreement about how a concept is to be applied. The intuition that the disagreement about whether the two songs are the same is merely a verbal dispute is driven by the fact that this seems like a fairly low-stakes issue; what do we care how the concept *song* is applied? Yet, as

Thomasson also notes, David Plunkett and Timothy Sundell have argued that disagreements about the application of concepts could be very substantive. Their example is the disagreement on whether waterboarding is torture. This is either a metaphysical disagreement about the nature of waterboarding and the nature of torture, or a meta-level disagreement on whether our concept of torture is to be applied to waterboarding. Nevertheless, this disagreement does not concern a trifling matter in either of its guises. (Plunkett & Sundell 2013) Another example would be whether two people of the same sex or gender could marry each other. This is a political question, but it is also partly a metaphysical question about the nature of marriage, and a meta-level semantic question about the concept of marriage.

These are all thoroughly normative questions, but I take all of them to ultimately be involved with a *pragmatic* ought. For instance, one might say "Let's just introduce a different concept for the union between people of the same sex or gender", but this distinction would be *impractical* given that we are heading toward a queer future. It would be a mistake like the one entailed by the concept *grue*. The conceptual structure of which a same-sex exclusionist concept of marriage is a part is one that does a very poor job of describing our social reality. This conceptual structure is outdated and obsolete.

Along this same line, Thomasson explains how her interpretation of metaphysical modality still allows for deep metaphysical disagreements:

> Such disputes, so conceived, are onto something deep — they are not concerned with how our parochial terms or concepts happen to work. Instead, they are about how we should live. This enables us at last to do justice to the feeling of "depth" in these debates — that we are not just aiming to find out about our concepts of persons (when we investigate personal identity), or freedom (when we enter the free will debate), or even works of art — but to find out something deeper, beyond parochial investigation of whatever our own concepts happen to be. But on this conception, what it is that we are arguing about is best understood normatively —in ways that

> are intimately tied up with large normative questions about how we should arrange our moral, legal, and other institutions and practices [...] In this way, the modal normativist view can lead to a kind of metaphysical normativism — that sees some of the most important work in metaphysics as involved in a normative project of working out what conceptual scheme we *should* be working with. (Thomasson 2020, 204)

This is very much in line with the general idea that philosophy is about figuring out how we ought to conceive of the world. How best to apply our concepts is a part of this project.

It could be argued that there is somewhat of a normative turn happening in analytic philosophy that exemplifies this methodology. As Thomasson argues, the principles of conceptual engineering epitomize this idea. (Cappelen 2018) Further examples are Sally Haslanger's ameliorative project concerning the concepts of race and gender (Haslanger 2012), as well as the shift in epistemology from the analysis of knowledge to a focus on how epistemic concepts are used from a social perspective, thereby introducing normative, and even explicitly moral (e.g., Fricker 2007), elements to the discipline.

The general concern is not with how we actually apply our concepts, nor is it conducive to think that there is one way of applying our concepts that is *objectively correct*, e.g., that carves nature at the joints. Instead, there are only divergent opinions on how we *ought* to apply our concepts. These disagreements follow the model of personal inquiry.

While I want to take many aspects of the normativist view on board, I ultimately do not think that the distinction between normative and descriptive can capture the distinction between personal and shared inquiry. First of all, the distinction between normative and descriptive must break down, in one direction at least, if philosophical talk is interpreted as normative, for it would be absurd to say that it is *merely* normative and not descriptive as well. This would be like defending a claim such as "Personal identity can persist over time", and when someone asks if this is actually the case, we would say "Well, no. But it *ought* to be the case." One would

rightfully respond by saying: “Who cares?! I am not interested in your opinion on what things ought to be like; I want to know what things are actually like!” If philosophical inquiry cannot even attempt to describe what is the case, what is the point? Therefore, we must maintain that the matter of how we ought to think about personal identity and the matter of what personal identity is actually like are one and the same. Philosophical talk is both normative as well as descriptive.

Yet, it would also be unwarranted to think of shared inquiry like science as purely descriptive. For instance, descriptivism about mathematics poses some of the same ontological and epistemological problems as the ones that motivated Thomasson’s normativism about metaphysical modality; it seems obscure to think of our mathematical claims as tracking and as being made true by an abstract realm. As Wittgenstein argues, mathematical propositions are better understood as ways of getting from one nonmathematical proposition to another. They may aid us in describing reality, but they aren’t descriptive of reality themselves. (Wittgenstein 1921, 6.211) Also, mathematical reasoning could neatly be interpreted as reasoning from a set of rules. Wittgenstein argues that it makes as much sense to think of these rules as *discovered* rather than constructed, as it does to think of the rules of chess as discovered. (Wittgenstein 1991, 374) Therefore, we could interpret claims like “2 + 2 = 4” as an expression and application of a rule rather than as a description of reality, just like we would interpret “Pawns cannot move backward” as an expression of a rule rather than a description of reality. If one is moved by the problems of descriptivism about metaphysical modality and finds normativism an appealing alternative, then one would probably have this same response to the case of mathematics.

Thomasson’s “normativist” conception of necessity is not entirely new. It has been advanced by, for instance, Robert Brandom, who has argued that modal claims are expressions

of inferences that conceptual contents allow or require us to make. (1994, 2008) Brandom himself traces it all the way back to Leibniz and Kant:

> Kant takes over from his reading of Leibniz the general idea of rules as what underwrite cognitive assessments of inferences and judgments. He understands such a priori principles, however, not as very general statements of fact (even metaphysical fact), but as rules of reasoning. They are conceived not as descriptive but as prescriptive — as (in Sellars's phrase) "fraught with ought." (Brandom 1994, 10)

However, Brandom, as well as Leibniz and Kant, apply this thought much more broadly to all kinds of necessity. Therefore, they are also normativists about nomological necessity and, thus, about causality. This is how Brandom's quote continues:

> This lesson dovetails neatly with the moral [Kant] draws from Hume's thought. [...] Kant's Hume recognizes that cognitive experience crucially involves the application and assessment of the correctness of the application of *rules*. For Kant, Hume's inquiry after the nature of the authority for this inferential extension takes the form of a quest for the nature of the necessity, understood as normative bindingness, exhibited by the rules implicit in empirical concepts. It is under this conception that Kant can assimilate Hume's point about the distinction between saying what *happens* (describing a regularity) and saying what is causally *necessary* (prescribing a rule) to his point about the distinction between saying what is and saying what ought to be. (Brandom 1994, 10-11)

Kant thus believes, according to Brandom at least, that all talk about modality is prescriptive rather than descriptive. It is expressive of the rules of reasoning. Thomasson herself also seems to believe this and recognizes the appeal of interpreting all modal talk as normative rather than

descriptive. (Thomasson 2020, 35) This Humean kind of normativism about causality has also been defended by Frank Ramsey. Ramsey argues that causal laws should not be thought of as infinite conjunctions that describe reality, but instead, he says that: "causal laws form the system with which the speaker meets the future". (Ramsey 1929a, 241) Ramsey believes that causal laws are prescriptive claims as to how best to relate ourselves to the world; they are not matter-of-factual claims. If this is the case, then scientific theorizing might be as much about expressing how we *ought* to think about reality (instead of describing reality) as philosophy is. The idea that the former is descriptive whereas the latter is not — that the former deals in facts and the latter in norms — breaks down. What, then, is the difference?

5.2 Objectivity

I have argued that much of philosophical inquiry is normative; it isn't about reality directly but rather about how we ought to conceive of reality. However, I have also shown that parts of scientific inquiry, such as mathematics, could be seen as similarly normative, which is why the distinction between norms and facts does not mirror the distinction between personal and shared inquiry. I shall argue now that there is a difference in the kind of normativity involved in personal and shared inquiry. This difference can explain why shared inquiry could be said to be realist.

In my discussion of peer disagreement in Chapter 3, I tried to uncover what one rationally ought to believe in the face of disagreement. A distorting factor that I ignored then is that there is only one thing one really ought to believe, which is the truth. I ignored this kind of *objective ought* because it does not give us a rule that we could follow in virtue of which we could be said to be rational. The rule of rationality must instead be something to the effect that we should believe what *seems* correct given the available evidence — i.e., that we should have

beliefs that are warranted by the available evidence — and disbelieve what *seems* incorrect, not that we believe what *actually is* correct and disbelieve what *actually is* incorrect. What determines the rationality of our doxastic attitudes is a *subjective ought*.

I have borrowed this distinction between subjective and objective ought from Allan Gibbard, who gives the example of flipping a coin to show this difference: suppose that the coin will land on heads, then it would be correct for you to believe that it will land on heads, so this is what you ought to believe, in the objective sense. Yet, what you subjectively ought to do is give equal credence to the belief that it will land on heads and the belief that it will land on tails. (Gibbard 2005, 340)

The objective ought is problematic because for it to be normative at all, it must give us a rule that we can meaningfully try to satisfy, not just accidentally, like by guessing heads. Nevertheless, we often make use of objective oughts. Boghossian gives an example of an everyday ought that seems out of reach: "You ought to buy low and sell high." (Boghossian 2005, 211) If only we knew how to follow this rule.

Gibbard argues that objective normativity is based on subjective normativity. It is the kind of thing I ought to do in the counterfactual case in which I know all the relevant information. Well, not quite, because, as Gibbard notes, there could be things that I objectively ought to do but that I would not need to do in the case in which I already possess all the relevant information. His example is that we could ask whether I objectively ought to go behind enemy lines to figure out the strengths of the enemy. It seems that if I am in this counterfactual position in which I can know that I won't get caught, and that doing this would give me certain very useful information, then I would judge that I objectively ought to do this. However, in this case, I would already know this information that I would be trying to get, so there would be no reason for me to go behind the enemy lines and get it.

Gibbard's response is to claim that the counterfactual me, who functions as the ideal judge, could not really be me at all, but they must be somebody who is distinct from me. Gibbard calls this judge "I^+". I^+ would have all the information but would also know of my ignorance and would therefore judge that if they were in my situation, in which they would not have all the information, they objectively ought to go behind enemy lines.

Gibbard then shows that what I *objectively* ought to believe is really what I^+ *subjectively* ought to believe were they in my circumstances. Yet, the only things I^+ would subjectively recommend me to believe are the things that actually obtain:

> I $ought_{ob}$ to accept that S iff I^+ ought to accept that S for my circumstance.
>
> I^+ ought to accept that S for my circumstance iff S obtains in my circumstance.
>
> (Gibbard 2005, 349)

This is how we get to the objective doxastic ought, which states that we ought to only have beliefs that are true (perhaps with the exception of a few useful falsehoods).

Yet, we do not need to rely on this mythical creature "I^+" to make this point. The difference between the subjective ought and the objective ought can be formulated socially. Suppose you think that your epistemic context is better than mine, and you believe *p*. In that case, you might judge that, on the one hand, I ought to believe *p* because it is true, but you realize that, on the other hand, this belief is not warranted by my epistemic context. Thus, all that is needed to make a distinction between the objective and subjective ought is a third-personal perspective. From a first-personal perspective, you could not make this same distinction, and you could only judge that you subjectively ought to believe *p* because it is warranted by your own epistemic context. Therefore, an objective ought is really our own subjective ought that we project onto people whose epistemic contexts we deem to be worse.

Of course, that hardly sounds objective. It is really a kind of indexical ought: I project what I ought to believe myself onto you.

In certain personal inquiries, such as the inquiry into matters of basic taste, we realize that our oughts only apply to ourselves, and we are reluctant to project them onto others. Therefore, it could be said that there is no *projected* ought at all in such inquiries. This is not the case for all personal inquiries, though. As discussed in Chapter 3, it is part and parcel of moral judgments that we judge that others ought to agree with us. However, this is a moral ought, not a doxastic ought; we do not usually think others ought to agree with us because our epistemic contexts are superior, but we think they ought to do so because having immoral beliefs is itself a moral failure. In philosophy, we might project our oughts and think, for instance, that the other party lacks crucial information about the view they are defending. In that case, we can deem our epistemic context to be superior and project our subjective ought onto them. Yet, it is not at all required that we always do so. We could realize that we just have different background commitments, which is why what I ought to believe may come apart from what you ought to believe. Therefore, in philosophy, one could judge that one ought to believe *p* themselves while judging, at the same time, that other philosophers are not required to believe *p*. For instance, one could say "As a consequentialist, I ought to believe *p*" and believe that those who aren't consequentialists do not have to believe *p*.

This is not the case for science. As mentioned in Chapter 4, given the push toward convergence by the truth norm, we believe that everyone must try to adopt the same beliefs. What it is to treat a subject matter as *objective* is to be committed to the idea that whatever we ought to believe should line up with what everyone else ought to believe. We could postulate, as a regulative assumption, an objective point of view, an Archimedean point, from which it is decided what *everyone* ought to believe, irrespective of their contexts. It is this regulative assumption of objectivity that is captured by the idea that we ought to believe what is true.

Philosophical inquiry lacks this objectivity. It has a certain kind of permissiveness that shared inquiry lacks: as long as your worldview is sophisticated and coherent enough, you are *allowed* to believe it. We could agree to disagree about what is the best way to conceive of the world. This is the extent to which this is a personal matter, and perhaps a matter of taste. Therefore, while science and philosophy may both be normative inquiries, they deal in a different kind of normativity. The scientific kind of normativity is objective, whereas the philosophical kind of normativity is, to a certain extent, subjective. It is this difference, rather than the difference between norms and facts, that follows from the difference between shared and personal inquiry.

5.3 How Can Philosophy Talk about Objectivity?

One worry with this kind of view of objectivity and philosophy's lack of objective truth is one of self-refutation: we want to say that science is objective whereas philosophy is not, but this is itself a philosophical claim which thus cannot be objective. Therefore, science is objective, but it is not objectively objective. Could we talk about objectivity at all without thereby implicitly trying to make objective claims? The problem vaguely reminds me of an example by Frank Ramsey of an absurd dialogue with a child:

> 'Say breakfast.' 'Can't.' 'What can't you say?' 'Can't say breakfast.'" (Ramsey 1929b, 268)

Luckily, this problem of talking about objectivity has a solution. The philosophical metalanguage need not be objective itself as long as it can be specified in this language what one must *do* to count as saying something objective. Brandom calls this relation between a language and an expressively weaker metalanguage "pragmatic expressivist bootstrapping".

(Brandom 2008, 26) He uses indexical language as an example. We need not be able to use indexical language in the metalanguage as long as we can say in the metalanguage what we need to *do* to engage in indexical discourse. Thus, we need the metalanguage to express practical rules that tell us what we need to do to correctly deploy indexical language. These are the rules, according to Brandom:

> If, at time *t* and place <*x*, *y*, *z*>, speaker *s* wants to assert that some property *P* holds of <*x*, *y*, *z*, *t*, *s*>, it is correct to say "P holds of *me*, *here* and *now*." …
>
> If a speaker *s* at time *t* and place <*x*, *y*, *z*> asserts "P holds of *me*, *here* and *now*," the speaker is committed to the property *P* holding of <*x*, *y*, *z*, *t*, *s*>. (Brandom 2008, 25-26)

Note that indexical language is not used in these rules; the indexical terms are merely mentioned.

The same holds for objectivity: I have developed a philosophical theory that tells us what we must *do*, i.e., how we must engage with our practices, to treat a subject matter as objective. The practical rule that makes discourse objective is that we aim for truth. When formulated as such, it is unclear what one must do to adhere to this rule, yet much of this thesis has been devoted to the difference that aiming for truth makes to our inquiry. On a practical level, what it is to treat the subject matter of an inquiry as objective is for us to add an *a priori* fault to our disagreements. One of the consequences of doing so is that we must generally conciliate in the face of peer disagreement. Philosophy's lack of objectivity does not bar us from articulating such practical rules.

6. Philosophical Progress

I have argued that we do not aim for truth in philosophical inquiry. I have shown how it follows from the lack of this aim that steadfastness could be a rationally permissible response to philosophical disagreement. In this current section, I argue that this explains why widespread persistent disagreement in philosophy need not be indicative of a lack of philosophical progress.

Let me start by sketching the opposite view, the view that widespread persistent disagreement does indicate a lack of progress. This view has been argued for by various philosophers. For instance, David Chalmers argues for this conclusion in his article "Why Isn't There More Progress in Philosophy?" (Chalmers 2015). In his article, Chalmers compares philosophy to science, and he argues that, relative to science, there is not a lot of progress in philosophy. The first question that needs to be asked is why Chalmers thinks that philosophical progress can be compared to scientific progress. Chalmers is a methodological naturalist about philosophy, which is why he supposes that philosophical inquiry is to be modeled after scientific inquiry, and, consequently, that philosophical standards must emulate scientific standards. Given these similar standards, philosophical progress could be measured along the same axis as scientific progress. Chalmers then claims that there is much more disagreement in philosophy than in science and, therefore, less progress.

Yet, I reject methodological naturalism and have argued above that philosophical inquiry and scientific inquiry are, for the most part, fundamentally different because the latter is characterized by the collective aim for truth, which the former lacks. I shall show that this difference has important implications for how we should view philosophical progress.

6.1 Disagreement and Knowledge

According to Peter van Inwagen (2004), the main problem of persistent widespread disagreement among philosophers is that, due to this kind of disagreement, philosophers lack adequate justification for any of their views. Therefore, philosophers cannot *know* the solution to a philosophical problem even if they have developed this solution themselves. This argument has been described in more detail by Daniel Stoljar in his book *Philosophical Progress: In Defence of a Reasonable Optimism* (2017). As the title of his book suggests, Stoljar rejects the pessimistic conclusion, and he does so by rejecting the empirical fact that, when compared to science, there actually is a lot of widespread disagreement in philosophy. I will, however, assume that there is plenty of widespread disagreement, and I will advert to Chalmers's paper (2015) for empirical evidence for this assumption.

Stoljar asks us to consider a case in which two philosophers disagree on whether materialism or dualism is true. They have both examined all the available arguments for either side, they are equally competent philosophers and regard each other as epistemic peers. The question is whether they should change their belief due to this disagreement. Stoljar believes that a natural line of thought is that they should suspend judgment on the matter. (Stoljar 2017, 131) After all, they know that at most one of them can be correct, and they have no good reason to believe that they themselves are more likely to be correct, so as a result they should give as much weight to the opponent's view as to their own. Given that they should suspend judgment, they are not justified in believing their views, which is why they cannot gain philosophical knowledge. Widespread disagreement gets in the way of philosophical progress by ruling out the possibility of philosophical knowledge.

As we now know, this argument implies the conciliatory view on disagreement. If steadfastness is rationally permissible, one can remain justified in believing their theory in the face of philosophical peer disagreement. Once we reject that philosophy aims for truth, and we

thereby accept that philosophical inquiry is personal, it becomes clear that the steadfast view on peer disagreement holds for philosophical inquiry and that, therefore, the mere presence of widespread disagreement does not automatically rule out the possibility of philosophical knowledge.

However, the above argument for why widespread disagreement gets in the way of philosophical progress does not only rely on the conciliatory view, but it also implies that progress is constituted by the accumulation of knowledge. This is a contestable implication, and there are many rival views of what constitutes progress. These other views are roughly as follows: 1) there is progress if and only if the developed theories get closer to the truth, 2) there is progress if and only if philosophical and/or scientific problems have been solved, and 3) there is progress if and only if there has been an increase in understanding. (Dellsén, Lawler and Norton, *forthcoming*)

Van Inwagen's conclusion that there isn't much progress in philosophy is a relative statement. To substantiate it, a comparison must be drawn between progress in philosophy and progress in another inquiry. As mentioned above, the comparison that has been made is between philosophy and science: there is not a lot of progress in philosophy when compared to science because there is more widespread disagreement in philosophy than there is in science. As it turns out, though, none of the alternative views on progress mentioned above could even accommodate this comparison between scientific and philosophical progress. Once we realize that philosophical inquiry is personal, whereas scientific inquiry is shared, it becomes clear that none of these accounts of progress can apply to both forms of inquiry.

(1) could be a good account of scientific progress, even though it does seem strange to say that the only thing that matters for progress is that our theories become truer, regardless of whether they could justifiably be believed. Yet, on my view of philosophical inquiry, it could

not be an adequate account of philosophical progress because this view on progress implies that the aim of the inquiry is truth, which, I have argued, it isn't in the case of philosophy.

(2) is a kind of anti-realist view on progress, which does not mention truth at all. This would make it suitable as an account of philosophical progress but unsuitable as an account of scientific progress. This is because truth is the aim of science, and, given that one makes progress only insofar as one gets closer to their aim, there should be some connection between scientific progress and truth, even if truth turns out to merely be a necessary condition.

Likewise, (3) could be a suitable account of philosophical progress, but it is unsuitable as an account of scientific progress. As argued in Chapter 5, one of the ways in which scientific practices differ from philosophical practices is in the mechanics of testimony. In the case of science, it is relatively easy to gain knowledge through pure testimony, without having to understand the facts that one is being told. This is relatively difficult to achieve in the case of philosophy. In philosophy, there is more of an obligation to understand the theories that one professes to know. If every individual scientist aims to understand everything, scientific inquiry won't get off the ground. Scientists need to be able to rely on the background research of their peers without understanding all the details. Of course, (3) could be interpreted as saying that the goal is not understanding at the level of the individual but at the level of the community, but this shared aim would then conflict with the personal nature of philosophical inquiry. Finally, (3) does not suggest that we aim for understanding *in addition* to knowledge but that we do so *instead* of aiming for knowledge. This is contrary to a realist view of science. It implies that we just aim to gain an understanding of the world around us, regardless of whether our theories are true. For example, Newtonian mechanics does a great job of helping us understand the world around us, but it is nevertheless a false theory. If our aim is truth, then a mere increase in understanding would not be enough to constitute progress.

Therefore, the only account of progress that seems to work for both science and philosophy is that of progress as the accumulation of knowledge. This does not mean that these other accounts are wrong, but they would be ineffective at getting the point across that Van Inwagen wants to get across with his argument, which is that, *compared to science*, there isn't much progress in philosophy. To make such a comparison at all, we must assume that progress is constituted by the accumulation of knowledge. Thus, since I want to argue that philosophical progress is of a different nature than scientific progress, there is only one account of progress that should trouble me: progress as the accumulation of knowledge.

That said, this also seems to be the account of progress that is most natural and intuitive (which is evidenced by the fact that Van Inwagen and Stoljar's arguments presuppose it). I think the accumulation of knowledge could be seen as the gold standard of progress. All the other views aim lower by saying that there is still some amount of progress if one develops a true theory, a solution to a problem, or an understanding of some phenomenon, even if one is not in a position to justifiably believe it. Views (2) and (3) could even fall prey to a relativist interpretation: the philosophers who came before us developed *different* philosophical theories to solve *different* philosophical problems (the problems of their times) and consequently developed a *different* understanding of the world, but in what way are their theories *worse* and ours *better*? My view is unique because it shows that one can accumulate knowledge in the face of disagreement, which arguably constitutes a more substantive kind of progress.

Thus, whereas others, like Dellsén, Lawler and Norton (*forthcoming*), try to counter Van Inwagen's argument by arguing that philosophical progress should not be seen as the accumulation of knowledge, I show that the argument fails to work even if we do think of philosophical progress as the accumulation of knowledge. It fails to work because it implies the conciliatory view, and it thus falsely implies that philosophical inquiry is shared and aims

for truth. Once we reject that philosophy aims for truth, we can see how widespread disagreement need not stand in the way of knowledge.

Granted, my account of knowledge for personal inquiry is not the same as that for shared inquiry because it involves the aim for practical certainty instead of truth. This is why philosophical progress is still different in nature from scientific progress, even if both could be said to consist in the accumulation of knowledge. Nevertheless, the aim for practical certainty is still an ambitious aim, especially given that other views of philosophical progress concede that one could not even justifiably believe one's philosophical theses. This is a concession that need not be made once we realize that philosophical inquiry is personal. Since philosophical inquiry is personal, we could still be rationally entitled to our philosophical beliefs in the face of widespread disagreement.

It is important to emphasize that this does not mean that philosophers could disregard disagreement altogether. What I am suggesting instead is that philosophers should examine all the arguments against their thesis and judge their strengths. If they deem all of these arguments to be weak or misguided, then they do not have to lose confidence in their view. It is in this sense that the mere presence of disagreement need not negatively affect the confidence with which one believes their own theory. This might seem like a rather modest conclusion since it would still put knowledge mostly out of reach, but then again, knowledge should not be too easy to come by.

6.2 Collective Knowledge

Chalmers does not actually endorse Van Inwagen's argument and, in spite of his naturalism, he agrees that the existence of widespread persistent disagreement need not prevent us from knowing a philosophical theory. However, he argues instead that even though philosophical

knowledge is possible, what is constitutive of philosophical progress is not just knowledge but *collective knowledge*, which is still impossible when there is persistent widespread disagreement. (Chalmers 2015)

The ideal of collective knowledge is probably inspired by Chalmers's naturalism. After all, it does seem right to say that collective knowledge is the aim of scientific inquiry. The aim for collective knowledge is the same as the aim for *shared* knowledge, which, I argued in Chapter 5, characterizes shared inquiry. However, philosophical inquiry is not shared, like science. This is why aiming for collective knowledge about philosophical matters makes as much sense as aiming for collective knowledge about the tastiness of ice cream or the fashionableness of elbow patches.

Given the personal nature of philosophical inquiry, the standard of progress must be personal, not shared. However, this seems rather relativistic, and it seems to betray our intuitions that some of our standards are shared. Nevertheless, I have argued that there are shared standards in the case of philosophy. Not all standards are as universally endorsed as the standards of science; they are more like the shared aesthetic principles of an artistic movement. Still, though, there are some very general theoretical standards that most of us try to abide by. In a previous section, I already mentioned Blackburn's examples: "information, sensitivity, maturity, imagination, coherence." (Blackburn 1998. 318) We can accept such theoretical virtues as shared standards, but it is important to focus on their open-endedness, and to deny that there are objective truths about which theories satisfy these norms. In a sense, these norms are indeed like aesthetic principles, and it could be a mere opinion which theory counts as more sensitive or imaginative.

These standards could help conceive of philosophical progress from the point of view of the community: there is progress whenever both sides of a disagreement work out the incoherencies of their views, even if, by doing so, they make their disagreements more and

more entrenched and intractable. This counts as progress because it allows us to state more and more confidently that the disagreements in question are merely due to a matter of taste. The more confident we get in our belief that our disagreement is due to a difference in taste, the more confident we are in exercising steadfastness. Thus, if anything, philosophical progress isn't usually made by resolving our disagreements (though this wouldn't be opposed to progress), but it is often made by making our disagreements more and more entrenched, until they become incompatible but internally coherent, mature, and informative worldviews.

Even though we are still disagreeing about whether to be nominalists or Platonists, consequentialists or deontologists, empiricists or rationalists, it seems clear to me that this kind of progress has been made in these debates, because the versions of the views on offer are a lot more informative, mature, coherent, etc. This holds especially for the current day and age, given that there are now so many people working in academic philosophy who try to fine-tune these philosophical debates. Of course, there are some serious concerns with our current philosophical practices; we might be jumping on and off hot topics too frequently, we are overproducing in a way that might make us unable to see the forest for the trees, the increasing need to get a lot of papers published may be detrimental to the quality of the work, non-analytic approaches to philosophical problems are being overlooked, etc. Yet, by the abovementioned criteria, current philosophical debates and contemporary philosophical theories are at a much higher level of sophistication than their predecessors.

Once we reject the comparison to science and realize that convergence is not indicative of progress, it becomes easy to see that there is indeed plenty of progress in philosophy. By crystallizing our disagreements, we do not tend to resolve them, but we do improve our respective views, iron out the kinks, and we can show that these disagreements turn on personal convictions and sentiments, which allows us to become more and more certain of our views, and thus brings us closer to knowledge.

7. Conclusion

This chapter started by showing how many of our philosophical practices display more similarities to the practices of personal inquiry than to those of shared inquiry. It then proceeded to compare philosophical inquiry to scientific inquiry in particular, arguing against the naturalist view that philosophical practices should emulate scientific practices. The positive view of philosophical inquiry developed here is a view that says that philosophy is about developing a *worldview*, a theory about how we ought to conceive of the world and our place in it. I have argued that this normative element could also appear in scientific inquiry, but that philosophical inquiry is nevertheless different; even if scientific inquiry is about how we ought to conceive of the world, it presupposes that this question has an objectively true answer, and thus that we all ought to conceive of the world in the same way. Philosophical inquiry, on the other hand, should not rule out the possibility that multiple worldviews are all internally coherent, and that a disagreement about which is the right one turns on a mere matter of taste or inclination.

Once we think of philosophical inquiry as personal rather than shared, it becomes easy to see that there has been plenty of progress in philosophy in spite of the existence of widespread philosophical disagreement. I have employed the lessons about peer disagreement from Chapter 3 to argue that there can be philosophical progress in the presence of widespread disagreement, even if we interpret progress as the accumulation of knowledge.

The claim this chapter, and this thesis in general, has argued for — philosophy does not aim for truth — is provocative and might seem, to some, too silly to be taken seriously. I hope, however, that, in addition to arguing for this claim by arguing *against* the naturalist view of philosophy, this chapter has also managed to sketch a positive view about the nature of philosophical inquiry and philosophical progress that shows why philosophy as a personal inquiry is a worthwhile enterprise.

7. Conclusion

This thesis has attempted to answer a rather straightforward question: Does philosophy aim for truth? One possible strategy for answering this question is by asking what truth is, and whether philosophical discourse is apt for truth. The answer to the latter question would presumably be yes, so once we have established that philosophical discourse is apt for truth, it seems that we are just a very small step removed from also establishing that philosophical inquiry aims for truth.

This is not the strategy I have taken in this thesis. Instead of launching headfirst into a metaphysical inquiry into the nature of truth, I have started with the question "What is it to aim for truth?" It might seem as if answering this question requires that we already know what truth is. For instance, according to the coherence theory of truth, to aim for truth is to aim for our beliefs to cohere, whereas for the correspondence theory of truth, the aim for truth is the aim for accurate representation. However, while these answers all operate on the theoretical level, I am interested in the practical level: What difference does the aim for truth make to our practices?

I was greatly helped by the work of Huw Price, who has already asked and answered this question for me. The aim for truth, Price argues, adds pressure to resolve our disagreements and to subsequently converge toward the same beliefs. (Price 2003) We use truth as a device for the coordination of our beliefs. However, I have shown that we do not always seem to use truth in this way, and that we *should* not always use it in this way. The utility of aiming for truth depends on the utility of the assumption of sameness, which is the assumption that we experience and process things in the same way and that, therefore, the beliefs that work best for me are generally those that work best for everyone.

This prompted me to make a distinction between those inquiries for which this assumption seems practically advantageous and those for which it does not. The former kind of inquiry, which I labeled *shared inquiry*, is thus the kind of inquiry for which it is practically advantageous to aim for truth, whereas the latter kind of inquiry, *personal inquiry*, is the kind of inquiry for which aiming for truth would be disadvantageous. For example, it seems pointless to aim to converge in our taste judgments since taste is personal, so it would not be productive to aim for truth when inquiring into matters of taste.

It is tempting to think of personal inquiry as dealing in mere sentiments that are wholly personal. In that case, disagreement about such matters would not be productive in any way; you have your inclinations and I have mine, and that's the end of it. However, I have argued that disagreement does serve a practical purpose in the case of personal inquiry. It aids us in structuring our beliefs in a productive way, and it allows us to keep each other in check. It is thus beneficial to engage in games of giving and asking for reasons, even if we do not aim for truth. Furthermore, truth still has a use in such games, not as a standard to aim for, but as a device that allows us to make doxastic allegiances and transfer our justificatory burdens.

While disagreement can still be productive in the case of personal inquiry, I have argued that its mechanics are different from those of disagreement belonging to shared inquiry. The aim for truth rationally requires us to conciliate when faced with peer disagreement. When this aim is absent from the inquiry, we could be rationally entitled to remain steadfast in the face of peer disagreement.

This contrast drawn between shared and personal inquiry is not exactly a novel one, of course; it is essentially the contrast between objective and subjective matters. Yet, the objective/subjective distinction is a metaphysical one that attempts to explain the practical differences mentioned above. I have argued that we should turn this explanatory order on its head. The objectivity of a subject matter is to be explained by the practical efficacy of treating

the inquiry into this subject matter as shared. Nevertheless, when we treat an inquiry as shared, we commit ourselves to an evidence-transcendent standard of correctness, which means that, even though this commitment to objectivity is downstream from a practical commitment, the resulting view is resoundingly realist.

This motivates the question: What could truth be for shared inquiry? I have argued that deflationism lacks the resources to show why shared inquiry aims for truth, whereas personal inquiry does not. This is roughly because deflationists must maintain that aiming for truth is the same as aiming to make correct assertions. While this might sound trivial, we also aim to assert correctly in the case of personal inquiry, and yet we do not aim for truth. The normativity of truth thus couldn't be reduced to the normativity of assertibility, which is why truth cannot be deflationary. Given that aiming for truth entails adding an *a priori* fault to our disagreements, and it thereby implies an evidence-transcendent standard, truth couldn't be epistemic either. I argue instead that we should think of truth as correspondence for shared inquiry.

Yet, I argue that truth could also be correspondence for personal inquiry, in a deflated sense. As long as we do not aim for truth, truth cannot be an evidence-transcendent standard because it wouldn't be used as a standard at all, and it, therefore, does not commit us to realism. It is the aim for truth that adds the metaphysical substance to our inquiry, not the truth property itself, which could just be seen as correspondence across the board.

The concept of truth could also be seen as unified. Truth's primary linguistic function is as a device for expressing agreement, so we could think of the concept of truth as such. However, there are different purposes for which we could use such a device. In the case of personal inquiry, we use it merely for the transmission of justificatory burdens. In the case of shared inquiry, we also use it as a norm, a standard, which conveys that, in addition to aiming for beliefs that work best, we must also aim for agreement. This may be classified as a pluralist account of truth, but only at the level of use.

The fact that personal inquiry engages in games of giving and asking for reasons suffices to show that personal inquiry is guided by reasons. However, I have argued for two kinds of rationality. As mentioned above, in the case of personal inquiry, we mustn't aim for convergence, and we must instead accept that, sometimes, we are just differently inclined. These different inclinations create different background commitments. These differences figure into our assessment of reasons. Consider the following: "This pasta is disgusting because it has cheese in it." This reason is admissible to those who believe that cheese is disgusting but not to those who do not. This is fairly obvious, of course, and this kind of difference in assessment can also be found in shared inquiry. Take, for example, the following claim: "The iron has melted because the temperature was over 1,538 ºC." This reason is admissible to those who believe that iron melts at temperatures over 1,538 ºC but not to those who do not.

The difference is that we are allowed to have different beliefs about whether cheese is disgusting, whereas we are not allowed to have different beliefs about whether iron melts at temperatures over 1,538 ºC. We might say, for example, that another is *personally* warranted to assert that the iron has not melted because we know that they believe that iron only melts at temperatures over 2,000 ºC. This would be a kind of *personal rationality*. Yet, in the case of shared inquiry, we must ultimately believe that our background beliefs ought to be shared. Therefore — assuming that our epistemic context is better than the person who believes that iron only melts at temperatures over 2,000 ºC — we must maintain that the belief that the temperature is below 2,000 ºC is ultimately not a good reason for believing that the iron has not melted. In doing so, we appeal to a different notion of rationality: *shared rationality*. In the case of personal inquiry, rationality is not shared. When we assess the views of another, we must adhere to their rules — or, rather, our interpretation of their rules — not ours.

A similar distinction also appears in the case of knowledge. I have argued that the aim for knowledge, from a personal perspective, amounts to the aim for practical certainty. When

assessing whether we *know* that *p*, we ask ourselves whether we are certain about *p*. This can be contrasted to the case in which we ask whether someone else knows that *p*. In that case, we do not care much whether they are certain about *p*, as long as they believe it. What we assess is mostly whether they are justified in believing *p*, and whether *p* is actually true. However, when assessing this second question, we are simply assessing whether we ourselves are doxastically committed to *p* as well. Therefore, truth only becomes relevant in the case of knowledge attributions, and it plays the same role it always does by being a device for expressing agreement. When I say "What *S* says is true" I express a shared doxastic commitment, and I thereby either take on *S*'s justificatory burdens or, in a different dialectical context, put my own justificatory burdens onto *S*. When I say "*S* knows that *p*" I likewise express a shared doxastic commitment to *p*, and I either take on *S*'s justificatory burdens or, in a different dialectical context, put my own justificatory burdens onto *S*. However, just like we can aim for truth, so can we aim for knowledge, in this shared sense, as the aim not just for practical certainty but also for shared doxastic commitments. This is what we do in the case of shared inquiry.

If our inquiry is personal, then we can know things in spite of the presence of widespread disagreement. For example, I could know that killing animals is wrong even though I also know that many people disagree with me on this. This would require me to become acquainted with all, or at least most, of the arguments usually raised against this view. If I do so and I judge all these arguments to be unpersuasive, then the disagreement would not undercut my justification anymore, and I would be rationally entitled to remain steadfast. I have argued that this is not the case for shared inquiry. However, while widespread disagreement is always problematic in the case of shared inquiry, when there is widespread *agreement*, we are in a position to learn and know things without being epistemically obliged to understand why these things are the case. I could know that $e = mc^2$ by noting the consensus

in the scientific community, without having to understand why this is so. It generally does not work that way for personal inquiry. For personal inquiry, it is important that we personally understand why things are the case, for if widespread disagreement need not be indicative of an inquiry having gone astray, widespread agreement does not necessarily indicate that the inquiry is on the right track either.

This allows us to draw the first comparison to philosophical inquiry. It certainly seems that understanding a philosophical theory is generally a necessary condition for knowing it to be true. This is why we try to foster understanding rather than knowledge in our teaching. We also seem to have a very high tolerance for disagreement in our philosophical practices, and a steadfast response often seems appropriate. When two philosophers disagree, we do not expect them to conciliate but we expect the disagreement to only become more and more entrenched. It is possible that this is just bad practice, and that we *should* conciliate. However, I have argued that doing so would not be practically advantageous. In the case of scientific disagreement, it is useful to suspend judgment and continue our inquiry because it is useful to believe that our continued efforts will turn up some new evidence that will push out some of the contending views and help us converge toward the same beliefs. This utility breaks down in the case of philosophy. Therefore, it does not only seem that we carry out philosophical inquiry as personal inquiry, there are also practical reasons why we should do so. The conclusion that can be drawn from this is that philosophy does not and should not aim for truth.

I have sketched a view of what philosophy might be once we give up on trying to compare it to science. Philosophy, I argue, is not about describing reality, but it is about building *worldviews*, which are normative theories that aim to tell us how we ought to conceive of the world and our place in it. This conception of philosophy undermines the distinction commonly drawn between theoretical philosophy, as the inquiry into the nature of reality, and practical philosophy, as the inquiry into how we ought to live. Instead, the whole of philosophy

is about how we ought to live; it is all practical. The project of figuring out how to conceive of the world is a part of the general project of figuring out how we ought to live. It should come as no surprise that this resonates with my commitment to pragmatism. The pragmatist would say that the normative character of philosophical inquiry is ultimately grounded in utility. When we ask how we ought to conceive of the world, we ask which conception of the world would best serve our practices.

I have argued that we can still think of philosophy as a highly disciplined activity, and not just as an arbitrary pick and choose of philosophical commitments that satisfy one's personal tastes or inclinations. Yet, instead of arguing that there should be one method, the scientific method, that dominates every aspect of our inquiry, I have shown that philosophy is more similar to the schools and isms of the artworld. There are many different philosophical methodologies, and methodological principles often run together with substantial philosophical commitments. There is no neat distinction to be drawn between philosophy and metaphilosophy. It is nevertheless possible to disagree with another's theory, due to the endorsement of different metaphilosophical principles, while also appreciating that this theory is internally coherent, given the rules it plays by. In that case, we would still disagree about which rules we ought to play by, but we must be mindful of the possibility that the disagreement is ultimately due to a mere matter of taste.

Once we give up on the idea that philosophy aims for truth, we can understand how there might still be philosophical progress in the absence of convergence. This is because, as argued above, the mere presence of widespread disagreement need not preclude the possibility of knowledge. Through the practice of disagreement, and the game of giving and asking for reasons, we can develop our respective theories and make them more and more refined and coherent. This usually does not lead to conciliation but rather to entrenchment. A large share of the activities of philosophy is the activity of insulating one's views from counterarguments.

This entrenchment would seem highly problematic if we were to take philosophy to be aimed at truth. However, philosophy does not aim for truth, which is why there need be nothing wrong with the practice of developing our theories in opposition to one another.

I hope this goes some way to showing how we could conceive of philosophical inquiry as a personal inquiry. However, I would like to conclude by stating the main thesis that I have argued for, which is not this positive view, but which is instead a negative thesis, a denial of an assumption that for many might seem trivial, and a radical departure from the idea that philosophical inquiry is scientific. *Philosophy does not aim for truth*.

Bibliography

Asay, Jamin. 2013. *The Primitivist Theory of Truth.* Cambridge: Cambridge University Press.

Ayer, A. J. 1936. *Language, Truth and Logic.* London: Gollancz.

Barnett, Zach. 2019. "Philosophy Without Belief." *Mind* 128 (509): 109-138.

Bernáth, László & Tőzsér, János. 2021. "The Biased Nature of Philosophical Beliefs in the Light of Peer Disagreement." *Metaphilosophy* 52 (3-4): 363-378.

Blackburn, Simon. 1998. *Ruling Passions: A Theory of Practical Reasoning.* Oxford: Oxford University Press.

Blackburn, Simon. 2005. *Truth: A Guide for the Perplexed.* London: Allen Lane.

Boghossian, Paul. 2005. "Is Meaning Normative?" In *Philosophy — Science — Scientific Philosophy. Main Lectures and Colloquia of GAP.5, Fifth International Congress of the Society for Analytical Philosophy, Bielefeld, 22-26 September 2003* edited by Christian Nimtz & Ansgar Beckermann. Paperborn: 205-218.

BonJour, Laurence. 1985. *The Structure of Empirical Knowledge.* Cambridge MA: Harvard University Press.

Bourget, David. & Chalmers, David. Unpublished Manuscript. "Philosophers on Philosophy: The 2020 PhilPapers Survey."

Brandom, Robert. 1994. *Making it Explicit.* Cambridge MA: Harvard University Press.

Brandom, Robert. 2008. *Between Saying and Doing: Towards an Analytic Pragmatism*. Oxford: Oxford University Press.

Cappelen, Herman. 2018. *Fixing Language: An Essay on Conceptual Engineering.* Oxford: Oxford University Press.

Carnap, Rudolf. 1959. “The Elimination of Metaphysics Through Logical Analysis of Language.” In *Logical Positivism* edited by A. J. Ayer. New York: The Free Press. 60-81.

Chalmers, David. 2015. “Why Isn’t There More Progress in Philosophy?” *Philosophy* 90 (1): 3-31.

Comesaña, Juan. 2012. “Conciliation and Peer-Demotion in the Epistemology of Disagreement.” *American Philosophical Quarterly* 49 (3): 237-251.

Davis, John. 2015. “Faultless Disagreement, Cognitive Command, and Epistemic Peers.” *Synthese* 192 (1): 1-24.

De Langhe, Rogier. 2013. “Peer Disagreement under Multiple Epistemic Systems.” *Synthese* 190: 2547-2556.

Dellsén, Finnur, Lawler, Insa & Norton, James. *Forthcoming*. “Would Disagreement Undermine Progress?” *The Journal of Philosophy.*

De Ridder, Jeroen. 2014. “Why Only Externalists Can Be Steadfast.” *Erkenntnis* 79 (S1): 185-199.

Dewey, John. 1929. *The Quest for Certainty: A Study of the Relation of Knowledge and Action.* New York: Minton, Balch & Company.

Douven, Igor. 2010. “Simulating Peer Disagreement.” *Studies in History and Philosophy of Science Part A* 41 (2): 148-157.

Elga, Adam. 2010. “How to Disagree About How to Disagree.” In *Disagreement* edited by Richard Feldman & Ted A. Warfield. Oxford: Oxford University Press. 175-186.

Elgin, Catherine. 2022. “Disagreement in Philosophy.” *Synthese* 200 (20): 1-16.

Englert, Alexander. 2022. “The Conceptual Origin of Worldview in Kant and Fichte.” *Journal of Transcendental Philosophy*: https://doi.org/10.1515/jtph-2022-0007.

Eriksson, John. 2016. “Expressivism, Attitudinal Complexity and Two Senses of Disagreement in Attitude.” *Erkenntnis* 81: 775-794.

Ferrari, Filippo. 2021. “Alethic Pluralism and the Value of Truth.” *Synthese* 199: 55-79.

Ferrari, Filippo. 2022. *Truth and Norms: Normative Alethic Pluralism and Evaluative Disagreements.* New York: Lexington Books.

Ferrari, Filippo & Pedersen, Nikolaj J. L. L. 2019. “Epistemic Peer Disagreement.” In *The Routledge Handbook of Social Epistemology* edited by Miranda Fricker, Peter J. Graham, David Henderson & Nikolaj J. L. L. Pedersen. London: Routledge.

Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing.* Oxford: Oxford University Press.

Gettier, Edmund. 1963. “Is Justified True Belief Knowledge?” *Analysis* 23 (6): 121-123.

Gibbard, Allan. 2005. “Truth and Correct Belief.” *Philosophical Issues* 15 (1): 338-350.

Goldberg, Sanford. 2013. “Defending Philosophy in the Face of Systematic Disagreement.” In *Disagreement and Skepticism* edited by Diego Machuca. London: Routledge. 277-294.

Grover, Dorothy, Camp, Joseph & Belnap, Nuel. 1975. “A Prosentential Theory of Truth.” *Philosophical Studies* 27 (2): 73-125.

Haslanger, Sally. 2012. *Resisting Reality: Social Construction and Social Critique.* Oxford: Oxford University Press.

Heney, Diana. 2015. “Reality as Necessary Friction.” *The Journal of Philosophy* 112 (9): 504-514.

Hopkins, Robert. 2000. “Beauty and Testimony.” *Royal Institute of Philosophy Supplement* 47: 209-236.

Horwich, Paul. 1998. *Truth.* Oxford: Oxford University Press.

Horwich, Paul. 2006. "The Value of Truth." *Nous* 40 (2): 347-360.

Howat, Andrew. 2020. "Pragmatism and Correspondence." *Philosophia* 49 (2): 685-704.

Hume, David. 1932. *The Letters of David Hume, Vol. 1: 1727-1765* edited by J. Y. T. Greig. Oxford: Oxford University Press.

James, William. 1907. *Pragmatism: A New Name for Some Old Ways of Thinking.* New York: Longmans, Green, and Co.

Jones, Karen. 1999. "Second-Hand Moral Knowledge." *The Journal of Philosophy* 96 (2): 55-78.

Kaspers, Tom. 2022. "Alethic Pluralism for Pragmatists." *Synthese* 200 (66).

Kelly, Thomas. 2005. "The Epistemic Significance of Disagreement." In *Oxford Studies in Epistemology, Volume 1* edited by John Hawthorne & Tamar Gendler. Oxford: Oxford University Press. 167-196.

Kelly, Thomas. 2010. "Peer Disagreement and Higher-Order Evidence." In *Disagreement* edited by Richard Feldman & Ted A. Warfield. Oxford: Oxford University Press. 111-174.

King, Nathan. 2012. "Disagreement: What's the Problem? *or* A Good Peer is Hard to Find." *Philosophical and Phenomenological Research* 85 (2): 249-272.

Kitcher, Philip. 2011. "Philosophy Inside Out." *Metaphilosophy* 42 (3): 248-260.

Kölbel, Max. 2004. "Faultless Disagreement." *Proceedings of the Aristotelian Society* 104 (1): 53-73.

Kornblith, Hilary. 2013. "Is Philosophical Knowledge Possible?" In *Disagreement and Skepticism* edited by Diego Machuca. London: Routledge. 260-276.

Kraut, Robert. 1990. “Varieties of Pragmatism.” *Mind* 99 (394): 157-183.

Kripke, Saul. 1982. *Wittgenstein on Rules and Private Language.* Cambridge MA: Harvard University Press.

Kripke, Saul. 2011. “On Two Paradoxes of Knowledge.” In *Philosophical Troubles: Collected Papers, Vol. 1.* Oxford: Oxford University Press. 27-51.

Lackey, Jennifer. 2008. “What Should We Do When We Disagree?” In *Oxford Studies in Epistemology Vol. 3* edited by Tamar Szabó Gendler & John Hawthorne. Oxford: Oxford University Press. 274-293.

Lane, Robert. 2018. *Peirce on Realism and Idealism.* Cambridge: Cambridge University Press.

Levy, Neil. 2021. “The Surprising Truth about Disagreement.” *Acta Analytica* 36: 137-157.

Lewis, Clarence Irving. 1923. “A Pragmatic Conception of the *A Priori*.” *The Journal of Philosophy* 20 (7): 169-177.

Lewis, Clarence Irving. 1946. *An Analysis of Knowledge and Valuation.* La Salle: The Open Court Publishing Company.

Lewis, David. 2000. “Academic Appointments: Why Ignore the Advantage of Being Right?” In *David Lewis: Papers in Ethics and Social Philosophy.* Cambridge: Cambridge University Press. 187-200.

Lynch, Michael. 2004. “Minimalism and the Value of Truth.” *The Philosophical Quarterly* 54 (217): 497-517.

Lynch, Michael. 2009. *Truth as One and Many.* Oxford: Oxford University Press.

MacFarlane, John. 2014. *Assessment Sensitivity: Relative Truth and Its Applications.* Oxford: Oxford University Press.

McDermid, Douglas. 2006. *The Varieties of Pragmatism: Truth, Realism, and Knowledge from James to Rorty.* London: Continuum.

Misak, Cheryl. 2004. *Truth and the End of Inquiry: A Peircean Account of Truth.* Oxford: Oxford University Press.

Misak, Cheryl. 2011. "2011 Presidential Address: American Pragmatism and Indispensability Arguments." *Transactions of the Charles S. Peirce Society* 47 (3): 261-273.

Misak, Cheryl. 2013. *The American Pragmatists.* Oxford: Oxford University Press.

Misak, Cheryl. 2015. "Pragmatism and the Function of Truth." In *Meaning Without Representation: Essays on Truth, Expression, Normativity, and Naturalism* edited by Steven Gross, Nicholas Tebben, and Michael Williams. Oxford: Oxford University press. 262-278.

Misak, Cheryl. 2018. "The Pragmatist Theory of Truth." In *The Oxford Handbook of Truth* edited by Michael Glanzberg. Oxford: Oxford University Press. 283-304.

Nickel, Philip. 2001. "Moral Testimony and Its Authority." *Ethical Theory and Moral Practice* 4: 253-266.

Pedersen, Nikolaj. J. L. L. 2012. "Recent Work on Alethic Pluralism." *Analysis* 72 (3): 588-607.

Peirce, Charles Sanders. 1877. "The Fixation of Belief." *Popular Science Monthly* 12: 1-15.

Peirce, Charles Sanders. 1878. "How to Make Our Ideas Clear." *Popular Science Monthly* 12: 286-302.

Plunkett, David. & Sundell, Timothy. 2013. "Disagreement and the Semantics of Normative and Evaluative Terms." *Philosopher's Imprint* 13 (23): 1-37.

Price, Huw. 1983. "'Could a Question Be True?': Assent and the Basis of Meaning." *The Philosophical Quarterly* 33: 354-364.

Price, Huw. 1988. *Facts and the Function of Truth.* Oxford: Blackwell.

Price, Huw. 1997. "Carnap, Quine and the Fate of Metaphysics." *Electronic Journal of Analytic Philosophy* 5 (1).

Price, Huw. 1998. "Three Norms of Assertibility, or How the Moa Became Extinct." *Noûs* 32 (S12): 241-254.

Price, Huw. 2003. "Truth as Convenient Friction." *The Journal of Philosophy* 100 (4): 167-190.

Price, Huw. 2011. *Naturalism Without Mirrors.* Oxford: Oxford University Press.

Price, Huw. 2022a. "Family Feuds? Relativism, Expressivism, and Disagreements about Disagreement." *Philosophical Topics*, 50 (1): 293-344.

Price, Huw. 2022b. "Global Expressivism and Alethic Pluralism." *Synthese* 200 (386): 1-55.

Putnam, Hilary. 1981. *Reason, Truth and History.* Cambridge: Cambridge University Press.

Putnam, Hilary. 1987. "Truth and Convention: On Davidson's Refutation of Conceptual Relativism." *Dialectica* 41 (1-2): 69-77.

Quine, Willard van Orman. 1951. "On Carnap's Views on Ontology." *Philosophical Studies* 2 (5): 65-72.

Ramsey, Frank. 1927. "Facts and Propositions." *Proceedings of the Aristotelian Society, Supplementary Volumes* 7: 153-170.

Ramsey, Frank. [1929a] 1931. "General Propositions and Causality." In *The Foundations of Mathematics and Other Logical Essays* edited by R. B. Braithwaite. London: Routledge & Kegan Paul Ltd. 237-255.

Ramsey, Frank. [1929b] 1931. “Philosophy.” In *The Foundations of Mathematics and Other Logical Essays* edited by R. B. Braithwaite. London: Routledge & Kegan Paul Ltd. 263-269.

Robson, Jon. 2012. “Aesthetic Testimony.” *Philosophy Compass* 7 (1): 1-10.

Rorty, Richard. 1979. *Philosophy and the Mirror of Nature.* Princeton: Princeton University Press.

Rorty, Richard. 1990. “Pragmatism, Davidson and Truth.” In *Objectivity, Relativism and Truth: Philosophical Papers*. Cambridge: Cambridge University Press. 126-150.

Rorty, Richard. 1995. “Is Truth a Goal of Inquiry? Davidson vs. Wright.” *The Philosophical Quarterly* 45 (180): 281-300.

Rosen, Gideon. 2001. “Nominalism, Naturalism, Epistemic Relativism.” *Philosophical Perspectives* 15: 69-91.

Schaffer, Jonathan. 2009. “On What Grounds What.” In *Metametaphysics: New Essays on the Foundation of Ontology* edited by David Manley, David Chalmers & Ryan Wasserman. Oxford: Oxford University Press. 347-383.

Schiller, F.C. S. 1907. *Studies in Humanism*. London: MacMillan and Co. Limited.

Šešelja, Dunja. 2019. “Some Lessons from Simulations of Scientific Disagreements.” *Synthese* 198: 6143-6158.

Shapiro, Lionel. 2021. “Truth’s Dialectical Role: From Friction to Tension.” *Inquiry*: https://doi.org/10.1080/0020174X.2021.1972834.

Shields, Matthew. 2022. “Truth from the Agent Point of View.” *The Philosophical Quarterly*: https://doi.org/10.1093/pq/pqac072.

Sider, Theodore. 2009. “Ontological Realism.” In *Metametaphysics: New Essays on the Foundation of Ontology* edited by David Manley, David Chalmers & Ryan Wasserman. Oxford: Oxford University Press. 384-423.

Sider, Theodore. 2011. *Writing the Book of the World.* Oxford: Oxford University Press.

Simpson, Robert Mark. 2013. “Epistemic Peerhood and the Epistemology of Disagreement.” *Philosophical Studies* 164 (2): 561-577.

Stoljar, Daniel. 2017. *Philosophical Progress: In Defence of a Reasonable Optimism.* Oxford: Oxford University Press.

Strawson, P. F. 1949. “Truth.” *Analysis* 9 (6): 83-97.

Thomasson, Amie. 2020. *Norms and Necessity.* Oxford: Oxford University Press.

Van Inwagen, Peter. 2004. “Freedom to Break the Laws.” *Midwest Studies in Philosophy* XXVIII: 334-350.

Vavova, Katia. 2014. “Confidence, Evidence, and Disagreement.” *Erkenntnis* 79 (supp. 1): 173-183.

Weatherson, Brian. 2013. “Disagreements, Philosophical and Otherwise.” In *The Epistemology of Disagreement: New Essays* edited by Jennifer Lackey & David Christensen. Oxford: Oxford University Press. 54-73.

White, Roger. 2005. “Epistemic Permissiveness.” *Philosophical Perspectives* 19: 445-459.

Williamson, Timothy. 1996. *Vagueness.* London: Routledge.

Williamson, Timothy. 2013. *Modal Logic as Metaphysics.* Oxford: Oxford University Press.

Wittgenstein, Ludwig. 1921. *Tractatus Logico-Philosophicus.* Translated by C. K. Ogden & Frank Ramsey. London: Kegan Paul, Trench, Trubner & Co., Ltd.

Wittgenstein, Ludwig. 1953. *Philosophical Investigations.* Oxford: Blackwell.

Wittgenstein, Ludwig. 1991. *Philosophical Grammar.* Edited by Rush Rhees. Oxford: Wiley-Blackwell.

Worsnip, Alex. 2019. "Disagreement as Interpersonal Incoherence." *Res Philosophica* 96 (2): 245-268.

Wrenn, Chase. 2020. "Alethic Pluralism and Truth-Attributions." *American Philosophical Quarterly* 57 (4): 311-324.

Wright, Crispin. 1992. *Truth and Objectivity.* Cambridge MA: Harvard University Press.

Wright, Crispin. 2001. *Rails to Infinity: Essays on Themes from Wittgentein's* Philosophical Investigations. Cambridge MA: Harvard University Press.

Wright, Crispin. 2006. "Intuitionism, Realism, Relativism and Rhubarb." In *Truth and Realism* edited by Patrick Greenough & Michael Lynch. Oxford: Oxford University Press. 38-60.

Wright, Crispin. 2007. "Rule-Following Without Reasons: Wittgenstein's Quietism and the Constitutive Question." *Ratio* XX: 481-502.

Wright, Crispin. 2021. "Alethic Pluralism, Deflationism, and Faultless Disagreement." *Metaphilosophy* 52: 432-448.