TILBURG ◆ UNIVERSITY

Tilburg University

Epistemic Justice

Geuskens, Machteld

Publication date:
2018

Document Version
Publisher's PDF, also known as Version of record

Link to publication in Tilburg University Research Portal

Citation for published version (APA):
Geuskens, M. (2018). *Epistemic Justice: A Principled Approach to Knowledge Generation and Distribution.*
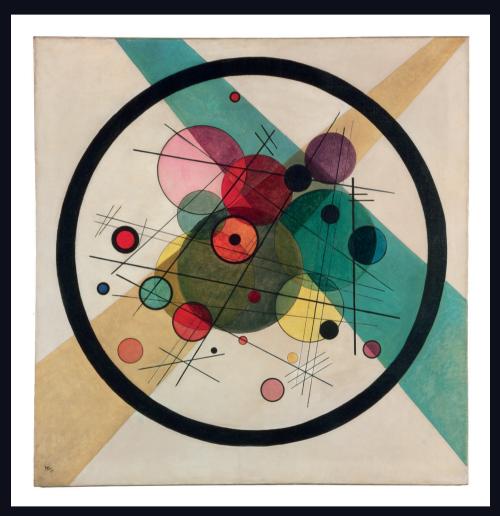[s.n.].

Epistemic Justice
A Principled Approach to Knowledge Generation and Distribution

Machteld Geuskens

Epistemic Justice

A Principled Approach to

Knowledge Generation and Distribution

Epistemic Justice

A Principled Approach to
Knowledge Generation and Distribution

Proefschrift ter verkrijging van de graad van doctor aan Tilburg University
op gezag van de rector magnificus, prof. dr. E.H.L. Aarts, in het openbaar te
verdedigen ten overstaan van een door het college voor promoties aangewezen
commissie in de Portrettenzaal van de Universiteit op

dinsdag 11 december 2018 om 16.00 uur

door

Machteld Clara Maria Geuskens
geboren te Leuven, België.

Promotor: Prof. dr. Alan P. Thomas
Copromotor: Dr. Filip A.I. Buekens

De overige leden van de promotiecommissie:
                  Dr. Paul Faulkner
                  Dr. Lisa Herzog
                  Dr. Philip J. Nickel
                  Prof. dr. Linda Radzik

Table of Contents

Preface

This thesis started when Alan Thomas emailed me to ask if I would be interested in pursuing a PhD, alerting me to an opportunity to obtain PhD funding from NWO. When I showed interest in doing so, we decided to "cook up a proposal", starting from my intuitions and convictions: a thesis in philosophy tends to be quite personal.

My strong conviction was (and still is) that knowledge should be shared much more widely than people often realise; that there is no point in keeping people in the dark for no reason, and, moreover, that doing so – granted that there really is no good reason – is both harmful for the person who is kept in the dark and an injustice. It is harmful in the sense that it can cause anxiety, frustration, or – when the epistemic neglect happens over a longer period – it can mean that the person comes to have less self-confidence or lowered self-esteem. It is an injustice, because there is no need to inflict that harm.

More positively put, my strong conviction is that sharing knowledge is empowering: that giving others knowledge that they may lack, even when unsolicited, is often more than a way of being helpful and kind. It is a form of inclusion, furthering trust in their abilities and trust in the nature of social relations. For, sharing knowledge is a way of entrusting others, of allowing others to be(come) one's equals, as it enables them to contribute their views too, and/or to better understand situations.

Having these intuitions and convictions is one thing; arguing for them in an academic context is quite another. Serendipitously, I received an emailed newsletter from Birkbeck College, mentioning a book entitled *Epistemic Injustice* by Miranda Fricker. That title struck me as so highly relevant – which indeed it became – that I could not help but work it into the proposal, wondering why its author had (apparently) not made the questions surrounding knowledge sharing more central to it. At the time, my dear friend Naomi Goulder who knew this literature strongly supported my idea that there was a lacuna. The proposal was written and the grant obtained. Here I would like to thank all of you who helped me, encouraged me, and accompanied me in the process of writing this thesis.

Alan, I feel incredibly lucky that you have been my supervisor. You were able to guide me in a process that was quite creative and daring. While I was keen to read Fricker's work, you pointed me to the works that I should study to start with. The breadth of your interests meant that you did not worry about my search for a fundamental guiding principle (i.e. epistemic justice) that could readily be applied to (a very wide range of) societal issues and discussions. Even as my topic related to anything and everything, however, and even as in philosophy every detail appears to matter, you managed to keep me from going astray, often using one-liners such as the time when you said it would be too time consuming "to go down every rabbit hole".

More fundamentally, you managed to keep me positive and inspired, throughout. Every meeting would end with a plan for action. I think it was your ability to inspire trust and reinforce positive commitment that helped to get this thesis materialised, from beginning to end. What helped, too, is your understanding that life is more than philosophy. Eating a cupcake was definitely the right way to spend time in York. So, I probably wrote more words here than I should have, but I hope we will continue to be in touch for a long time to come.

Filip, thank you for taking on the role of co-supervisor and always being inspired to talk about philosophical works, and interested to hear what I was working on.

Amanda Cawston, thank you for taking on the role of additional 'daily supervisor' this last year, which importantly included helping me monitor my progress.

Dear Committee Members, thank you for your time reading and assessing my thesis. Thank you, too, for your words that you look forward to reading publications based on this thesis.

Dear Matteo Colombo and Alfred Archer, you are role models to me in terms of doing serious and relevant philosophy. Thank you both for accepting the role of Opponent at my defence.

Dear (extended) family, I have started this foreword with information about my thesis. That is because I know many of you only read the foreword – and because I consider the strong intuitions that I started out with as part of our shared family heritage. There is a strong conviction that has run in our family for decades that all sorts of injustices and harms can be prevented through education and instruction. The drive to tell others how things ought to be done (or not done) has led to quite some hilarity when visitors read the notices that are put almost everywhere inside the family house in Koksijde. But, there is also a spirit of optimism that people are susceptible to arguments, such that explaining one's views – and clarifying one's intentions –is always in order. (If only people would communicate more, mutual understanding may be improved). Of course, people can fail to listen. My father understood that space is needed for being stubborn; he patiently repeated any points that he considered important, sometimes for years. It seems to me, then, that my convictions did not arise out of thin air. I hope and trust I am doing all of you justice by saying so.

Dear Colleagues (past and present), from both TiLPS at Tilburg University, and from STeM at Erasmus University: thank you for being wonderful colleagues and for becoming friends, too, as so much happened in the past six and two years, respectively. I have been very fortunate to work with you, even as I cannot thank each of you individually here for what you have meant.

One learns most from one's peers, but I think that is even more true of one's roommates. Thank you dear roommates – from Tilburg: Chiara Raucea, Zeynep Fortis, Dominik Klein,

Lieve mama, dankjewel dat je altijd voor me klaarstond en klaarstaat. De toewijding van jou en papa aan ons alle drie – en aan de kleinkinderen – is iets wat voor ons altijd vanzelfsprekend was, en wat altijd bij ons blijft. Het is de basis voor elke vorm van geluk. Ik zie je graag, heel graag, en ik hoop dat we samen nog veel mooie en goede dingen zullen doen en beleven, ook nu papa er niet meer is.

Lieve Maarten en Goedele, een betere broer en zus kan ik me niet wensen. Dat ik heel veel aan jullie te danken heb, daar sta ik nauwelijks bij stil... Wisten jullie dat jullie me allebei 'geluk' gaven op een heel praktische manier, deze laatste paar jaar?

Lieve Kees en Marja, dankzij jullie kon ik elke maandag zorgeloos aan de slag, ook toen Hannah nog maar net 10 weken oud was. Dankjewel dat jullie konden bijspringen ook, waar nodig, en geduldig en vol vertrouwen hebben gewacht tot er op een dag het nieuws zou komen dat mijn thesis af was.

Lieve Hannah, je grote enthousiasme voor boeken, al toen je twee jaar oud was, was voor mij heel inspirerend. Ik wilde spontaan ook meer boeken lezen. Ooit zul je dit boek misschien lezen – maar dat hoeft niet, want alles wat bruikbaar is geef ik aan je door. Ik kan niet anders dan mijn overtuigingen en wereldbeeld met jou te delen.

Lieve Niels, dankjewel dat je me altijd ruim baan gaf om aan mijn thesis te werken, in avonden, weekenden, en werkvakanties, en dat je me, tegelijkertijd, wist te motiveren om de afronding van mijn thesis niet te lang te laten duren. Wat ben ik blij dat we getrouwd zijn waar mijn vader bij was, en dat we samen de toekomst tegemoet gaan. Zoals jij het verwoordde: "Er is leven aan het einde van de tunnel" …. zowaar en zó waar! :)

Introduction

There is no shortage of discussion of the importance of knowledge sharing. Practical topics such as the ethics of whistleblowing in organizations; proposed changes in the law of privacy because of the impact of technology; the scope and limits of both responsible journalism and the need for government secrecy – in all of these contexts arguments for the sharing of knowledge are presented, contested, and defended. What all of these discussions have lacked, to this point, is a foundational grounding in a rigorous account of the norms that govern the proper distribution of knowledge. In this thesis I seek to address this issue, starting from the following philosophical question:

> Is there a (moral or epistemic) norm that knowledge ought to be shared *that follows from the nature of knowledge?*

If one excludes the part in italics, then this question may strike one as a question of applied normative philosophy. However, my aim is to inquire, fundamentally, into what *any* moral doctrine would need to say about knowledge. So, my inquiry is not one about how to apply a specific moral doctrine to knowledge, whereby we start from assumptions about the role knowledge plays in human life. Instead, we should start with what knowledge is (even while that question is hard to answer), and derive the norms that pertain to it from there. At this stage, we also do not preclude the answer that the normativity in question is epistemic rather than moral – hence the brackets. This approach creates room for an argument that the normativity has a *dual* moral-epistemic nature, and indeed it will be found and argued that the moral and epistemic are non-coincidentally related, as Miranda Fricker has argued (Fricker, 2007, p. 1, p. 17).

Of course, the ultimate aim of my work is a practical one. For my aim is not only to understand why, and to what extent, knowledge yields questions of a moral and political nature, but also to establish a framework that can be used to answer them. Once there is clarity about the rationale for or against knowledge sharing, then we can inquire what we should do, from a normative perspective with some particular item of knowledge, or with putative knowledge claims (claims that pass for or that are presented as knowledge). To address my main question, I have formulated the following sub-questions, which will be answered in subsequent chapters:

1. What is knowledge according to our philosophical understanding?
2. Is the objectivity of truth an obstacle to a social conception of knowledge; and how can that hurdle be overcome?

3. What normative rules and principles are needed for the generation and preservation of knowledge on a social conception of knowledge?
4. What is knowledge sharing, and how is it possible, according to philosophical conceptions of knowledge and of testimony?
5. How can the norms and responsibilities that we find in testimony and the nature of testimony itself best be explained?
6. What are the norms and principles that should structure our social epistemic practices, notably that of testimony, which ensure the generation and preservation of knowledge (on a social conception of the latter)?

The argument is developed in three stages. I begin with the most fundamental discussion over whether or not the concept of knowledge is definable and social. I agree with "knowledge first" theorists, such as Michael Welbourne and Timothy Williamson, that the concept of knowledge is basic, indefinable, and irreducible (Welbourne, 1986, p. 14; Williamson, 2002, p. v, pp. 3–5). I think it is a strength of Williamson's view that the characterization of knowledge as the most general *factive* attitude allows us to understand in what way the concept plays an important social role (Williamson, 2002, pp. 39–40). That is because the single most important obstacle to appreciating this role is the view that the psychologically "narrow" or "internal" mental state of *belief* is conceptually prior to knowledge and forms part of the correct analysis of the latter (Welbourne, 1986, pp. ix–x, 6, 44). At the same time, Williamson's focus is on the metaphysical nature of knowing, and not on the social nature of knowledge. His insistence that knowing is a factive mental state is compatible with, but not necessary for, the social view of knowledge which I defend according to which it is a socially produced status.

The traditional idea that the concept of belief is prior to the concept of knowledge has the implication that any subsequent account of the social role of the concept has to proceed on the basis that it is only "weakly" social – only the common property of individuals in the sense that any two individuals can token the same belief type and content. By contrast, Welbourne argues that it is distinctive of knowledge that it is "essentially commonable" (Welbourne, 1986, p. 1). I argue that this insight can be combined with those of other epistemologists who also investigate this social role, namely, the genealogists for knowledge and truthfulness – Edward Craig and Bernard Williams (Craig, 1990; Williams, 2002). Their work gives insight into the social role of the concept of knowledge and into the social system in which it is embedded – Williams calls this the "belief-assertion-communication" system (Williams, 2002, p. 84).

This leads on to the second stage of my argument where it is argued that existing approaches to *testimony* are also disadvantaged by an assumption of the priority of methodological individualism. That methodology reduces the explanandum, which in our case is (the acquisition of) knowledge, to a claim about individuals (Longino, 2002, pp. 47, pp. 150–151; Thomas, n.d.). That is why the literature on testimony restricts itself

to two questions: is testimony a source of knowledge reducible to ordinary induction or not? And: is there a special problem of credulity in the acceptance of testimony? It is argued that even the most sophisticated evidentialist and assurantist views are committed to this individualism (e.g. Faulkner, 2011, p. 4; Hinchman, 2005; Moran, 2006; McMyler, 2011). That is why my approach, which develops Welbourne's view that knowledge is essentially commonable and "may often be had for the asking" (Welbourne, 1986, p. 1), makes progress by abandoning this individualism in a way that is more radical than Welbourne himself advocated (Welbourne, 1986, p. 44; see Kusch, 2002, pp. 60–62).

Testimony can be transmissive of knowledge between individuals, but I argue that it can also be generative as is argued by Martin Kusch (2002) who makes the same criticism of Welbourne. The role of testimony in generating knowledge is part of why "the problem of cooperation" (Faulkner, 2011, p. 4) which Paul Faulkner casts as "the problem of explaining the rationality of testimonial cooperation" does not arise (Faulkner, 2011, p. 7). Instead, in so far as we have a problem which requires cooperation and trust, the problem is that of achieving and sustaining knowledge through testimony and of validating only those claims which properly should qualify as knowledge. I argue that it is the social generation of knowledge which requires cooperation amongst all who engage in testimonial practice and that such cooperation requires that testimonial practice be structured in terms of trust.

The point that trust has a role because, and in so far as, there is cooperation which is needed to secure knowledge and the point that the knowledge system is a mutual assurance game have been made by Williams (2002, p. 89) and Thomas (Thomas, n.d.). That the aim of testimonial cooperation is knowledge has not been recognized, so I argue, by Faulkner.[1] Because Faulkner focuses only on testimony and not on knowledge, his ethical norms for testimony, including those of trust, can be equated with proper norms for communication (Faulkner, 2011, p. 172). Indeed, Faulkner reduces the ethical trust-based norms of testimony to those given by Grice for successful communication (Faulkner, 2011, pp. 180–181; Grice, 1989, pp. 26–30). I argue that this is a mistake: the norms of trust in testimonial practice, viewed as a knowledge-generative practice, are to be spelled out by the principle of *epistemic justice*. So, I do not think the "problem of cooperation" as Faulkner has described it arises, and I also do not endorse Faulkner's (dis)solution of it.

However, I think Faulkner is right in arguing that current accounts of testimony do not capture the idea of testimonial "uptake". Epistemologically individualistic accounts of testimony lack space either for reasonable trust or for acquiring knowledge on the basis of trust altogether. He is right that what we need is a notion of uptake that satisfactorily explains why we should exercise rightful trust (Faulkner, 2011, pp. 11, 19).

---

1 As I explain in Chapter V section 5, Faulkner attributes to Williams a reductive theory of testimony, and describes Williams's aim as that of providing a theory of how testimonial knowledge is possible – as opposed to Williams's actual aim, which is to show how knowledge is possible.

This is another way of saying that the current theories of how testimony works do not connect truth and trust in the right way.

That connection, I argue, is an ethical connection. However, it is a connection that relates to knowledge, and not to communication. To achieve it, we need more than ethical norms to regulate cooperative communication. We need ethical norms to regulate the successful generation, validation, and dissemination of knowledge, which are, so I argue the norms of epistemic justice (cf. Fricker, 1998/2011, 2007). That ethical norms are needed is explained by Williams's genealogy of truthfulness (Williams, 2002, Ch. 3–5). And those norms are an important part of explaining both the possibility of having knowledge on trust and of why critical uptake is required. They explain why there is always a standard in play that one's trusting another has to be rightful or deserved. I will explain this ethical dimension of knowledge in Chapter III when I discuss Bernard Williams's genealogy of truthfulness as an important corollary of Craig's genealogical account of knowledge. There I also argue that the needed norms require an underwriting that is broader than truthfulness, On the one hand, this is so because our intuitions of truthfulness otherwise get stretched too much. On the other hand, it is so because our social epistemic practices have a political dimension brought out by the work of Miranda Fricker (1998/2011, 2007). So, I argue that the needed norms are underwritten by a positive conception of epistemic justice (see Fricker, 1998/2011, 2007).

In the third stage of my argument I develop an account of what is ethically and politically required for a positive theory of epistemic justice. I consider what norms would follow from epistemic justice to regulate our epistemic behavior and what social conditions are needed to sustain this. Epistemic justice, then, conceived as a positive, normative principle can be used to guide and evaluate our epistemic conduct, but also to critically assess our social epistemic structures. Investigating what conditions are needed for the norms to be sustained, I look at the arguments presented by Helen Longino, who argues that in order to generate knowledge an epistemic community has to be sufficiently inclusive and diverse (Longino, 1990, 1995, 2002). This provides a starting point for what a society well-ordered by justice will look like. It does not provide a comprehensive theory of justice, nor even of epistemic justice for the reason that epistemic justice is only one form of justice. On my pluralist view of justice, achieving justice in one sphere requires that other forms of power, and hence other spheres of justice, are upheld too. This is a point also defended by Michael Walzer who proposed a pluralist theory of justice, even as his theory – which centers on power as a social good – is purely distributive in its focus (Walzer, 1983, pp. 10, pp. 19–20).

I argue that epistemic justice is a vital element for any society that aspires to be a just one. This is because it is part of what justice requires and because, as Miranda Fricker points out, injustices of a non-epistemic kind can be addressed only when those who are mistreated or who suffer injustice of a structural nature are heard (Fricker, 2013, p. 1324). This not only requires that people enjoy some basic epistemic power. It also

requires that the concepts needed for people to express themselves are available to them (Fricker, 2007, pp. 150–152).

While I believe it is important to spell out the notion of epistemic justice in positive terms, I also recognize that both "knowledge" and "justice" are terms that may be appropriated by the socially powerful, who tend to have the most influence epistemically, too, in determining the criteria for the use of such terms. So, there must also be constant vigilance when it comes to the application of the criteria that are used to assess knowledge claims and indeed vigilance with regard to claims that something is a matter of (epistemic) justice (Young, 2000; Douglas, 2009, p. 128). On my view, epistemic justice means that all human beings can contribute to the shaping of social meanings and that all putative sources of information are included, i.e. that all are able to influence what comes to pass as knowledge – whereby the criteria for the relative amount of authority someone has, to wield more influence than others, should be epistemically as well as morally grounded. Furthermore, these criteria themselves are to be regarded as open to contestation.

Chapter I | Knowledge, its Definition and its Sociality

Introduction

This chapter explores the question of what the sociality of knowledge amounts to and how it is incorporated in our understanding of the everyday concept of knowledge. It does so for two reasons: first, in order to speak of – let alone to try and establish – knowledge sharing as a putative moral norm, we need to know what knowledge is. Second, if it is true that knowledge is social, then – depending on what that means – we may find that there is an argument that proceeds from the nature of knowledge to the existence of a norm that knowledge should be shared.

I will begin by examining how theorists have approached the concept of knowledge in order to understand its nature. Given that the concept of knowledge has given rise to a number of competing accounts, I will examine which of them is compatible with the idea that knowledge is social in a minimal sense, and which of them allow knowledge to be social in a much stronger sense, and whether or not there is a good case to be made that knowledge is, indeed, social in this much stronger sense.

My aim in this chapter is not to explain and contribute to the extensive literature in epistemology that followed Gettier's seminal paper arguing that justified true beliefs may fail to be knowledge (Gettier, 1963). Rather, my aim is to show that there is a paradox concealed in the very idea of knowledge as something which individual humans possess and, simultaneously, as something objective. That problem seems to me deeper than the problem which is raised by claiming that knowledge is social, even as it is generally regarded a more pertinent problem when that point is made. For the idea that knowledge is social appears to directly contradict what the objectified account of knowledge tries to salvage: that truth would be reduced or relativized to human needs or interests because, and in so far as, knowledge is relative to a community.

In recent years, however, it has become increasingly common in epistemology to acknowledge that "knowledge is social" (Welbourne, 1986, p. viii; Kusch, 2002, p. 1; Goldman, 2011, p. 3). At the same time, especially given the apparent contradiction between subjectively "possessed" knowledge and the objective truths to which knowledge is answerable, it is important to inquire what is meant by such a claim. And here the meanings of the claim that "knowledge is social" diverge in particular cases, such that it remains, as yet, unclear what the sociality of knowledge means. Below is a

list of possible meanings, some of which overlap, which I have ordered on a scale from low to high contestability, starting with the rather uncontroversial at 1 and ending with the more radical claims (especially 11 and 12, which I do not myself endorse):

1. Knowledge is only properly ascribed to humans (Craig, 1999, p. 35; cf. Foucault, 1972; cf. Lackey, 2008; cf. Popper, 1979, p. 115).
2. Scientific knowledge is social "in the uses it serves", which are that "it is available to use in support of other theories and hypotheses and as a basis for action" (Longino, 1990, pp. 75–76; Popper, 1979).
3. Knowledge is socially generated (Kusch, 2002; Longino, 1990, p. 69, 2002, pp. 128–129; Popper, 1979).
4. Knowledge is commonable, i.e. essentially shareable (Kusch, 2002, Longino, 1990; Welbourne, 1986, p. 1).
5. Knowledge is a social concept, it relates to a social phenomenon (Craig, 1990, p. 3).
6. The standards for ascribing knowledge are (in part) socially generated (Craig, 1990; Kusch, 2002; Longino, 1990, 2002).
7. Knowledge is a social status: it is neither a natural kind nor an artefact (Kusch, 2002; Longino, 2002; Welbourne, 1986; cf. Kornblith, 1999, p. 161; Kornblith 2002, p. 11, p. 29).
8. Knowledge requires discursive interaction (Kusch, 2002; Longino, 1990, p. 216; Longino, 2002, p. 129; Mill, 1989).
9. Knowledge is (typically) had by a community of knowers: while an individual may have knowledge, the typical subject of knowledge is a community (Kusch, 2002; Longino, 1990, 2002; Welbourne, 1986).
10. Knowledge of any community is local, partial and perspectival (Haraway, 1988; Longino, 1990).
11. Knowledge is relative to a community (Kusch, 2002, Part III; see Longino, 2002, p. 138, where she explains why this is *not* her view).
12. Knowledge is a matter of social convention.

This list is neither complete nor exhaustive of the meanings that can be given to the idea that knowledge is social. I am providing it only so as to indicate that, when we say that the social nature of knowledge has been acknowledged, it is not clear what claim is being made. The claim that knowledge is social can have a minimal interpretation that knowledge may be passed on or generated socially even while it is done so via individual beliefs, or a much more radical one, that knowledge can only be generated and is typically had by a community. Kusch and Longino allow that individuals can know, even if the knowledge is socially created/validated/sustained.

The claim that knowledge is a *social concept* needs some explanation about what claim

is being made. If the claim is about the concept having some social use, such as referring to something which we want to be able to say, then it seems trivially true: then every concept is social. This has led Craig to make the point that this truism is not what he means (Craig, 1990, p. 3). If the claim is about the concept's referent, then we are really talking about knowledge as a social phenomenon, and then it would appear somewhat misleading or inaccurate to claim that the *concept* of knowledge is social. Why should we stress that point?

Here, it is important to realize is that conceptualizing a social phenomenon requires more work than conceptualizing a natural kind phenomenon. First, the conception of a social phenomenon is more elusive. In trying to understand the concept of a social phenomenon, such as the concept of knowledge, we try and spell out what aspects and aims of social practice it captures (e.g. Welbourne, 1986, pp. 75–76; Craig, 1990, p. 16). Second, the question of what conception of knowledge, understood a social phenomenon, is correct is more important. For that conception may *redefine* the phenomenon, even if the purported purpose is to describe it – e.g. Welbourne (1986, p. 65), who complains of "conceptual adjustments". Craig acknowledges this, to some extent, when he claims that the purpose of a practice can help us define the extension of a concept in cases where that extension is unclear (Craig, 1990, pp. 2–3). Even so, Craig's purpose is a descriptive one: to explicate the concept we actually have (Craig, 1990, p. 3). This is something he emphasizes, because the method of explication was invented by Carnap who used it for normative purpose – of constructing the notions fit for a unification of science (Craig, 1990, p. 8).

On this first point, while I think conceptual analysis always had as its aim the clarification of practice, namely of how we – correctly – use a term, in order to help us sustain and question our specific applications of it, the realization that knowledge is a social phenomenon leads to expectations that should be lowered and a method that cannot claim to be as rigid as it was once believed to be. It is not just that capturing the relevant aspects of a social practice that institute a social phenomenon is hard. It is that it is not possible to codify the concept, while room will continue to be needed to re-appraise and contest the meaning of a concept, as it may have changed over time.

On the second point, we find that when we try and capture a notion that relates to a social phenomenon, our task is inherently not only descriptive but also normative: the emphasis placed on the importance and relevance of particular aspects of the practice has an impact on what the phenomenon is conceived to be. It shapes the criteria for its application. Accordingly, and importantly, the task of "defining" a concept, when it concerns a social phenomenon, is not a technical analysis, of merely describing norms. It is, instead, *already* a normative enterprise where not only linguistic intuitions are attended to but whereby also, and more importantly, values and social interest or stakes are involved.

What I argue, throughout this thesis, is that a re-appraisal of the type of concept that

knowledge is, whereby it is a social concept in that it relates to a social phenomenon, has deeply normative consequences. That it should have these consequences is clear, for one thing, from the perceived threat posed by the idea that the concept of knowledge could be understood in a way that leaves out the requirement that it pertains to truth, or that objectivity is required for knowledge (Williams, 2002, p. 5, p. 59, p. 65, and p. 147).

The idea that knowledge is social indeed invites thoughts that knowledge is relative to a community (Kusch, 2002, pp. 3–4), or even that it equates to what a community regards as knowledge (Schantz, 2011). In such a reading, truth is relativized to a community, too. It constitutes a danger to view knowledge in this way because it validates claims made by those who claim truth is on their side too. They do so while they seek to push their agenda, not only by weakening the standards for the denial of "other people's" truths, but also because it serves their promotion of "alternative facts" where truth is not the only relevant criterion, even as they may claim that their opposing claims are to be regarded as equally true and more apt. When knowledge is conceived as relating to a social phenomenon, then, scientific work as well as everyday claims invite questions as to who may define the concept of knowledge, and about the proper moral and epistemic conduct required to arrive at knowledge. My account of sociality of knowledge will try to avoid the pitfalls inherent in this reductive view of knowledge as strongly social – a reductionist view.

2. Is Knowledge Social?

It could be argued that the sense in which knowledge is social is a very minimal one with which no one need disagree. Indeed, many extant theories – and typically those which proceed by identifying conditions for knowledge ascriptions – focus on whether or not some person, A, knows that p (Goldman, 1967; Dretske, 1971; Nozick, 1981, p. 170; Sosa, 2007). Bernard Williams accused epistemology in this vein as too focused on what he called "the examiner situation":

> (In) the examiner situation … I know that p is true, this other man (sic) has asserted that p is true, and I ask the question whether this other man really knows it, or merely believes it. I am represented as checking on someone else's credentials for something about which I know already. That of course encourages the idea that knowledge is belief plus reasons, and so on. But this is far from our standard situation with regard to knowledge; our standard situation with regard to knowledge (in relation to other persons) is rather that of trying to find somebody who knows what we don't know; that is, to find someone who is a source of reliable information about something. (Williams, 1970/1973, p. 146)

I will return to the prescience of Williams's final sentence, as developed in Edward Craig's genealogy for knowledge, later. I will focus here on what Welbourne also calls

"belief-theories" of knowledge (e.g. Welbourne, 1986, p. ix–x) that tend to focus on the individual knower by making knowledge a function of individually obtained justification (see Welbourne, 1986, p. 5; Longino, 2002).

This family of approaches may well be social in some quite minimal sense, such as that we must have at least have an ascriber and a person to whom knowledge is ascribed. According to Schmitt, this is an assumption to which all varieties of individualism can subscribe (Schmitt, 1999, p. 355). Further, one can advert to the fact that knowledge is primarily ascribed to a human being (see Craig, 1990, p. 35), while the very idea of ascription entails that knowledge is a status which one person receives from another – and rightly or wrongly so. At the very least, then, *that* social element, i.e. the idea that knowledge is a status in some sense, even if it may not be reduced to purely a social status, is implied by any theory which seeks to analyze knowledge through whether or not it would be "properly" ascribed – even if we assume that "properly" relates to social norms that regulate knowledge ascription within a community, such that these are based on the criteria which the theory sets out to uncover. My interest, however, is whether or not there is a good case for a stronger conception of sociality: one that could ground a more substantive norm than this minimal and uncontested (and uncontestable) sense of "sociality".

3. Current Theories of Knowledge and Their Aim

On an orthodox approach to epistemology it is the task of philosophers to find individually necessary and jointly sufficient conditions for knowledge (Welbourne, 1986; Craig, 1990; Strawson, 1992; Greco & Sosa, 1999; Williamson, 2002, p. 5; Henderson & Greco, 2015, p. 3). This aim characterizes the approach to epistemology – and philosophy more widely – that has been called the project of analysis. It proceeds by considering the "application conditions" of a concept in order to try and define it via its component parts. The proper application conditions are arrived at by reflecting on linguistic use and linguistic intuitions, which are "tested" in different circumstances, using a variety of cases and circumstances, to validate the hypothesis that the concept is defined by the conditions we would intuitively think successfully describe and predict its application (Craig, 1990, p. 1).

In the case of the concept of knowledge, too, philosophers have sought to analyze the concept into component parts. In the orthodox terms, they have been using intuitions regarding the concept's extension to capture the concept's intension; this is to say, that by considering intuitions concerning that to which it refers, or more generally its application conditions, we try and discover the intuitive meaning of the term.

The task of analysis is primarily seen as demystifying the concept through formulating

its correct analysis. This is the project of the "JTB" epistemologist, concerned to solve the equation that knowledge is "justified true belief" plus an additional factor that makes this conjunction of conditions explanatory or non-accidental (Craig, 1990, p. 45; Henderson & Greco, 2015, p. 3). Plausible further conditions include a tracking condition, as on Nozick's account, whereby beliefs "track" the truth, which means they are subjunctively sensitive to it (e.g. Nozick, 1981 p. 178; Dretske, 1983) or a safety condition, according to which the belief is not one that easily fails its aim of being true, or whereby there is no "veritic luck" involved when the belief is true, as Sosa and Pritchard respectively explain it (Sosa, 2007, p. 25; Pritchard, 2003, p. 119). The identification of this further condition is a response to the challenge posed by "Gettier cases".

Edmund Gettier famously argued that justified true belief does not suffice in every case for knowledge as there can be counterexamples where it is merely accidental that the belief is both justified and true (Gettier, 1963, p. 123). In such cases, the relation between the person's justification for the belief (which we would accept as a good justification) and its being true nonetheless appears wholly accidental, making it inappropriate to ascribe knowledge in such a case. So, as Gettier argued, a belief being justified and true is not yet sufficient for knowledge unless the truth and the justification are connected "in the right way" (Nozick, 1981, pp. 169–170; Pritchard, 2003, p. 113).

The recent history of epistemology has seen many philosophers attempt to identify some condition that could fix this problem such that the relation between the justification of a belief and its truth is no longer accidental in a way that bars knowledge ascriptions (Nozick, 1981, p. 179; Dretske, 1983; Prichard, 2006; Sosa, 2007; Greco, 2009). But it seems to me that no such condition has been found that would be true without circularity (i.e. without incorporating the concept of knowledge itself), on the one hand, and that would also withstand more elaborate "Gettier-style" counter-examples, on the other (see Craig, 1990, pp. 51-52, Ch. 9; Miracchi, 2015; Williamson, 2002, pp. 4–5, p. 30).

The failure to find necessary and sufficient conditions for knowledge ascriptions has led some epistemologists to accept contextualism about knowledge (DeRose, 1995; cf. Lewis, 1996, p. 550, p. 564). According to contextualism, it depends on the context of a given case whether or not knowledge ascriptions are true because the epistemic standards for knowledge can vary depending on the circumstances (DeRose, 1995, p. 4). The circumstance that some knowledge is crucially important in a given case, such that a false knowledge ascription would have serious consequences, raises the bar for ascribing knowledge to someone (DeRose, 1992, p. 915). The idea of contextualism is that these contextually variable standards can be used to explain why intuitions about knowledge ascriptions can differ. On their view, Gettier cases make salient certain special features of the context, which we ordinarily do not take into account, such that we decide against a knowledge ascription which we ordinarily would have given (Lewis, 1996, p. 557).

Hence, contextualists can vindicate the tripartite analysis, but at the cost of introducing variable standards. These standards have the implication that one may lack

knowledge when relevant circumstances of a case are being made salient – or when the stakes are raised (DeRose, 1995; Williamson, 2000; M. Williams, 2001). In my view, the contextualist approach foreshadows two aspects of the problem of capturing the social nature of knowledge depending on how one considers what the contextualist approach centrally involves.

On the one hand, one may think that Gettier cases are avoided by contextualism only because the sceptic is granted to be right to some extent: we cannot get at the nature of knowledge via its ascription. For, if one wants to view knowledge as a phenomenon that is not itself contextual, then one's response to contextualism is to grant that knowledge ascriptions are not an adequate basis for understanding the concept of knowledge. Then one may opt to give up the project of linguistic analysis – and aim to get at the phenomenon of what knowledge is more directly, instead. The (consistent, complementary) alternative way to put this is to say that in so far as the contextualist is right, what we seek to know is not about *knowledge* but only about *knowledge ascriptions and the correct (social) criteria we set for them.*

On the other hand, one may accept that the contextualist provides a true account of the nature of knowledge. But if one does that, then one accepts that knowledge is relative. For, we would then equate knowledge ascriptions with knowledge and accept that a knowledge claim is true in one set of circumstances and false in another. Moreover, if knowledge entails truth then it appears that we are committed to saying that P is known and hence (regarded as true) is true in one set of circumstances, while that P is known (and hence is true) is false in another set of circumstances.

Given these two options, it appears we can say that contextualists – in so far as we find their account convincing – show that *either* we should give up on conceptual analysis of the concept of knowledge via knowledge ascriptions (and focus instead on the latter) *or* if we think it is an actual analysis of knowledge that we can arrive at in this way, and hence that knowledge is contextual itself, then we should accept that both knowledge and truth are relative. I would argue for the former, not the latter, if I were to have to choose. But neither option is satisfactory: a focus purely on knowledge ascriptions, would move us too far away from the normative question, since knowledge is only properly ascribed when someone possesses it.

But what is interesting, and my reason for making this point, is that the contextualist gets us to these conclusions about knowledge by focusing – albeit indirectly – on the role of the knowledge ascriber and the social conditions that pertain to knowledge claims, such as what is at stake for the humans who seek to acquire knowledge or validate a knowledge claim. The social dimension, and the reflection on the (social) circumstances when a knowledge claim is made, i.e. the act of appraisal, so I would argue, is what makes contextualism susceptible to the same kind of dilemma that also surfaces when we seek to understand the social nature of knowledge.

Another way to try and avoid the Gettier problem – and arguably its opposite, as

it is metaphysical as opposed to social – is to abandon the internalist requirement on justification; to go descriptive rather than normative, when trying to capture the notion of knowledge (even while recognizing that it is a success term and that it is a normative notion for which we try and codify or capture the criteria. This avoids our problem, as the idea is to provide an account that explains that knowledge ascriptions are only rightly made when true, i.e. when the justification suffices for the knowledge claim from some external perspective to suffice for the knowledge claim.

Nozick developed a paradigmatic form of this alternative when he developed a truth-tracking account of knowledge (Nozick, 1981, p. 179). Nozick argued that a person has knowledge if he or she would "track the truth", in the sense of his or her belief being counterfactually sensitive to falsehood (Nozick, 1981, p. 178). But since there is always some far-fetched possibility of falsehood, Nozick claimed that on his truth-tracking account of knowledge, a person should be sensitive only to *relevant* alternatives in order for their true belief to be knowledge (Nozick, 1981, p. 175).

This is a subjunctive conditional of the form "if p were true, then q would also be true" is true when in all close by possible world where p is the case, q is the case (Nozick, 1981, p. 173). As such, the subjunctive conditions of his truth-tracking account have a criterion for relevance "built-in"; and as such, Nozick claims that they may be used to provide that criterion for alternative causal accounts of knowledge (Nozick, 1981, p. 175). His reason for resisting causal accounts, however, is that they do not readily accommodate mathematical knowledge and ethical knowledge given that in these cases there is no apparent causal relation (Nozick, 1981, p. 170). The truth-tracking account fares better as there is no direct causal connection assumed, but only a counterfactual sensitivity between world and belief (Nozick, 1981, p. 173).

An important point about Nozick's account is that the relation of tracking the truth, between the truth of P and the subject's believing that P, is that it is relative to the method used. Even while Nozick's account is externalist, he makes clear that not all ways of coming to hold a belief will be as good, in terms of satisfying the subjunctives, as others. The example he gives is that of a (forgetful) grandma who may be told that her son is alive by being told, even while she would be told the same when he is not, while she could come to know the same fact by seeing her son.

At first, Nozick's account appears promising in circumventing the Gettier problem as it arises for JTB accounts. Its power appears to be that it is descriptive in its spirit, while an externalist notion of justification is at work. The descriptive aspect is that Nozick provided a subjunctive criterion for when knowledge claim would be appropriately made, namely when it is made in the right type of world/circumstances, such that the justification – or rather the method of acquiring the knowledge in question – is only relevant to the extent that it has delivered the individual with the belief (and may be deemed better or worse for the fact of its tracking the truth or its failing to do so).

Yet, the problem this stipulation gives rise to, I argue, is circularity of the form any

reductive analysis of the concept is supposed to avoid. This circularity arises because it has been stipulated by the Nozickean account that P is true while the tracking conditions state that the person who has knowledge that P *only believes that P is true under the right circumstances*. But one cannot rule out that there are circumstances where he person whom we take to know that p does not know it according to Nozick's tracking conditions, unless we restrict what alternatives are relevant. An example given by Nozick explains this problem.

So, consider the "Great Bank Robbery" case of a man who robs a bank wearing a mask, but is recognized by a bystander from "wanted" posters as Jesse James, because the mask slips. That Jesse James is the culprit would not have been known if his mask had not slipped, which is arguably a close possible world. Nozick has argued, in response to the case, that in a world where the mask doesn't slip, recognizing Jesse James is not the *method* we would use for gaining knowledge of who the culprit was. So, he believes that the apparently close by possible world can be dismissed as a relevant alternative – which is needed if the recognition of Jesse James is to qualify as a way of arriving at knowledge, which intuition about the concept's application Nozick wants to validate (Nozick, 1981, p. 193).

Yet, Craig argues that this is a "fudging" maneuver on Nozick's part (Craig, 1990, p. 22). There is a difference between the method we use and the evidence it yields, as was pointed out by Forbes (Forbes, 1984, pp. 47–48). The method we did pick, and the (quality of the) evidence *it* yielded, is relevant for assessing whether we have arrived at knowledge. Accordingly, as Craig rightly argues, the scenario of the mask not slipping *ought to be irrelevant*. And so there is something wrong with Nozick's account if the scenario is relevant for assessing whether the testifier has knowledge, because in a close by possible world the method would not have been available.

On Nozick's account, the alternatives which we may disregard because they are irrelevant are so in virtue of our (antecedently) accepting that it is indeed knowledge that we may ascribe in such cases. So Nozick's account of relevance presupposes an account of that which is already known. Nozick's account does not, then, illuminate the nature of the concept of knowledge, as this concept is invoked *in order* to decide on the *relevance* of alternatives.

Similarly, the contextualist response to Gettier cases either yields to skeptical intuitions that no knowledge is had in some given circumstances, or yields to dogmatic conviction that there is knowledge in a given case, because we ascribe it. It cannot provide conditions for knowledge which are independent of the intuitions we have that are part of a prior grasp of the concept of knowledge – even if we may point to high stakes in a case where knowledge is lacking, just as we may, on Nozick's account, that in a case were knowledge is had we can discard a world that appears relevant by saying that it is not in fact close by, due to the method not being available there. These two strategies, then, for renovating the tripartite analysis of knowledge do not succeed in doing so without circularity.

Now, circularity is only a problem, one may think, if the project is to find necessary

and sufficient conditions for knowledge. If the idea is to model the ascriptions in a way we find roughly acceptable, as a descriptive exercise, circularity need not be a problem: the model may help to isolate some aspect of the concept, e.g. simply the appropriate conditions for knowledge ascriptions, better. Yet, the objection of the Jesse James case goes some deeper: the account has to yield the right outcomes; even if it features as a proxy model, it has to be acceptable as the best in range. Here Craig claims that his genealogical account of knowledge can more aptly capture (if not predict) what properly qualifies as knowledge and what does not. Indeed, his account of knowledge is one which is key to an understanding of the role of knowledge and its social nature, so I argue in the next chapter. As a prelude to that discussion I will describe another form of skepticism towards the JTB approach that seems equally opposed to the project of genealogy but which I will argue is complementary to it. That is Williamson's "knowledge first" epistemology.

4. The Knowledge First Approach

This approach holds that the concept of knowledge is unanalyzable, even if further aspects of our use of the concept can be illuminatingly characterized (Williamson, 2002, pp. 3–5). Williamson writes:

> The working hypothesis should be that the concept *knows* cannot be analysed into more basic concepts. But to say that is not to say that no reflective understanding of it is possible. (Williamson, 2002, p. 33)

Instead of explaining what knowledge is by analyzing it in supposedly more basic terms, Williamson wants us to focus on elucidating the nature of knowledge from a different perspective. Williamson's approach to understanding the concept of knowledge is a metaphysical one. He analyzes knowing as a state of mind which ought to be central to the philosophy of mind – as well as to epistemology (Williamson, 2002, p. 6).

One motivation for Williamson's dissatisfaction with analyses and for his slogan "knowledge first" is akin to Craig's and Welbourne's, whose works I discuss later. As Williamson explains, an analysis, especially one that is complex enough to deal with Gettier cases – if it were possible to give it – does not help us achieve a better understanding of the concept (Williamson, 2002, p. 31). In understanding it, we seek to understand what role the concept has in our lives, e.g. why it should enjoy such a widespread use in our lives, and that project goes beyond a codification of its criteria (Welbourne, 1986, p. 35; Craig, 1990, p. 2).

Similarly, Williamson notes that we may be puzzled as to how, if an elaborate analysis is given, it can be at all a match to the concept as we tend to think of it, given the importance we attach to it – why it is so important to us that some statements are qualified

as knowledge while others are not? Williamson comments as follows:

> Even if some sufficiently complex analysis never succumbed to counterexamples, that would not entail the identity of the analysing concept with the concept *knows*. Indeed, the equation of the concepts might well lead to more puzzlement rather than less. For knowing matters; the difference between knowing and not knowing is very important to us. Even unsophisticated curiosity is a desire to *know*. This importance would be hard to understand if the concept knows were the more or less ad hoc sprawl that analyses have had to become; why should we care so much about *that*? (Williamson, 2002, pp. 30–31)

The importance of the concept – its role in our lives – however, is lost in analyses given of the concept, which do not show why we should care to have knowledge as opposed to justified true belief (Williamson, 2002, Preface p. v). Indeed, a consequence of such analyses is that it appears that from an epistemic point of view, knowledge – whether you know or not – does *not* matter (Kaplan, 1985).

In *Knowledge and its Limits*, Williamson builds on an externalist understanding of mental contents. His argument is that mental states (or more generally, mental attitudes, both stative and non-stative), can be externally individuated, too (Williamson, 2002, p. 6, Chapters 1 and 2). Williamson argues that adopting the view that some mental states are externally individuated makes it tenable to claim that knowing is a mental state, namely a factive one.

He makes clear that that his claim is that there is a mental state being in which is both necessary and sufficient for knowing (Williamson, 2002, p. 21). He thereby denies the orthodoxy that knowing must be understood in terms of the mental state of believing (Williamson, 2002, p. 21). On that orthodox view, knowing cannot be a mental state, because it is "factive" in that "truth is a non-mental component of knowing" (Williamson, 2002, p. 22). That truth can be part of the individuation of the mental state is Williamson's externalist stance on the nature of the mental attitude in question.

While it challenges Cartesian internalism in the philosophy of mind, Williamson's position can also be viewed as a way of countering Karl Popper's claim that scientific knowledge and common knowledge are categorically distinct.[2] (Popper, 1979, pp. 108–109). This is noteworthy, because the rejection of a categorical difference between what may be called "scientific knowledge" and "everyday knowledge" is also a natural conclusion of philosophers who maintain that knowledge is social (Longino, 2002, pp. 2–3). Of course, the denial of a categorical difference is not the denial that there may be a larger and more robust amount of evidence backing what may be deemed "scientific knowledge". So, even those who maintain that our conceptions of scientific knowledge and our everyday knowledge are continuous may nevertheless hold that there is a difference between them. They can add that claiming scientific knowledge requires a certain kind of understanding (Welbourne, 1986 p. 53), or more methodological rigor as the criteria for knowledge are set by the scientific community (Longino, 2002, p. 162), or that the standards of accountability are higher when scientists claim knowledge, because scientists enjoy greater authority and

are better able to foresee the consequences of possible errors (Douglas, 2009, pp. 82–83).

Williamson shares Popper's view that, on the everyday conception of knowledge, it is understood as an attitude or mental state ascribed to (individual) humans (Popper, 1979, pp. 108). However, unlike Popper, he does not think that because and in so far as it is conceived as mental attitude, knowledge must be subjective, i.e. non-objective (Popper places knowledge in his "world three" of abstract and objective entities (Popper, 1978, p. 144-145; Popper, 1979, pp. 73-74, pp. 106-107). Accordingly, he manages to explain why we care for having knowledge, even given the everyday sense in which knowledge is a mental state that is ascribed to us. The explanation for why we care is that it is a mental state that links us to truths and hence means that if we are in a state of knowing, we possess truths:

> On this account, the importance of knowing to us becomes as intelligible as the importance of truth. Factive mental states are important to us as states whose essence includes a matching between mind and world, and knowing is important to us as the most general factive stative attitude. (Williamson, 2002, pp. 39–40)

In making this point, Williamson also accords a special role to knowing. Of all the factive mental states, it is the most general one, i.e. one that is entailed by other factive mental states, such as perceiving, remembering, and hearing (if the latter is a direct perception of an event) (Williamson, 2002, p. 34, p. 37). Theoretically, then, knowing takes a pivotal explanatory role in explaining the mental, and our actions (Williamson, 2002, p. 7 and chapter 3).

Even so, Williamson also claims that we do not always know whether or not we actually have knowledge: the state of knowing may be indiscriminable from the "inside" from a state which resembles it (Williamson, 2002, pp. 24–26). Admittedly, these two positions in combination superficially yield a paradox. The paradox is that if knowing is a state which we cannot internally ascertain, then the fact that the mental state of knowing has a factive nature, i.e. that it links us to truths, does not help us to explain why we should care for it (Williamson, 2002, p. 34). For it could be that we merely *think* that we are in that mental state.

That is, the "non-luminous" nature of a state of knowledge makes knowledge – or the state of knowing – seem less relevant for us. Yet, Williamson claims that even while it is not the case that we always know that we know, we care to know precisely because it is metaphysically a different situation from the one which obtains when we fail to know, i.e. when we are wrong or ignorant (Williamson, 2002, p. 34). In addition, so he argues, we may *know* that we know, even if he thinks that this will not always be the case (Williamson, 2002, p. 23). He argues that typically we do not know through introspection what state we are in when it concerns a mental state that is non-trivial (Williamson, 2002, p. 25, Ch. 4).

So, on the one hand, Williamson argues that knowing that one knows (KKp) is not a requirement for knowledge, which means that he is not susceptible to skeptical arguments

that undermine our existent knowledge – as, Williamson claims, Descartes erroneously thought when he confounded subjective and objective certainty (Williamson, 2002, p 23). On the other hand, Williamson argues that his account *is* compatible with our being in a state of knowing even when we cannot internally distinguish that state from a state of believing, as a disjunctive analysis for veridical perceptual states shows (Williamson, 2002, p. 44; see Child 1994, pp. 143–64; Dancy, 1995; Martin, 1997; McDowell, 1982; Snowdon 1980–1, 1990; cf. Hinton 1967, 1973). That second reply means that the paradox is less pertinent: it explains that the importance of knowledge is not undermined by arguments that seek to show that we can never ascertain that we know "from the inside".

Williamson observes that the idea that knowing is factive, while belief is not, and that belief is best explained by reference to knowledge, parallels disjunctive approaches in the theory of perception, where we understand the non-veridical cases in terms of the veridical ones (Williamson, 2002, p. 48). For Williamson, this is a desirable result because he contends that knowing is the most general factive mental state, i.e. the state which is entailed by all other factive mental states (Williamson, 2002, p. 34).

Perceiving is also, on his view, a factive mental state, but it only delivers one kind of knowledge: perceptual knowledge. Remembering is a factive mental state that yields another kind of knowledge. Just as belief is a mental process that aims at knowing that p, then, other mental processes aim at factive mental states too: perception aims at perceiving that p, memory aims at remembering that p (Williamson, 2002, p. 48).

Examples of mental states are: believing, desiring, feeling in pain (Williamson, 2002, p. 27). But not everything qualifies as a mental state. A revealing contrast, Williamson claims, is between knowing and believing truly (Williamson, 2002, p 27). While knowing is a mental state, there is no state picked out by believing truly. The latter would be a state with this feature: if a person is "in" this state, this would be necessary and sufficient for believing truly. There is, of course, the mental state of believing, but there is no mental state "in between" believing and knowing such as "believing truly"; the latter does not exist as a mental state (Williamson, 2002, p. 27). As he later argues, believing truly is composite and therefore not a mental concept: it is part mental and part non-mental (Williamson, 2002, p. 28). He maintains that mental states are conceptualized as such, i.e. that they are captured by at least one mental concept: "a state is mental if and only if there could be a mental concept of that state" (Williamson, 2002, p. 28).

While knowing is factive, it is nonetheless captured by a mental concept; it is not a mental component plus something external added to it (Williamson, 2002, p. 28). Knowing is what he calls "prime", a non-composite concept (Williamson, 2002, p. 66; in *Chapter 3 Primeness*). The alternative conception of knowledge fits with an ordinary "pre-theoretical understanding" of its nature, as we do consider the state of knowing to be a cognitive state (Williamson, 2002, p. 22). As Williamson emphases, this alternative to the JTB-analysis also has a history in philosophy that predates Gettier style counterexamples (cf. Pritchard, 1950; Williamson, 2002, p. 23).

A central tenet of the view that knowing is a mental state is that even if the mental state of knowing is introspectively indiscriminable from the mental state of believing, there is still a difference between these two kinds of mental states as they possess distinct individuation conditions. Knowing that p is a "factive mental state", which entails the truth of p; in contrast, believing that p is a non-factive mental state (Williamson, 2002, pp. 21–22, p. 34).

In arguing that knowing is a mental state Williamson aims to rid epistemology of an undue burden and rehabilitate the importance of its subject matter (Williamson, 2002, p. 5). The concept of knowledge, according to Williamson, is an important focus for understanding the nature of mind as it has an ineliminable explanatory role when we seek to understand the relation between world and mind (Williamson, 2002, p. 1). Knowledge (or better, for Williamson: the mental state of knowing) is the basis of action (Williamson, 2002, pp. 7–8; and chapter 3), questions of evidence (Williamson, 2002, p. 9; and Ch. 9), and judgment and assertion (Williamson, 2002, pp. 10–11; and Ch. 11). This is Williamson's take on why knowledge *matters* (Williamson, 2002, p. 31).

There is no need to postulate that the concept of knowledge builds on the concept of belief metaphysically. Williamson illustrates this last point as follows:

> It may be a priori that being crimson is sufficient for being red, but that implication need not be explained by an analysis of one colour concept in terms of the other. One can grasp either concept without grasping the other, by being shown examples of its application and non-application. (Williamson, 2002, pp. 43–44)

So being a specific kind of red may be considered sufficient for an object to be red, but it is not necessary to invoke the one concept to make people grasp the other concept: what is needed is some experience with how these concepts are applied. This position fits with Christopher Peacocke's theory of concepts, which is that the nature of a concept is given by a correct account of how a user who has mastered the concept would apply it (Peacocke, 1992, p. 5).

As Williamson points out, to suppose that knowledge can be analyzed in terms of belief presupposes that believing is somehow more basic than knowing, i.e. that belief is conceptually prior to knowledge. For one can only understand an analysis if one has a prior grasp of the concepts which feature in the analysis. But he asks "Why should we suppose that belief is conceptually prior to knowledge?" (Williamson 2002, p. 2). There are no arguments for this view, apart from the fact that knowledge entails belief. However, this entailment is compatible with its being the case that we understand the concept of belief in terms of (or with reference to) the concept of knowledge. The "knowledge first" claim, understood in this way, is the opposite of that which JTB-analysts assume: we understand the concept believing by reference to that of knowing, not vice versa (Williamson, 2002, p. v, pp. 2–3).

The "knowledge first" approach implies a reorientation of our philosophical

understanding of the relation between mind and world. It takes knowledge as the primary notion to explain the direction of fit from world to mind – while previously belief had been regarded as playing this role. Belief, in turn, is understood as aspiring to knowledge.

Williamson maintains that knowledge, as the object of philosophical study, should not be seen as belonging exclusively to the domain of epistemology. If knowing is a state of mind it also, forms the subject matter of the philosophy of mind – which focuses on the relations between mind and world. Here Williamson maintains that philosophy of mind "cannot afford" to ignore knowing as a mental state as it is in fact at the core of its subject matter (Williamson, 2002, p. 6). On the other hand, he admits that the epistemological questions surrounding knowledge remain important. That is, epistemological concepts such as "justified", "caused" and "reliable", which do not feature in Williamsons' account of knowing as a mental state do have important connections to knowledge (Williamson, 2002, p. 41).

As Williamson recognizes and stresses "any adequate account of knowledge should enable one to understand these connections" (Williamson, 2002, p. 41). Knowledge features in philosophical accounts of belief, action, evidence or justification, and assertion. He argues that knowledge cannot be eliminated from an explanation of action over time (Williamson, 2002, pp. 7–8). Further, he argues that "knowledge is the evidential standard for the justification of belief" (Williamson, 2002, p. 49) and that "the fundamental rule of assertion is that one should assert p only if one knows p" (Williamson, 2002, p. 11). These explanatory roles of the concept he views as supporting his argument that knowing is a mental state.

5. Knowledge First and the Social Role of Knowledge

While the philosopher most readily associated with the knowledge first approach is Williamson, one of the first to advocate the primacy of knowledge, and hence a "knowledge first" approach was Michael Welbourne (Welbourne, 1986). Welbourne's work is significant for this thesis because he combines the intuitions that motivate the "knowledge first" approach with a conception of knowledge as strongly social. Welbourne also held that the belief-condition on knowledge is a mistake (Welbourne, 1986, p. ix). Importantly for my argument about the social nature of knowledge, he also thought that the belief-condition barred our understanding of *the way in which* knowledge is social (Welbourne, 1986, p. ix). For Welbourne argued that the social practice of testimony is central to an understanding of the social nature of knowledge; I think that is a genuine insight that I will be developing in this thesis.

While people had started to recognize the social aspect of knowledge in the sense that

knowledge "seems to get about in conversations and through newspapers, books and so on" (Welbourne, 1986, p. ix) and had also started to recognize the importance of that aspect for understanding knowledge, Welbourne argued for a more radical suggestion. On his account, knowledge is "essentially commonable" (Welbourne, 1986, p. 1). By this he meant that it is in the nature of knowledge that it is "transmissible or communicable" (Welbourne, 1986, p. 1). What Welbourne thus stressed as a crucial feature of the concept of knowledge, which often goes unrecognized in philosophy, is that knowledge "may often be had for the asking" (Welbourne, 1986, p. 1).[1]

While he recognized that the genesis of some knowledge may be hard and require effort, his claim is that when something is knowledge it can be acquired without much effort: "knowledge *as such* is not hard to come by" (Welbourne, 1986, p. 1). This is a radical break with concurrent thinking about knowledge because we tend to think that knowledge is hard to come by and that it requires effort on the part of someone who acquires it. And indeed, even as we recognize that knowledge may be spread from one party to another, we tend to want to hold on to this idea that knowledge is something which an individual acquires and for which they should use their own "on-board sources" to do so (Craig, 1990, p. 11).

Welbourne pointed out that once we take an individualistic notion such as belief as the core notion of our conception of knowledge, it is hard to understand how knowledge can be transmitted from one person to another. This is the essential point that connects his "knowledge first" and social intuitions concerning knowledge. He explains the individualistic nature of belief by the fact that a belief is conceived of as a psychological entity that is "in" the individual mind of the person who has it. As such, two beliefs can have the same propositional content, but two people are not seen to have "the same belief" (Welbourne, 1986, p. 6).

The two beliefs differ, because belief is individuated, at least in part, as the psychological state of the person to whom it pertains. Yet, so Welbourne noted, the common conception of knowledge as a person's belief which meets some certain additional requirements, builds on such individually held beliefs, which are only type-identical and not token-identical.[2] Welbourne pointed out that this makes for an important disparity between knowledge and belief, since we regard two individuals as holding two different beliefs, with the same propositional content, but we do not speak

---

1 Welbourne's work may be regarded as defending a theory of testimony, rather than as a theory of knowledge. The point which Welbourne makes, and which I draw inspiration from, is that we should reconceptualize knowledge by better understanding the role and nature of testimony. Moving beyond Welbourne, I argue – similar to Kusch (2002) - that this requires we reconceptualize testimony as *generative* of knowledge, in Chapter IV and V.

2 According to Welbourne, this leads to problems in explaining not only how individuals can confer their own knowledge to others, but also in understanding how there can be a common knowledge at all. Indeed, as Welbourne sees it, the notion of objectivity can't be explained when there is no way to share a common outlook (Welbourne, 1986, p. 6). The best there is, is convergent beliefs. But how can such convergence be achieved? He thinks it is impossible for belief-theorists (those who hold beliefs are the basis of knowledge on the part of individuals) to explain how the ideal of objectivity, even as an aim, could have arisen at all.

of their having two "knowledges" (Welbourne, 1986, p. 6). Instead, two people are said to share the same piece of knowledge, when they both know that p. Welbourne uses this linguistic fact to argue that it is prima facie plausible that knowledge is not the "possession" of any specific individual: it is *commonable* (Welbourne, 1986, p. 6).

Welbourne further notes that this picture of knowledge as acquired and possessed by individuals is associated with classical empiricism. For the empiricists, notably Locke, it was important to stress precisely that each individual needs to understand the content of a judgement in a way that grounds the knowledge (Welbourne, 1986, p. 51; cf. Locke, 1788). This has generated a psychological, subjective condition on knowledge, regardless of whether that was Locke's intention. In addition, Descartes' quest for certainty yielded a demand for self-sufficiency – both epistemic and ethical (Welbourne, 1986, p. 6; Zagzebski, 2012, Ch. 1).

Here the idea that humans are fallible led to the idea that one should rely only on oneself, as one can ensure one's own methods for acquiring knowledge are reliable, while others may be sources of error (Hume, 2000, Section X). Our concept of knowledge became bound up with a questionable ethical ideal of self-sufficiency (see Taylor, 1995, p. 8).

One idea taken from Descartes' *Meditations* is that in order to obtain knowledge, one must try and minimize the chances of going wrong: to undergo a "once and for all" project of acquiring all and only true beliefs (Descartes, 2010; Williams, 1978). To minimize the chances of error, one must thus avoid depending on others, who could not only willingly lie and deceive, but may also be incompetent and unreliable as a source. The upshot is that a person must ascertain for himself or herself that some potential/purported piece of knowledge properly qualifies as such in order to come to also possess the knowledge which someone claims to have ascertained and would try to "share".

So, in his denial of the individualistic nature of knowledge, Welbourne seeks to make room for deference to another person as an epistemic authority. That is, Welbourne holds that knowledge can be shared or transmitted without any epistemic work being required on the part of the receiver of a testimony – aside from believing the speaker (Welbourne, 1986, p. x, p. 5). This "aside" is an important one. As Welbourne makes clear, the hearer is to exercise his judgement and can "withhold belief when he is told something" (Welbourne, 1986, p. 66).

Arguably, it makes the receiver responsible for acquiring knowledge after all. Indeed, I will maintain that the hearer should be critical, a point which Welbourne also stressed in the passage quoted (Welbourne, 1986, p. 66). Here, it is tempting to take Welbourne's point that the hearer should be duly critical as *shifting* the entire epistemic responsibility from the speaker to the hearer (unless one regards the hearer's "believing" the speaker as merely a moral matter, in which case the epistemic responsibility remains with the speaker). However, the very idea that we should accord

the full epistemic responsibility either to the speaker or to the hearer, is something which may be denied on a social conception of knowledge, as I will argue in Chapter IV and V.

That is, the idea that epistemic responsibility is shared appears to follow rather directly from Welbourne's work. For Welbourne maintained that an essential feature of knowledge is that it *passes as knowledge* – or at least that it *can come to pass* as such. On his view of what its essential "commonability" comes to, the concepts of knowledge and testimony are closely linked, and there is, strictly speaking, nothing puzzling about gaining knowledge through testimony, because of this link: knowledge is commonable by its nature.

Yet, in so far as Welbourne only commented on cases of "telling" – i.e. of knowledge transmission, where it is assumed that the speaker should speak from knowledge – Welbourne's account of testimony would appear to be a classic version of assurantism inspired by Austin (Austin, 1946). This is somewhat striking, for the reason that assurantism has often been presented as engaged with a puzzle which has an individualistic conception of knowledge as its premise. For it offers a "non-reductive" solution to questions relating to the status of knowledge as a self-standing source of the warrant necessary for testimonial knowledge (whereby what is being explained is how knowledge can be obtained by one individual from another).

A question then is why Welbourne's account of testimony is not more "radical". For, if passing as knowledge within a community is required for knowledge, as Welbourne claimed, it would appear logical to maintain that testimony is the practice whereby we test the claims that are putatively knowledge for their commonability. That is, testimony can be seen as a practice, consisting of multiple testimonial exchanges, whereby claims that deserve to pass as knowledge are socially validated as such. The view that testimony is a knowledge generative practice is held by Martin Kusch – who appears right to claim that he is taking Welbourne's views, about the nature of knowledge and the centrality of testimony, to their logical conclusion, that:

> (…) pushing Welbourne's analysis to its natural conclusion, and thus further than he seems willing to go, (…) 'knowledge' is a social status like money, and thus it only exists in so far as there are items upon which we are willing to impose the status. And the collective imposition of the status happens in and through local transmissions, first of knowledge-claims, and later of accredited knowledge-items. (Kusch, 2002, p. 62)

Now, while one way to make sense of Welbourne's position is Kusch's, another way is to take seriously that knowledge is held, by Welbourne, to be indefinable, even while the conception of knowledge may be elucidated through the other concepts to which it is linked.

On the score of elucidating knowledge, Welbourne writes, first of all, that he believes the concept of knowledge and belief have evolved simultaneously with the system

of communication (Welbourne, 1986, p. 4). Welbourne also links knowledge, more specifically, to the idea of objectivity. Hence, he writes, on the one hand, that:

> Our idea of knowledge, to put it roughly, is the idea of communicable information, information as to the facts, information which is objective in the sense that it is not dependent on any particular point of view, but is available to anyone at all, with the capacity to understand the utterances in which it is embodied. (Welbourne, 1986, p. 5).

He also maintains, more radically, that:

> The commonability of knowledge is one and the same with the possibility of a common, public and objective world.
> (Welbourne, 1986, p. 6)

What is meant in this last passage is brought out more clearly by his rejection of a solitary knower as prototypical. We arrive at our conception of knowledge and of truth, so Welbourne claims, because we engage in (testimonial) "practices of telling, accepting what we are told and so on" (Welbourne, 1986, p. 84). There, so the thought is, we have to verify claims and decide whether or not to accept them. Importantly, Welbourne thinks – and this is the point that Kusch alludes to – that a given piece of knowledge requires a community (that recognizes it as knowledge in testimonial practice):

> The supposed solitary knower whose say-so is not accepted is rather like a man who with a £10 note in a community which has no use for money.
> (Welbourne, 1986, p. 84)

However, Welbourne is not equating knowledge with a social status: he maintains, merely, that there are conceptual links which we should take seriously, and that it is senseless to try and understand these concepts outside the practices that embed them. So the citation continues:

> It is senseless to insist that what he has is worth £10 in his context. And it would be equally pointless to insist that solitary Percy knows, when his say-so is universally rejected.
> (Welbourne, 1986, p. 84)

While Welbourne refused reductively to define knowledge, it may be thought that he did not take his own view to its logical conclusion, as we just saw. As Martin Kusch argues, the implication of Welbourne's view is that testimony is not only transmissive of knowledge but also generative; it is when a bit of testimony is accepted that it comes to pass as knowledge, while, that it passes as knowledge and can continue to pass as such is a test for its truth for it is deserving of that status. But this is not a definition of knowledge; Welbourne, like Williamson, can consistently deny that the knowledge is definable at all (Welbourne, 1986, p. xiii, p. 85). Thus, in Kusch's view, individual bits of knowledge are generated by being socially accorded the status of knowledge. He views the idea that

knowledge is essentially commonable as the idea that it is a necessary condition for knowledge that it passes as such (even if not a sufficient one), where Welbourne left this somewhat vaguer; and saw "commonable" as a term indicating potentiality, as opposed to actuality.

In his own, social-communitarian, account of knowledge Kusch spells out a conception of testimony as a practice. He rejects the individualistic view that testimony is a one-off encounter between two individuals, which analytic philosophers have taken it to be, and which is based, so he claims, on a legalistic conception of testimony (Kusch, 2002, p. 16). The idea that Kusch defends is, moreover, that it is a practice that it is *generative – or potentially degenerative* – of knowledge claims. A charge against Kusch has been that he is changing the subject from testimony to something else (Lackey, 2006, p. 36).

Admittedly, Jennifer Lackey makes this point mainly because she believes that performative speech acts do not belong in the domain of testimony; because they are in the business of creating facts rather than (transmitting) knowledge. While it would indeed be questionable if we create knowledge by announcing it, the mechanism of performative speech acts is used, by Kusch, mainly to illustrate the role testimony plays in conferring *epistemic statuses* to claims. This is, I think, a point Lackey's argument fails to see: while speech acts create status, as well as contents, *the former* is the focus when we are explaining the role of testimony in generating knowledge. Here, quite notably, creating a status is in large part a performative act – even if one should heed the right criteria for doing so. Hence, it is not to be equated with the idea that one creates warrant or reasons to believe as assurantists have it, even if one may give others reasons to believe, indeed, by offering assurance, but that is something else from one's implicit endorsement of a claim when accepting it and when using it in reasoning and when passing it on to others. Notably, Lackey herself claims that "the words of speakers, not their beliefs, are what matter for an epistemology of testimony" (Lackey, 2008, p. 71). That point is entirely compatible with the point Kusch makes, if one understands one's use of words – and one's passing on certain claims - as the implicit endorsement of these claims as knowledge.

I believe, however, that redefining the subject in the case of testimony may indeed be necessary if we want to free it from the individualistic assumptions that lie behind recent discussions of testimony – such as the issue of whether testimonial encounters instantiate the vice of credulity. The questions that testimony yields are ones that do not arise on a social conception of testimony, as Welbourne claimed. Of course, Welbourne's merely claiming that point is not enough. To show why the problems do not arise, one needs to make plausible why that is so. This includes, to my mind, a need to provide an error theory of why such problems might have been thought to arise, and what proper place in epistemology can be reserved for the concerns that sustain that line of thought.

In so far as Welbourne was less radical then Kusch, when presenting his account of knowledge, we as yet lack an idea of why and how the puzzle of testimony can be

dissolved by a social conception of knowledge. This too, is why Welbourne's conceptions of knowledge and of testimony are my starting point, but not the end point. For indeed, if testimony is central to knowledge, we should understand why this should be so, and in what sense this is so.

Kusch's idea that testimony is generative of knowledge is what provides the missing connection here, so I believe – even while I do not think we should then understand knowledge in terms of agreement in a reductive way. That is, while we can explain the social status of knowledge by reference to the generative nature of testimony, we cannot explain the generative nature of testimony through claims about the social nature of knowledge. For that to be possible, we would need to know that knowledge is a social status – a conceptual point as opposed to a practical one.

On my view of testimony, the speaker and hearer are each responsible for their roles. They together determine what qualifies as knowledge – given that knowledge is social and hence any item of knowledge must be such that it is accepted by others as such, for only if it is accepted as knowledge can it be shared, and without being shared it cannot qualify as knowledge. This is so even if we would want to see that all and only truths are accepted as knowledge, and hence that all claims that are shared as if they are knowledge properly qualify as such.

So, I will argue that the responsibility of the hearer is to be critical in his or her uptake: believing a speaker is a moral-epistemic matter. At the same time, I will argue that the claim – held by, amongst others Elisabeth Fricker and Williamson, and contested by Williams – that it is a precondition for any successful testimonial exchange (which could be defined narrowly as an exchange that leads to shared knowledge between speaker and hearer) that a speaker should speak "from knowledge" is mistaken (E. Fricker, 2015, p. 59; Williamson, 2002, Ch. 11; cf. Williams, 2002, Ch. 4, sections 3 and 4). Rather, at best, the speaker should speak from actually possessing a truth. On my view it will take away too much of the responsibility for knowledge to say the speaker already should properly qualify a claim as knowledge.[3]

Attributing or allowing the status of knowledge to a proposition is a responsibility that must be shared between speaker and hearer. I will argue that a testimonial exchange is dynamic in its nature, and that both parties to it have responsibility for the ensuing knowledge between them. As I see it, testimony does not "transmit" knowledge. It transmits knowledge claims, and depending on the veracity and convincingness of these claims and how these claims are received, it generates knowledge (Kusch, 2002, p. 18,

---

3 At least, in the context of the social process of the validation of knowledge claims, which is part of knowledge generation. There may be a point where some claim is transmitted as knowledge whereby it is the sole or main responsibility of the speaker that what they present as knowledge – and hence as sufficiently socially validated – is indeed knowledge. In such a case, the hearer is off the hook, even as there remains room for the hearer to be critical and investigate the claim made nonetheless. If there were no such cases there would be (*almost) no use for authority and no point in talking of testimony as a form of transmission (as well as a generative source of knowledge claims and of knowledge). * I write "almost" as some sources may be deemed more reliable than others and thus be given a heavier weight.

p. 62; Lackey, 2008, p. 71). I will discuss these issues further in chapters IV and V below.

On Welbourne's view then, knowledge is essentially non-individualistic. A piece of knowledge (and not merely its propositional content or the propositional content of an individualistic belief) can be shared (Welbourne, 1986, p. 6). So, according to Welbourne there is nothing puzzling about testimony. The question how it is that someone can gain knowledge on the basis of mere "say-so" is not a problematic one. It is seen to be problematic only when we fail to pay proper regard to the social nature of knowledge: when we fail to see it is "essentially commonable". Mistakenly, then the (sole) focus in attempts to analyze knowledge has been on third person singular ascriptions of the verb "to know", targeting (only) an individual's personal psychological state and the conditions which that state needs to meet to qualify as a state of knowing (Welbourne, 1986, p. 73).

The social dimension of knowledge is relevant if we are looking at a concept where the social dimension is not merely a fact of us using the same term, but where the applicability of the concept (after we have fixed its meaning) requires additional social judgement. There are, after all, cases where it is debatable whether or not something is knowledge. And these debates are substantive: they are not debates about the interpretation of some social or linguistic norm, but about the social status being conferred upon some proposition or not.

6. Knowledge First and the Practice of Knowledge

Both Welbourne and Williamson have one striking view in common: that there is a knowledge-norm of assertion, according to which those who speak informatively should do so from a position of knowledge (Welbourne, 1986, p. 19; Williamson, 2002, pp. 10–11, Ch. 11).

But here it should be noted that we may not know that we have knowledge: for whether or not something is knowledge is partly determined by whether it is accepted (on Welbourne's view). Likewise, Williamson holds that we do not necessarily know that we know, because whether or not something qualifies as knowledge depends on the way the world is given that it is a *factive* mental state. He takes pains to argue that even while knowledge is a mental state it is not "luminous" – i.e. not self-evident which state we are in, one of knowing or one of believing – while he also seeks to make clear that we can know that we know, in quite a number of cases.

These conceptual links, however, from knowledge to testimony and to objectivity and community and to intentional action and to what qualifies as evidence, and the various other roles that they attribute to knowledge, do not yet help us to see the point of the concept of knowledge, i.e. why we have the concept of knowledge in the first place. For that we need an explanatory story about why we should have the concept; what its main

function is.

Here the genealogical account by Craig helps (Craig, 1990). And, interestingly, it is an account that emphasizes our epistemic co-dependence, much along the lines Welbourne suggested. For on Craig's view, the concept of knowledge arises out of the situation that we seek to acquire truths from others, and for that reason, have a need to identify reliable sources (as Williams implied in his critique of the "examiner situation"). Knowledge is not defined by Craig and he is clear that a functional definition ought to be resisted. However, the concept is explained by reference to a hypothetical function that originated it: namely, to signal the status of reliability, which is a social status in the sense that the criteria for it are socially set, even as these criteria should be set such as to serve epistemic success.

A thing to note about Craig's genealogy, which I discuss in the next chapter, is that it does not provide a definition of knowledge even if there may be an urge to try and define knowledge functionally on the part of Craig's critics. But the idea of a genealogy is that we can use it to show why concepts that appear not to have a function, and which are valued intrinsically (for their own sake), in fact have a hidden function that comes into view by reflecting on their (speculated) origin. It is because the function of such a concept is not apparent that a genealogy does not – and cannot – provide a fitting functional definition while the concept may also have evolved further such that it has features and functions that do not bear on its speculated reason of coming into existence – though, for a genealogy to have some plausibility, the concept's proposed initial purpose should make sense, and should help us explain how the concept that we currently could have could have evolved from it.

In Craig's original conception of the concept there is a link not only to social reliance but also to testimony, something which is clear from the fact that he chooses as the proto-concept which would arise given the need for people to acquire more truths than they can themselves muster, the concept of "a good informant". This is someone who is able to inform another person reliably and truthfully and who is accessible to that other person (Craig, 1990, p. 85). While a good informant is also someone who is to be readily identifiable as such, we may also trust others to be better placed to identify those who qualify as a good informant, in cases where we do not know who to turn to – or need to be referred to someone who knows the truth about a given matter. This is why the concept of knowledge arose, according to Craig: to make such referrals possible (Craig, 1990, p. 88).

It is also why the criteria for being a knower are set high "so as to satisfy the highest denominator" but are not set so impossibly high that we would all need to be sceptics (the fact that the concept exists means that it is to have a use) (Craig, 1990, p. 91). What is also important to realize is that the criteria which are set for the most reliable sources, as well as the markers that serve to identify the most reliable sources – the "indicator-properties", so to speak– are socially determined, even as those properties must correlate

with the informant's being a highly reliable source.[4] That being so, these criteria may, potentially, be the wrong ones, if we fail to identify the sources that are reliable as to the truth on the given, particular, matters. Indeed, as Miranda Fricker has pointed out, the fact that indicator properties will be used can also mean that the criteria come to be rigged by the socially powerful (Fricker, 1998, p. 62).

But even while there may be moral and political questions that surround the proper assignment of rational authority to people, that point should not deter us from recognizing the social nature of knowledge as Welbourne rightly insisted. He argued that problems which are of an ethical nature – e.g. related to a threat of authoritarianism, which might arise from a conception of knowledge which gives deference to authorities a role – should not be resolved conceptually. To do so is in fact, according to Welbourne "positively dangerous" as it obscures the ethical issues (Welbourne, 1986, p. 65).

On Craig's genealogy of knowledge the point that the concept of knowledge serves is to help a community achieve epistemic success, i.e. truths. Here, however, it is worthwhile to point out that Craig does not equate knowledge with a community's epistemic success – nor does Welbourne. Instead, both recognize that a community may set the wrong criteria and may, thereby, come to regard claims as "knowledge" while those claims fail to be true. A point that needs to be emphasized, then, is that the genealogical account provided by Craig does not eliminate or relativize the (need for) a concept of truth.

Conclusion

If it is true that knowledge is social, then, can we provide a philosophical account of the concept of knowledge that can make sense of that claim? As I have argued in this chapter, the social nature of knowledge as part of some definition of knowledge appears to lead us to problem(s). The main concern is that it would appear to relativize knowledge to a community and hence curtail the idea of objective truths. Another concern is that it is not possible to determine when knowledge is had: when knowledge depends on social practice as well as epistemic criteria, it becomes a matter of degree whether or not some item is knowledge in – if not "for" – a given community. Even so, there are reasons to think that knowledge is social – quite aside from the ethical/political aspects that can be explained when that thesis is true, and that we can make sense of if and when knowledge is social.

In this chapter, I began my argument by addressing the question of whether an approach completely opposed to the project of analyzing knowledge as justified true belief – the "knowledge first" approach of Timothy Williamson – can incorporate any

---

4 The point of correlation is stressed by Craig, but is in fact a criterion on the indicator property we set. He does explicitly alert us to the social nature of the indicator property.

acknowledgement of the social nature of knowledge. To show how it can, I started with a philosopher who reverses the order of those questions: Michael Welbourne who argued, first, that knowledge is essentially "commonable" before going on to establish conclusions that anticipate, in many respects, key claims of knowledge first epistemology. In the next chapter I will explain why there is no inconsistency between the position of these "knowledge first" epistemologies and the genealogical approaches of Craig (1990) and Williams (2002).

Chapter II | A Genealogy for Knowledge

Introduction

As I noted in my opening chapter, the traditional analysis of the concept of knowledge has been rejected by two groups of epistemologists: the "knowledge first" epistemologists, Welbourne (1986) and Williamson (2002), and the genealogists for knowledge, Edward Craig (1990) and Bernard Williams (2002). I have also noted two interesting aspects of the way in which these rejections of "JTB" epistemology interact: that there is no inconsistency between claiming both that the concept of knowledge is unanalyzable and yet also has an insightful genealogy; and that there is no inconsistency between being a knowledge first epistemologist and emphasizing the social function of the concept of knowledge. It is, indeed, Welbourne's combination of a knowledge first approach with an acknowledgement of the strongly social nature of knowledge that makes him so important for the arguments of this thesis.

The point of a genealogy of knowledge, as opposed to an analysis, is to ask this question: suppose that we postulate the existence of a group of people very similar to us in respect of their characteristic human capacities, and suppose, also, that they live in a "state of nature" with needs and interests similar to our own: what would lead these people to introduce the concept of knowledge, were they to lack one?

The point of the genealogical inquiry is to investigate why the concept of knowledge is useful and important to us and why it is thought to play a central role in our lives. Craig's "state of nature" story starts out from the viewpoint of an individual inquirer (Craig, 1990, p. 12). Even so, there is nothing in the very idea of a genealogy that fails to acknowledge that this need can be a social need. The fact, however, that a shared social need is central to Craig's genealogical account – while Craig also acknowledges that the concept relates to a social phenomenon – is in my view obscured by the way in which Craig actually formulates many of his arguments in *Knowledge and the State of Nature*. His perspective is – superficially – an individualistic one. This may be explained by the fact that Craig sought to show that his account is plausible. As such, he emphasized that the genealogical account yields a conception of knowledge that vindicates the belief-theoretical analyses that were given – up to a point –and hence do justice to the intuitions that the individualistic belief-theoretical sought to capture. It does this, even as Craig claims his approach can capture and explain more, namely also our conflicting intuitions when it comes to applying the concept in unusual circumstances (Craig, 1990, pp. 2–3, pp. 6–7). But this ought not to be interpreted as some kind of concession that the JTB account can be repaired as being partly true, or true of some parts of our epistemic practices, but not all of them.

At the outset, Craig is non-committal on the question of whether belief is a necessary condition for knowledge while that idea is central for JTB-theorists (Craig, 1990, pp. 12–14). In this chapter I explain how Craig makes a convincing case that sociality is irreducibly part of our understanding of the concept of knowledge, such that any individualistic understanding of the notion has to be abandoned. Given the irreducibly individualistic nature of belief, then it follows that Craig's argument gives us strong reason to abandon any version of the JTB project.

On my interpretation, the outcome of Craig's work is that we should move to understanding the social role of the concept of knowledge and its social nature. After discussing the genealogical method and Craig's genealogical account, I argue that it yields a result that is close to Welbourne's conception of knowledge as "essentially commonable". The plausibility of Welbourne's thesis is greater if it can be shown that Craig's genealogy has provided an alternative route to the same conclusion. I also defend Craig's genealogical account from objections that have been raised to it.[1]

2. The Genealogical Method

Craig's genealogy seeks to explain how we could have come to have the concept of knowledge that we actually have. The genealogy works to illuminate our actual, ordinary, concept of knowledge by telling a fictional story of how a concept that is sufficiently much like our concept of knowledge could have arisen in the first place, in a "state of nature", which tries to assume as little as possible and yet to be plausible in its assumptions as pertaining to any more basic community of people. The genealogy elucidates our actual concept only if the concept which is constructed through the genealogy can then be shown to match it, by yielding intuitions regarding its intension and extension that correspond to the intuitions yielded by our actual concept of knowledge. In Craig's own words:

> We take some prima facie plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what conditions govern its application. Such an investigation would still have an anchorage point in the everyday concept: should it reach a result quite different from the intuitive extension, then, barring some special and especially plausible explanation of the

---

1 Craig's work has been discussed in the Greco & Sosa volume, 1999, by a number of authors, who, however, do not give an in-depth discussion of his work. In the Greco & Henderson 2017 volume, Elisabeth Fricker "reworks" Craig's approach as well as its outcome. In this chapter I will argue her approach is misguided. In that same volume, Michael Williams compares Craig's genealogy with works by Robert Brandom and uses their insights in conjunction. Very recently, a paper was published listing and providing a line of answers to all objections to Craig's method, by Martin Kusch & Robin McKenna (2018). I have not been able to incorporate their points, but it is fair to say that I agree with the points they made; they too have found that there is no incompatibility between Williamson's account of knowledge and Craig's, even as they also note that Craig has "undercut" one main point of Williamson, namely the suggestion by Williamson that his account is the only good alternative to traditional analysis (Kusch & McKenna, 2018).

mismatch, the original hypothesis about the role that the concept plays in our life would of course be the first casualty. For it is not the idea to construct an imaginary concept, but to illuminate the one we actually have, though it be vague or even inconsistent; and to illuminate it by showing that a concept with the hypothesized role would have characteristics closely resembling those that it exhibits itself. (Craig, 1990, p. 2)

The method of a genealogy is not new in philosophy. It is a method that is commonplace in political philosophy (Nozick, 2013). Craig's novel idea was that the method of genealogy may be applied in epistemology to elucidate the concept of knowledge. In political philosophy "as-if" stories are told, which may or may not have historical plausibility, concerning how concepts such as those of the state arose in the first place – where it is understood, from the start, that the target concepts are social institutions, which is something that people do not realize about the concept of knowledge, namely, that it is embedded in a social institution (if we interpret the word "institution" broadly) (Nozick, 2013; Williams, 2002, p. 31). Craig sought to treat the concept of knowledge as embedded in such a social institution too and see what that kind of understanding that enterprise could bring.

While in political philosophy a genealogy is often used normatively, Craig notes that his use of a genealogy is descriptive (Craig, 1990, p. 9). Craig uses it to validate our intuitions concerning third-personal knowledge ascriptions. At the same time however, his genealogy aspires to do more: it also seeks to explain these intuitions. That is, it seeks to explain why the intuitions we have concerning knowledge-ascriptions make sense if and when they do, what rationale can be given for the intuitions we have regarding knowledge-ascription, and how they vary in different cases.

Notably, in the final step of matching the constructed concept to our actual one, there is room for applying Rawls' methodological step of reflective equilibrium (Rawls, 1951, 1971/1999). According to this method, a philosophical inquiry based on our basic intuitions in which we have a high degree of initial confidence, such as the inquiry into the principles of justice, may be revisionary of our intuitions as well as being descriptive of a rationale that is based on (most of) them. So, if we believe that the constructed concept can be applied more consistently than the actual concept we have, we may adjust the linguistic intuitions we have of our actual concept, even as we seek to construct a concept that makes sense of most of our actual intuitions so as to describe our actual concept. Craig writes:

But should our intuitions prove indeterminate or elastic, this type of investigation might reveal constructive ways of stretching them, and the rationale behind the stretch. With luck it might also reveal the sources of indeterminacy or elasticity which dogged the attempts to answer, or even to ask, the first familiar question. (Craig, 1990, p. 2)

Prima facie, then, understanding the nature of knowledge, through understanding the concept of knowledge in terms of that concept's function makes sense, because and in

so far as the concept does not refer to a natural kind, a rival view defended by Hilary Kornblith (1999, p. 161; 2002, pp. 11, 29), but rejected by a number of philosophers, e.g., Longino (1990), Zagzebski (1999), and Kusch (2002). If the concept of knowledge were to refer to something, X, which is knowledge, then knowing what function the concept serves (namely, referring to object X) would not illuminate the nature of X.

So, Craig's project is predicated on the fact that knowledge is a social and not a natural kind: it is a concept embedded in a social institution or system. Craig writes at the outset that he cannot justify this assumption, other than by dogmatically stating that knowledge is, according to him, a social phenomenon. But he also points out that his method assumes less than the method of conceptual analysis, and he invites us to bear with him, as "the proof of the pudding" is in the eating (Craig, 1990, p. 3). At the very least the concept can be treated as a socially embedded, since it gives rise to, and is shaped by, social needs. On the other hand, especially the final step in reflective equilibrium of matching the constructed concept to our actual one, still assumes that the way the concept is used in our language – by a language user who grasps its application conditions – will help us to understand the kind of thing that the concept is a concept of (Peacocke, 1992).

A question is what makes for a good genealogical elucidation: how we can affirm that it has captured the core of the concept fitting its prototypical extension and intension. According to genealogists, their genealogical story can be useful even if, through historical contingencies and for other reasons, our actual concept does not appear to us to have been called into existence for a social purpose. It may be used for phenomena that are such that we do not think of them in terms of their function (Williams, 2002, p. 31). Craig remarks that if Williams is right in this, then "some functional phenomena must be good at covering their functional tracks, so to speak." (Craig, 2007, p. 197). Both Williams and Craig point out that a genealogy may be used, and is useful precisely in such cases: the hope is, as Craig puts it, that "a genealogy will show us how they do it" (Craig, 2007, p. 197). The way in which a genealogy may be used then, then, is to provide the missing link: a story about the concept's origin that serves to bring out the concept's purpose. A genealogy is particularly suited to do this, as it can do it without reducing the phenomenon, or our conception of it, to the purpose it serves. Hence, the genealogy gives insight without yielding a reductive functional definition of the concept.

Yet, to make the genealogy plausible, we should in such cases also understand through the genealogical story that and why these concepts are so good at not revealing the work they do, the purpose they serve. Williams provided such a story, as he argued that a conceptualization in functional terms of some phenomena can serve to undermine their ability to serve their function. Overt function would, in fact, undermine what these concepts do for us. This is true if and when the concepts are to be valued intrinsically as opposed to instrumentally. So, he maintains that, in the case of

justice, knowledge and truthfulness, we actually need a genealogy to understand their nature, because only a genealogy can elucidate them without reduction or unwanted instrumentalization – a point to which I will return (Williams, 2002, pp. 35-36).

The point where Craig most clearly deviates radically from the JTB- theorists is that he sees himself as discovering truths about what he calls a "prototypical notion" (Craig, 1990, p. 15). Accordingly, Craig is not attempting to give a list of necessary and sufficient conditions, which would yield an account that should work in every case – i.e. accommodate all our intuitions about knowledge ascriptions. (And which would, in turn, be vulnerable to counter-examples.) Instead, proceeding from considerations about the point of the concept, Craig seeks to describe the concept's core function, such that it understood that variations of the concept's application can occur which may also be explained by reference to the purpose the concept serves. Craig aims to illuminate the concept by explaining both why the concept applies when it does, and not when it does not: "We are asking not so much: when is the ascription of a certain concept correct, but rather, why is it applied?" (Craig, 1990, p. 14).

This approach allows for variations in the concept's application conditions, as long as these can be explained by reference to its purpose. Importantly, such variations are not to be seen as counterexamples to an analysis of the concept. That is because the idea that all of the conditions characterizing the concept must apply in every case is dropped. Craig's reasonable plea is that in exceptional cases or as he puts it "in freakish circumstances" the concept can behave differently in order to serve its purpose (Craig, 1990, p. 14). This means that cases that qualify as counterexamples to any putative conceptual analyses are not necessarily problematic for the genealogical account.

Counter-examples are fatal to any analysis, and set off a regressive search for counter-example proof revisions to the analysis, but this problem does not arise for the genealogist. The account is successful, and not undermined by would-be counterexamples, if the application conditions it yields are correct for the typical case and if exceptions to it can be explained. Craig argues not only that we should give up on the method of analysis – at least in the case of the concept of knowledge – but also that we should give up on the pervasive orthodoxy that concepts can only be descriptively characterized in terms of conditions which may each be shown to be necessary and that are jointly sufficient (Craig, 1990, p. 14).

This change of perspective also distinguishes Craig from knowledge first theorists, since his claim is more radical than saying, as knowledge first defenders do, that no good analysis of the concept can be given because the conditions which we believe are necessary for characterizing the concept do not altogether add up to it – i.e. are not sufficient. Knowledge first theorists tend to deny that belief is a necessary condition for knowledge for two reasons. First, we need no prior grasp of the concept of belief in order to understand the concept of knowledge (Welbourne, 1986, p. xiii; Williamson, 2002,

pp. 43–44). Second, the concept of belief is not metaphysically prior to that of knowledge as it would be if belief were constitutive of knowledge as its mental substratum (Williamson, 2002, p. 6). But in saying this, knowledge-first theorists do not yet make the move which Craig makes: to argue that conceptual analysis in terms of necessary (and jointly sufficient) conditions is too rigid since "greater flexibility is required in the description of concepts", and hence that – at least for the concept of knowledge –it ought simply to be dispensed with (Craig, 1990, p. 14).

In saying that he seeks to characterize and illuminate the core of the concept, Craig takes up Wittgenstein's point that in characterizing a concept we should not look for its essence, or essential features, but examine how the concept is used and hence how it features, and what its different features are, in different contexts (Craig, 1990, p. 14; Wittgenstein, 2007 §67). So Craig finds support and inspiration in Wittgenstein's point that a concept is characterized in terms of its features just as a family resemblance is characterized by overlapping similarities of facial features, none of which are individually necessary.

Craig also thinks that a focus on analysis and an obsession with repudiating or accommodating counterexamples has only served to confound our understanding of the concept of knowledge. It has led JTB-theorists to ever more complex analyses that make it hard to see why a concept like that should have such a central role in our lives (Craig, 1990, p. 2). So the way to proceed is descriptively to characterize the concept directly via its role and without a need to regard its nature as if it is unchanging and codifiable by some list of necessary conditions that jointly uniquely capture it.

Admittedly, Wittgenstein denied that there is even such a thing as a concept's core, if that core is understood in essentialist fashion. We should abandon the idea of an essence or essential features. Instead of wishing to characterize the supposed essential core of the concept, we should look at the concept's use (Wittgenstein, 2007 §43). We should view its behavior from one context to another as alike and different in the way that is familiar to us when we compare members of a family, where there are overlapping similarities between faces without any of these being the essential feature or set of features. The idea then is that even as we may wish to characterize a concept's core – its typical way of featuring in most contexts or in the most familiar contexts – we need not view any or all of the conditions that we may find necessary for understanding the notion. Nor indeed should we aspire to achieve a list that gives all the conditions that are jointly sufficient. Craig appears to understand the notion of a core in the same way. For instance, he is prepared to be non-committal about the notion of "belief" at the heart of the concept (Craig, 1990, pp. 12–14). My view is that we can be more committal than that, on the basis of Craig's own genealogy: the essentially commonable concept of knowledge cannot contain belief as an individualistic component.

However, there are distinct ways of accommodating Wittgenstein's point about "family resemblances", as I have noted, and Craig is committed only to one of them. His model of using a genealogy to identify a prototypical "core" of a concept is one way in which to interpret this generic idea. So, the success of his proposal depends not on this abstract

methodological point, but on whether a successful genealogy can actually be constructed for any particular concept.

3. Craig's Genealogy for Knowledge

Craig's genealogy claims that our concept of knowledge exists in order to meet a social need that would arise in any linguistic community, namely a need to identify good informants within that community. The story starts out from a "primitive" community (Craig, 1990, p. 4) – I prefer the term "more basic" – which is advanced enough for its members to share a language, even if the language as yet lacks the concept of knowledge. (The problem with the term "primitive" is that it misleadingly implies that the genealogist is engaged in an actual ethnography or history.) Even so, vis-à-vis evolutionary epistemology, Craig notes that the development of our cognitive abilities will not impact on the nature of the concept. This is because, as he explains, the concept depends on "certain very general facts about the human situation", such that even when we cognitively develop the concept's role is the same (Craig, 1990, p. 10).

The subsequent genealogy has three steps. The first step is the claim that, in any community with a language, there will naturally arise a division of epistemic labor. A corollary of this claim is the sub-claim that in any such community there will be a system for working out who are the reliable informants on whether or not p. The second step is the claim that the concept of knowledge that may be hypothesized to have arisen out of this need. Then, in a third step, it is argued a constructed concept, i.e. a concept that would answer the need we have hypothesized exists, matches up in reflective equilibrium to our actual concept of knowledge.

This final step is important as it is the claim that the concept which the genealogical story yields is sufficiently alike the concept we actually have in order to be able to explain the intuitions that we have about the actual concept's intension and extension. The aim of this, as in the parallel step in Rawls's use of the method of reflective equilibrium, are to show that the genealogy is a non-debunking account of our existing concept and that it gives us insight into it (Williams, 2002, p. 36).

In the first step Craig uses the following assumptions, which are held to be true of any community in which the members share a language to communicate:

a)   Having true beliefs is advantageous for survival, and

b)   Members of the community have a need for more true beliefs than each can acquire individually, with their onboard resources, e.g. perception, reason.

c)   All members have some true beliefs which others lack – based purely on their situational circumstance, such as their position.

These assumptions together are enough for the reasonable speculation that:

H1: (Reciprocal) information sharing is advantageous.

Notably, it is not clear who shares with whom, or for what reasons information is shared. It is assumed however, as we saw in (c) that all members can contribute to sharing information, if they wish to do so. So, according to Craig,

H2: Within the community there is a cooperative enterprise of increasing the number of true beliefs, in which all members of the community may take part.

The concept of knowledge then arises because it is reasonable to assume further that:

H3: Not all informants are equally reliable with regard to a given matter.

This last assumption (H3) coupled to the second hypothesis (H2) will mean that the community has to develop a mechanism for filtering true from false beliefs. Thus, Craig concludes:

There is a need to evaluate sources of information, i.e. to 'flag' reliable sources as 'approved sources of information. (Craig, 2002, p. 11)

It is in order to meet this need that a new concept is required. That which arises in the State of Nature, Craig claims, is the concept of a good informant. The concept of a good informant as constructed from the viewpoint of the individual inquirer is as follows:

(..) as well as his having the right answer to my question,
1. He should be accessible to me here and now
2. He should be recognizable by me as someone likely to be right about p.
3. He should be as likely to be right about p as my concerns require.
4. Channels of communication between him and me should be open.
     (Craig, 1990, p. 85)

In order to show how this concept could have evolved into our concept of knowledge, Craig emphasizes the stage of objectivization, as part of concept formation of concepts of artefacts, in general and as such in the case of the concept of knowledge too (Craig, 1990, p. 88).

It is in the process of objectivization that the members of a community converge on using the same concept so that their "individual" conceptions come to have the same extension. The example Craig provides is of the concept of a chair. He thinks that even as we might each have different subjective conceptions of something like a chair, it is when we want to refer to it, so that others may give us something to sit on, there was a need to converge on the criteria for what objects qualify as a chair (Craig, 1990, p. 84).

Craig argues that the concept of a chair too, is primarily understood by us through its function. In the case of the concept of a "good informant", sharing the same concept ensures we can socially refer to reliable informants, such that we can pick out those who are to count as such on behalf of others, and ask others to pick them out for us. The social pressure for an objectivized conception is that others may be in a better position than me to pick out good informants on some particular topic (Craig, 1990, p. 88).

I think is it important to stress that in Craig's account all the subjective aspects of the concept of the good informant are dropped via this process of "objectivization". This helps us to see why the constructed concept can be the concept of knowledge, as we know it, with its demanding standards. For the community comes to have a shared standard for determining who are to count as reliable informants on a given matter. Craig conjectures that the standard for conceiving of someone as a reliable informant will be pitched high, as this person has to satisfy even the most demanding inquirer (Craig, 1990, p. 91). He argues that convergence on a high standard of reliability is a function of the context, namely the fact that it is unknown what is practically at stake for a receiver. So, while there is a contextualist element to Craig's account it is "transcended" in this process of objectification; contrast, for example, the "gatekeeping" contextualism of David Henderson which differs from Craig's view by introducing context-relative standards (Henderson, 2011).

Reliably to recommend a person to someone, not knowing what the stakes are, demands that the standards of reliability are set high. The concept which helps us recognize and flag (or advertise oneself as) a highly reliable informant, then, so Craig posits, is the concept of knowledge. So, the claim here is not that the application conditions of the concept of knowledge are inherently contextually determined. The claim is that the concept of knowledge has come to be as demanding as it is in order to accommodate even the highest common denominator, because the most demanding inquirer plays a determinative role in setting the common standard (Craig, 1990, p. 91). It is, therefore, a contextually invariant view.

Just as condition (3) becomes demanding, Craig notes that condition (2), on the other hand, becomes less demanding, in the sense that there should just be someone who can recognize the informant. It is not necessary that all within the community can recognize the informant, by some property which he should have that makes it likely he has a true belief. Indeed, if all could recognize the informant, then the concept of knowledge would have no use. Here it should be noted that while condition (2) comes to be less demanding because of what we may regard as a division of labor, where only some in the community need to recognize the property of what makes someone likely a reliable source on some particular topic, condition (3) does not rely on this division of labor but, instead, on the "highest common denominator" idea.

The decisive factor which determines how the process of objectivization works out, vis-à-vis these subjective conditions, appears to be whether a community has a use for the

concept. Condition (1) is also dropped in the process of objectivization: the whole point is that we can recommend informants to others for their future reference so they do not need to be in the immediate proximity and immediately available – even if of course we must assume that these informants are available to them at some point. Now, the interesting question for this thesis is how condition (4) can be further explicated.

Should the channels of communication between him and me become a channel between the informant and just anyone? What does a "good informant" in the objectivized sense come to? Does it mean that one need not act as a good informant to anyone, but that one need only act accordingly if one chooses to, to just some people? Or should it mean, as a demanding enquirer may want, that everyone is in principle entitled to know everything at all? Would we have to think along the lines of someone being available at all to at least some one other person, or rather again on the model of what a demanding inquirer may wish for, who would want access whoever he is, whatever it concerns, and whatever the relationship? What Craig has to say on this score appears to be the former claim:

> (T)he question whether or not channels of communication are open isn't a simple Yes or No matter, but depends on facts about the particular inquirer and his relationship to the possible informant. (Craig, 1990, p. 92)

So, Craig does not consider the question of whether or not it would follow from his view that being a knower would give rise to a norm that knowledge must be shared in principle. Arguably, such is the case if the concept is to be useful and that is the argument I will develop in this thesis. Craig appears to believe that it is enough if the knowledge of one person is made available in a kind of "trickle down" fashion via trusted relationship only to those who actually need it. He points out that there are cases, when it concerns confidential information of a student, for instance, where it is enough that the student informs a particular audience whom she trusts, such that this person can then act to disseminate or to just take the appropriate kind of action.

Craig is right of course that much depends on the context and particularly the relationship between the informant and the inquirer. Furthermore, as Craig implies with his example, whether or not channels of communication should be open depends on the type of knowledge in question. However, the suggested answer appears to not get us far enough, as it seems to invoke an exception rather than the general case. It would seem that the genealogical idea of flagging informants as a basic role of the concept of knowledge entails that there should be at least some proneness to inform others when the status of a knower applies to one, at least in the prototypical case – and conversely, that there should be a default expectation that one who seeks to gain knowledge, or who needs it, will (upon asking as well as proactively) be informed by the person to whom the status of knower applies.

What Craig has achieved, rather than giving any definition of knowledge or equating it in terms of another concept, is that the nature and role of the concept can be explained

through a fictive, even if by my lights plausible enough, genealogy. What his genealogy shows, and is important to capture, is that the concept of knowledge is a social concept, because the standards that are embedded in the concept are shaped by social convergence as to what standards are required, for the concept to function. This happens in the stage of objectivization. A separate point is that it follows from the genealogical understanding of the concept that these social standards are upheld through social, as opposed to individual, judgement. This is because according to the understanding yielded by the genealogy of the concept, knowledge is fundamentally a social status.

So, the concept is social in the deeper sense that it refers to a socially constructed reality: that of mutual ascriptions of status. It is social, then, in the strong sense that its application depends on social judgement as opposed to individual judgement. In a later chapter I will explain that this does not yet relieve any of us of our individual epistemic responsibility, nor indeed of understanding knowledge as entailing an objective notion of truth. It is a crucial aspect of my thesis that people have individual epistemic responsibilities.

Paradoxically, the fact that knowledge is essentially social makes it easier to see both that and why we (should be seen to) have such responsibilities towards one another. This is a paradox, because it might appear that an individualistic approach can explain why we should always exercise our own judgement. But such a sweeping claim is simply false. That knowledge is social does not relieve us of our individual duties as members of a community of knowledge, or lessen our responsibility. As Welbourne points out, "the more people who defer to authority the fewer the checks on what is published as knowledge" (Welbourne, 1986, pp. 66–67). Even while it is true that the responsibility comes to be shared, we are also responsible for the epistemic states of our peers. A generative account of testimony, whereby we see how social interactions matter, which I shall provide later, will help explain this, while I will also provide an account of what knowledge requires both ethically and politically in the next chapters.

4. Informants and Sources of Information

In this section, I will consider the issue that Craig chooses to elucidate the notion of knowledge via the notion of an informant, rather than that of information. Why is this so, and what is the distinction?

The distinction, as Craig admits, roughly converges with that of the distinction between humans and non-humans, as people are said to know, and knowledge is not ascribed to inanimate objects, such as trees. The basic idea here is that knowledge presupposes possession of truths by a human being. It is clear that we do not ascribe knowledge to inanimate objects such as trees in other circumstances. Trees can serve,

however, as sources of information from which information is gleaned. Insightfully, Craig notes that, conversely, there are cases where humans are not regarded as informants, but instead only as sources of information. He makes clear, however, that this deprives them of something peculiarly human. For when a person is regarded solely as a source of information, she is regarded as a means, and not as a "cooperative member" of the species, i.e. not treated as an end in themselves (Craig, 1990, p. 36).

This is an important insight which converges with Miranda Fricker's later stance that it is a form of what she labels "epistemic injustice" when a person is "wronged" in their "capacity as a knower", the essential epistemic harm of which she relates importantly to one's human dignity (Fricker, 2007, pp. 44, 54; Thomas, 2018). She maintains, too, that this can happen when a person is treated merely as a source of information, or worse, as not even being that. For this discussion, I focus on the fact that being an informant is something that is predicated of humans, while humans and objects alike can be sources of information, as Craig noted (Craig, 1990, p. 35).

So, the choice of "informant" in Craig's account rather than "information" highlights already that the concept of knowledge, according to him, is rightly only predicated of humans. Craig notes that we don't speak of a tree knowing how old it is, not even metaphorically (Craig, 2002, p. 35). However, Craig argues that more is needed in order to explain why the concept of knowledge should relate only to humans, other than a mere description of how we use our language. What explains why we are willing to attribute knowledge to people, and not to inanimate objects, especially given that the former can be regarded as sources of information at times too?

Craig argues that there is a point to this. His explanation is that humans can be actively helpful when they inform others. And thus – like Moran in his account of the special nature of a speech act of telling – Craig focuses on the voluntary, intentional aspects of humans' ability to communicate so as to inform others (Moran, 2006, esp. pp. 299-300). Here he considers that what makes informants stand out from other sources of information are two things:

> (i) (T)he convenience of an informant as opposed to a 'mere' source of information, and (ii) the special psychology of team-work in a community, something which is involved in the use of informants but not in the use of sources of information. (Craig, 1999, p. 36)

Craig elaborates upon both points in turn:

> The point about convenience is that anyone who understands both the question (a test which the inquirer can hardly fail) and the language he speaks or gestures he uses can get information from an informant, whereas getting it from a 'source of information' will often call for various degrees of more specialised knowledge, perhaps for a special ability to evaluate evidence. (Craig, 1990, p. 36)

And with regard to the second point:

> What I have in mind is the special flavour of situations in which human beings treat each other as subjects with a common purpose, rather than as objects from which services, in this case true belief, can be extracted. (…) an informant is a co-operating member of our species. That means that he can often empathise with the inquirer, and react not just to the question but to the presumed purpose of asking it, so giving the inquirer useful information that he didn't know he had need of (…) Mere sources of information, on the other hand, though they may often be extremely useful, are never actively helpful. (Craig, 1990, p. 36)

Thus, according to Craig, the main thing that distinguishes an informant from another source of information is convenience, since no background knowledge is required if one wants an answer from someone who has the relevant expertise. This is coupled, on the other hand, to the thought that humans, who are approached as cooperative members of the species, have an asset that sets them apart from inanimate objects. They can be actively helpful in finding the answers that one needs, because of they are able to empathize with the inquirer. As such, they may well have a further insight into what knowledge the inquirer may additionally need, in addition to what he had asked.

The idea of human sources of information, i.e. "informants", being central to Craig's account of knowledge, then, is because only human sources can communicate and provide the information to the (sometimes veiled or buried) questions which others want answered. What explains and justifies our focus on human informants is that, when we seek to explicate the concept of knowledge, we find that human communication plays a role that is of special significance.

I do not think that JTB-theorists can find fault with the assumption on Craig's part that knowledge is rightly to be attributed only of humans as they make this same assumption. However, what is interesting is that Craig's idea that we turn to other people to learn truths about the environment, and that informants should be central to understanding our concept of knowledge, would appear to vindicate assurantist theories of testimony over evidentialist ones.[2]

However, this impression is not wholly right, for two reasons. First, within the debate between assurantists and evidentialists, testimony remains a source of "second-hand" knowledge. But, as I will argue, only if testimony is a generative source, i.e. a source of first-hand knowledge, can we understand the way in which testimonial exchanges are formative of knowledge, and as much is required if and when we interpret Craig's genealogy as one that explains why the social standards for our (social) concept of knowledge are so high.

Secondly, Craig speaks not of an "informant" but of "good informants". Here "good" is an idealized notion – as also noted by Miranda Fricker, who points out the good informant is "a near idealized testifier" (Fricker, 2012, p. 251). In the stage of constructing

---

2 I will discuss this debate in more detail in Chapter IV.

the concept, both in the subjective stage and the stage of objectivization described by Craig, the central question is what conditions a good informant should meet. In answering that question, he insists that we have some rather high demands: a good informant is whatever an informant needs to be a reliable source, and yet also accessible and identifiable (Craig, 1990, Ch. 10).

As such, the genealogical account actually remains compatible with the two rival philosophical positions concerning knowledge transmission through testimony: evidentialism and assurantism, which I will discuss in Chapter IV. For it may be that anyone who is a good informant needs to satisfy all of the criteria an evidentialist would set for being a reliable, trustworthy source. This prima facie compatibility is a good thing: it shows that testimony requires a better elucidation than the current mainstream rival conceptualizations of it provide, while the understanding that I shall put forward is one that is supported by the social nature of knowledge and the genealogical argument that acknowledges this point. Even so, by making the notion of an informant so central in the genealogy, testimony is central to Craig's elucidation of the concept of knowledge, and this is a typical move for assurantists, who argue that testimony and knowledge are closely wedded concepts. On the genealogical account, and indeed on the social account of knowledge it serves to bolster, the social practice of testimony is basic to our understanding of the concept of knowledge and to any actual cases of knowledge (Kusch, 2002, p. 71; Welbourne, 1986, p. 81).

In addition, for my thesis, Craig's emphasis on the active helpfulness of an informant is of interest too. In so far as we may wonder, later on, whether people should take an active role in sharing knowledge, it is relevant that we have a special capacity, of proactively acting as an informant, that may shape others' expectations that – given a spirit of cooperation – we should do so.

5. Criticism of Craig's Genealogical Method

Elisabeth Fricker has argued that Craig's genealogical account lacks coherence. She believes that it could not be possible that the concept of knowledge came about in the way that Craig suggests it did (E. Fricker, 2015, p. 46). Moreover, she faults Craig's theory because she believes that mere possibility does not suffice for giving a convincing theory of knowledge. Her requirement is even more demanding: it ought not to be simply possible, but historically plausible, that the concept have evolved this way (E. Fricker, 2015, p. 47, pp. 52–53). Ideally, of course, the most plausible historical account would be the historically true account. Her argument for claiming this is that if the genealogical story is not (at least) historically plausible, then the story does not contribute to the truth/elucidation of the concept. In such a case, a merely analytical elucidation of the concept and its function would suffice.

Before arguing that there is no historical plausibility to the story, Fricker argues that the analytical elucidation of the concept that she attributes to Craig is not convincing. She argues that, for the analytical elucidation to work, it needs to be the case that knowers are (putative) good informants in the typical cases that Craig identifies (E. Fricker, 2015, p. 53). But she finds this implausible (E. Fricker, 2015, p. 54). On the one hand, some people have knowledge, but cannot communicate this. On the other hand, people who have knowledge need not (be willing to) ever communicate their knowledge. She cites Kelp's example of a priest who may never recount what he has heard in confessionals (Kelp, 2011). She further adds the case of babies and animals cited by Jennifer Lackey, and, in addition, intelligent life-forms unrecognizable to us, who do not have a language. Conversely, she believes that one can be an informant without having knowledge. Here, assuming the JTB conception of knowledge, she mentions Lackey's case of the creationist teacher who informs students of evolutionary theory but without believing it, and who therefore, on the JTB conception at least, does not know. She also invokes Lackey's example of a group that may serve as a reliable source for information, without perhaps being adequately ascribed knowledge of what they promote (E. Fricker, 2015, p. 54; Lackey, 2008, p. 58; Lackey, 2014, p. 263).

Note that a presupposition of every one of these counter-examples is that they are exactly that – the kind of counter-example that would falsify any analysis in terms of necessary and sufficient conditions. However, I have already explained Craig's rejection of this model. I have also noted his aim of capturing the prototypical, generic, form of a concept such that its instances can have "outliers": unusual, a-typical cases of precisely the kind captured by Fricker's lists of silent priests and skeptical teachers of evolutionary theory.

Fricker's view is that the notion of communication is separate from the concept of knowledge and cannot be used to individuate that concept, or even to merely elucidate the nature of the concept. She concedes that knowledge and communication are linked in an important way: knowledge is a norm for assertion. But even as Fricker holds that knowledge is crucial for communication, she resists the idea that the concept of knowledge – or rather, of being a knower – would entail that of being a (good) informant. She envisages the connection between knowledge and communication as the reverse of the one Craig supposes. Rather than communication being central to knowledge, whereby exchanges of information are assumed to be more primitive than knowledge, knowledge is central to information exchanges and hence more primitive than communication – hence the title of her paper "Know first, tell later" (E. Fricker, 2015, p. 59).

So, for her, knowledge has a role in enabling communication of the assertoric type, but this role does not yet capture the concept of knowledge. Fricker would have to agree that knowledge is in part individuated by its role in enabling assertion. However, consistently with this, she can maintain that it is not the only or most important role of the concept of knowledge, so that more needs to be said about the concept if we are to understand its nature.

In her own explanation, Fricker emphasizes two main roles of the concept of knowledge

that she thinks are more fundamental than the role of knowledge in assertion. First, it features as a norm in practical reasoning: *"If and only if one knows that P, may one with epistemic propriety deploy P as a premiss in practical reasoning that may lead to action"* (E. Fricker, 2015, pp. 79–80).

Second, the concept of knowledge is used for those beliefs which are stable, and which can be stored to rely upon later: *"If and only if one has evidence sufficient to know P may one without further enquiry accept P into one's store of settled facts, without further research"* (E. Fricker, 2015, p. 80). So knowledge will yield an end to research, providing the "settled facts" which can be stored and propositions that can serve as premises because the possessor is sure of them. As a consequence, the concept of knowledge signals that someone who is in possession of knowledge is epistemically positioned to act as an informant, and specifically to use the special speech act of telling others that P. So, the main two roles of knowledge on her account are that of signaling rationality and epistemic security on the part of the knower himself or herself, whereby a third role which follows from this is that knowledge also signals this to others.

In countering Craig's account, she appears to wish to undermine the competitive advantage his account has over mere functional accounts, namely that it can claim to pinpoint the most important function of the concept by way of pointing to the fact that there is a plausible need which gave rise to the concept. This is an advantage not only because it singles out and explains why some function is pivotal to understanding the nature of the concept, but also because the genealogical approach allows that the concept has come to deviate somewhat from its original need, such that exceptions can be allowed and accounted for (e.g. exceptions in terms of some people not being willing or able to inform others, but nonetheless still being regarded as knowers).

Fricker argues that Craig's State of Nature story could not possibly have happened, because it is incoherent to suppose that people can inform one another without having the concept of knowledge (E. Fricker, 2015, p. 59). She claims that both to be able to inform another person and to be able to interpret someone's assertion as intended to convey information, one needs to possess the concept of knowledge. The reason is that only knowledge, on the side of the speaker, warrants presenting information to another as the truth as to whether or not p, and the audience can only appreciate an assertion as such, if and to the extent that they possess this concept.

A direct counter to this is given by Michael Williams (2015). He clarifies that even without language we can intimate what it is we need, so we don't need the concept of knowledge to start off with, as Fricker supposes. Here he points to the work of Robert Brandom, whose Sellersian view is that concepts are "a linguistic affair". That would appear to commit him to the same view as that of Fricker that language comes before conceptual grasp. Yet, as Williams points out, Brandom explains that epistemic deference is something which we practice prior to having the epistemic language to do so, while "epistemic vocabulary is a kind of secondary vocabulary" (M. Williams, 2015, p. 169).

This idea is supported by the point that we can defer without the use of the word "knowledge", which is the case in one of Craig's examples, as Williams points out – namely the example where I ask Fred, because he is up in the tree and can see where the lion is (M. Williams, 2015, p. 169). For what this shows is that we could use a word which is not yet had in that situation; even while, strictly speaking, this is a situation where we are in direct contact with our informant, and hence have no need for the term of "knowledge" as Craig has explicated it, for that term comes to be required when we need to find out whom to ask. In these kinds of situations, it could well be that there are attempts to express what we need. That is, there can be, as Williams clarifies borrowing an idea from Michael Oakeshott "intimations" of the more sophisticated way of talking: that allow us to say what we could otherwise only show (M. Williams, 2015, p. 170).

On the basis of this reflection, Williams comes to the deeper insight, worth noting, about why we have the word "knowledge". He notes that it facilitates the identification and evaluation of sources – which is precisely the point on which Craig insisted – but adds to this – as is much in line with the topic and line of argument of this thesis – that we need the *word*, over and above the concept, from the perspective of the usefulness of having it as an evaluative term available for the explicit purpose of spreading knowledge (claims):

> The question now arises: what is the point of having the word "know," if the concept of knowledge was always already implicitly present? The answer must be that explicit epistemic vocabulary makes the practice of information sharing more flexible and powerful. Here is an important way in which this happens: with explicit epistemic vocabulary in hand, we can share information not just with informants but about them; and we can investigate their trustworthiness in a more self-conscious way. (Williams, 2015, p. 170)

An indirect counter to the Fricker's objection is that it relies on the claim that knowledge is the norm of assertion – while the speech act of telling is a subclass of that of assertion, and involves a speaker giving assurance to her audience (E. Fricker, 2015, p. 62).

The second view is related to the first, if and as far as it is assumed that an act of telling requires – according to social, moral or epistemic norms that sustain assertoric practice – that one is well enough positioned epistemically to give the assurance in the first place. My view as I will present it in the next chapter runs counter to this. I hold that testimony can be and often is generative of knowledge, and so I think that we often make statements that only become knowledge at a later stage. So, we do not need to speak from knowledge, in order for our utterances to come to qualify as such. This position is also defended by Martin Kusch (Kusch, 2002, pp. 72–73).

Fricker's "historical critique" can be answered. One way to do so, which she has not discussed in her paper, is to hold that the norms about assertion develop only alongside or after the concept of knowledge has arisen in the way described by Craig. That is, Fricker claims that there can only be cases of informing others if the norms of

assertion have been instituted. While Fricker is clear that she does not think there can be communication without such norms, I believe that she is assuming a certain kind of communication – namely a practice of telling which is governed by a mutuality of normative expectations – which may well come into being only after there is a shared basis of trust in the first place, and a shared understanding of what epistemic cooperation requires, which arises precisely because of a need to cooperate epistemically. If that is the case, then the fundamental assumption of Craig's genealogy is not in question. As a genealogy need not have historical plausibility, but only explanatory power to elucidate the concept in a way that we find instructive, the genealogy is jeopardized by Fricker's historical critique, or the point that there need to be normative expectations, as she too has recognized that the latter arise.

Thus, while Fricker appears to treat Craig's genealogy as a failed attempt at a reductive analysis, stating explicitly that his methodological strategy is a reductive one, that interpretation is a mistake (E. Fricker, 2015, p. 47). A genealogy's success does indeed depend on how well the constructed concept matches both the intension and extension of our concept of knowledge, but the constructed concept should not be taken to be the equivalent of our concept of knowledge. It is the convincingness of its conjectured purpose versus the extension that role would yield, for the constructed concept, and the match between the constructed, objectivized concept as explaining the concept of knowledge, that is to do the work – not the reduction of the (non-objectivized) function of the concept.

6. Situating the Objection: The Genealogy versus JTB

Welbourne's "belief-theorists", as well as defenders of conceptual analysis, may hold that Craig's genealogical account is really to be cast as a functional analysis of the concept. One reason why a functional analysis is so tempting, especially for JTB-theorists, is that Craig's supposed individualistic starting point may naturally be expected to yield an individualistic end-result as well. The outcome would then be a description of "a knower" as "a good informant". This result fits perfectly with the individualistic outlook of JTB-theorists. So, I argue that it is not an accident that we cannot reduce the genealogical account of knowledge to a functional analysis.

However, as I see it, this move that assumes an individualist epistemology is compatible with the genealogy is rightly resisted by Craig and by all those who take the social nature of the concept of knowledge seriously, because looking at being a knower in terms of "a good informant" is only one half of the story and neither its end-result nor most illuminating part, which concerns the fact that being an informant is a status that is socially created and ascription of which has social use.

Moreover, while we can understand the concept by reference to considering why people who lacked the concept would have invented it, the way we understand the concept is not typically in terms of its function. This "original use" for the concept is hidden from us when we think of it (Craig, 2007, p. 197). And perhaps this is how it should be. Bolstering his genealogical account of the concept of truthfulness, Williams has argued that if we were to understand truthfulness as a functional norm, we would never have succeeded in living up to its demands. Because "no society can get by with a purely instrumental conception of the values of truth" (Williams, 2002, p. 59).

So, it is part of Williams's story that the values of truth should be treated and should come to be considered as having intrinsic as opposed to a merely instrumental value. Precisely because it appears impossible to explain how something has intrinsic value without giving an account that may be seen as reductive (as it is valued for instrumental reasons; Williams, 2002, p. 90) or non-naturalistic, i.e. metaphysical or Platonic, appealing to higher qualities that a thing has as viewed from or as pertaining in some other realm (Williams, 2002, p. 18), Williams argues that:

> (I)t is in fact a sufficient condition for something (for instance, trustworthiness) to have an intrinsic value that, first, it is necessary (or nearly necessary) for basic human purposes and needs that human beings should treat it as an intrinsic good; and, second, that they can coherently treat it as an intrinsic good. (Williams, 2002, p. 92)

The special feature of a genealogical explanation, as opposed to a merely functional explanation of a concept, then, is that it can "represent" intrinsic value "as arising from more primitive needs and desires" such that "when we reflect on that story, we can find the value intelligible without at the same time losing our hold on it" (Williams, 2002, p. 90). So, we can understand the value then as intrinsic, because we understand both that and why it is intrinsic, and can do so without a need to conceive of the value as functional, which would be to reduce it to an instrumental value, even if it did originate to meet certain desires and needs.

Similarly, one can say that knowledge needs to be an idealized notion, as it contains the aspiration of an ideal, namely the possession of something which is true. The notion is ideal in its nature, because and in so far as we additionally assume human fallibility. The fact of fallibility, as a corollary of an objective notion of truth, makes for an ideal conception of knowledge, since it entails that we can never know that we know, even if it does not speak against our having first-level knowledge (Welbourne, 1986, p. 83).

For first-level knowledge it suffices that what passes as knowledge really is knowledge. Given, however, that there is always a possibility of error, a functional understanding of the concept would yield a conception of knowledge which is unstable. For the idea that the concept of knowledge is really only what we socially regard as our best attempt at truth would undermine the very use of the notion of knowledge, which notion is to apply

only to what really is true. The notion would be undermined because we would either 1) give up on trying to claim that we possess knowledge, if we are skeptically inclined and realize that we cannot make good our claim to truth, or 2) we would be debunking the notion of knowledge as objective and stable, and replace it as a notion of something social and changing, while the former understanding explains how knowledge, as a shared concept, is a cornerstone to our understanding of objectivity.

For us to have the concept of objectivity – for the very idea of there being truths that we can all grasp to make sense – we need the idea of the possibility of transcending the various different individuals' subjective viewpoints. The fact that we seek to learn truths about the world from one another is presupposed by Craig and is what gives content to his (fictional) story about the pull towards the "objectivization" of our concepts. I also noted the importance of a presupposed, shared, objective world to Welbourne's account; it is an assumption embedded deeply in the argument of Williams's *Truth and Truthfulness*, shaping as it does Williams's (Davidsonian) claim that without the concept of truth natural languages would not be learnable (Williams, 2002, p. 5, p. 284 footnote 9; Davidson 1973; see also Stevenson, 1993).

So, I have argued that while Craig's genealogy starts out from the viewpoint of an individual inquirer, that individualistic perspective is transcended when we move to the stage of objectivization. In that stage, we see that the criteria for knowledge are set at a social level. Even so, it does not immediately follow from this that individuals will have to rely on each other's judgement for who qualifies as a reliable source, since the idea of a good informant can be internalized, after the standards have been socially determined. However, I believe that it may be argued that the questions of what qualifies as knowledge and who qualifies as a knower is no longer a matter of individual judgement, but has become a shared responsibility of all who are part of the community (or of its "main" members).

The upshot of Craig's account is that we actually need a community that confers epistemic titles – and not just only a community that sets these criteria. The latter may be a reading of Craig which is compatible with all he says: all Craig explains is that the standards for our concept of knowledge are set high as we want the term to pick out reliable resources and have converged on the conception of what a good informant is. What I want to add in addition – which is something which I think Welbourne adds, and I believe Martin Kusch too shares this view – is that we need the community (or the relevant members of a community) to apply the social criteria if we want there to be actual cases of knowledge as opposed to merely a shared conception, such that the concept comes to have content (Kusch, 2002, pp. 71–72).

For all Craig has said about the concept, it could appear that individuals can handle the criteria for knowledge on their own perfectly adequately, and that no social process is needed for an individual to possess knowledge. Certainly, it may be thought that an individual can entitle herself as a knower and that she can qualify some propositions as

knowledge for himself or for herself. However, what I want to say about such a case is that it is the exception which is made possible by the stable social norms that govern knowledge ascriptions. In fact, as I see it, even such an exceptional case is social to the extent that piggybacks on other, related, actual social judgements being made regarding what passes as knowledge and what not. An individual may "create" knowledge, even if it is not socially validated, but only validated by himself or herself, but only in so far as there is a possible audience with whom it may be shared and if there is no social contestation to it available in the form actual cases of knowledge that have been socially validated and that undermine or contradict it.

Conclusion

As I interpret Craig, the outcome of his genealogical argument is that we should move to understanding the social role of the concept of knowledge. One objection tries to reduce his genealogical account of the concept to a functional analysis, which could be thought to be possible by running it backwards via "reverse engineering", i.e. by ignoring the story and simply taking at face value what the genealogical story yields. Against this tempting move, I argue that it is not an accident that we cannot reduce the genealogy Craig provides to a functional account of the concept nor that JTB-theorists would seek to do so. The reason why such a reduction does not work is that it ignores the second part of the genealogical account, whereby the concept of knowledge is formed. In this stage, which Craig calls the objectivization of the concept, the perspective of the individual which was Craig's starting point is overcome. Instead, the concept is one that comes to embed shared social standards. I have added that it follows from this that the concept's application thereby comes to be a social affair, too. While the concept refers to a socially constructed reality, its instantiations are validated by social judgement too.

Chapter III | Knowledge, Truth, Epistemic Justice

Introduction

The aim of this chapter is to revisit and answer the main difficulty raised by the social nature of knowledge. This difficulty, which we encountered in the last chapter, is that a social conception of knowledge appears to jeopardize the notion of objectivity by relativizing truth to a community's agreement or instrumentalizing truth in so far as it serves as a criterion for knowledge. The reason is that social factors are now seen to have a role in establishing what comes to be validated as the truth. This is problematic to the extent that our conception of the everyday concept of knowledge depends on a robust or "realist" understanding of the notions of truth and objectivity where it is constitutively independent of evidence (Strawson, 1977, p. 20).

In this chapter I explain how Williams's genealogy of truthfulness offers an important way of understanding why we care for a conception of knowledge that does not reduce or relativize truth. In turn, so I argue, this is a double vindication for the genealogical understanding of knowledge offered by Craig: both for its *content*, which shows how knowledge is social and yet does not reductively define the concept, and for the genealogical *method* which makes such an understanding of the concept possible (and which Williams also defended). While the main aim of the chapter will have been served by explaining how the social nature of knowledge is elucidated and accounted for by way of the two genealogies, the chapter also takes an additional step.

That step is that I argue that our social cooperation does not require only ethical behavior but *also* the right political structures. This is to move beyond the notion of truthfulness, in so far as that notion is intended to capture only ethical requirements and focuses only on individual agents and their conduct even if it is emphasized that their ethical behavior is required because of the social conditions of life, and particularly the fact of social (inter)dependence. So, I argue, additionally, that the more central and fundamental notion to understanding the possibility of knowledge is that of epistemic justice, which has both an ethical and political dimension (Fricker, 2007, pp. 1–2, p. 8).

My conclusion then is that if knowledge is social, then in order to understand how knowledge is possible, we should turn to the ethical and political aspects of knowledge generating practices – hence of, as Fricker puts it, our "epistemic practices", i.e. those "human practices through which knowledge is gained or again lost" (Fricker, 2007, p.

vii).[1] In saying this, the value of Williams's work is not at all diminished: in fact, his insights concerning the importance of truthfulness for knowledge and concerning the relation between truthfulness and trust are used better to understand the role and nature of testimony and why there is a need for epistemic justice.

Even as I will explain the nature of epistemic justice more fully later, I argue here already that epistemic justice is the proper alignment of truthfulness and trust, since the need for such a balance is implicit in Williams's account. Williams pointed out that while truthfulness is required for knowledge, trusting relations are required for sustaining truthfulness as the shared ethos between those who cooperate to sustain knowledge. Trust in this context is secured when the expectation that the trust one bestows in others is rightful: when it is a reasonable conjecture that such trust will not be betrayed. As Williams made clear, when contributing to the knowledge had by all, each person expends time and effort and thus makes an investment, which is based on a mutual expectation that all will contribute in a like manner. That is, there is a shared ethos, and there are both sanctions and/or useful remedies that ensure all will comply.

What I contribute as a friendly amendment in my rendering of Williams's genealogy is that the investment made by all is not simply that of being *truthful*, but that of being *epistemically just*, while the shared ethos is not one of *truthfulness* but one of *epistemic justice*, as the latter is the broader term, and is both ethical and political in nature – or so I will argue.

2. Ethics as a Requirement for Knowledge: Williams's Genealogy

Craig's genealogy, discussed in the previous chapter, started from the stipulation that epistemic cooperation is a basic component of his genealogical account. In any human society, epistemic cooperation will have developed in the light of the need to enhance chances of survival. Bernard Williams developed this point and then looked at which strategic motivations might come into play when people cooperate epistemically.[2] Williams was particularly concerned with the explanation of the "value of truth" (in

---

1 Examples by Fricker are "conveying knowledge to others by telling them, and making sense of our social experiences" (Fricker, 2007, p. 1). Note that the term "practice" is defined by MacIntyre as:

   "any coherent and complex form of socially established cooperative human activity through which goods internal to that form of activity are realized in the course of trying to achieve those standards of excellence which are appropriate to, and partly definitive of, that form of activity, with the result that human powers to achieve excellence, and human conceptions of the ends and goods involved, are systematically extended." (MacIntye, 1981, p. 175).

2 That Williams's genealogy of truth and truthfulness is a corollary to Craig's genealogy of knowledge is reflected in the likeness of what Williams describes as the starting point of his genealogy of truthfulness to the account given by Craig on how the concept of knowledge arises, namely by a need for epistemic cooperation in the form of sharing – or as Williams's terms it "pooling" everyday truths (Williams, 2002, p. 43), and from Williams's positive appraisal of Craig's genealogy (Williams, 2002, p. 21, p. 32, p. 284, footnote 2).

the sense that he explains, namely "as a shorthand for the value of various states and activities associated with the truth"; Williams, 2002, p. 6).

If people do not uphold this value, the project – increasing the stock of knowledge held by the group by working together – will fail (Williams, 2002, p. 88). This is not only because the joint endeavor will need to yield beliefs that are likely true – after all, it could be that the majority of the beliefs *are* true. The project would also fail because "a necessary condition of co-operative activity is trust, where this involves the willingness of one party to rely on another to act in certain ways" (Williams, 2002, p. 88). In the case of epistemic cooperation, this comes to a willingness to trust that others act truthfully when it comes to speech, and to their trustworthiness as a source in the sense that when inquiring, they must be motivated by the truth, resisting fancy and wishful thinking (Williams, 2002, p. 127).

But, according to Williams, people may well have reasons to free ride when they can expect to benefit from others' contributions and are not themselves forced to expend effort and contribute to the shared pooling of information (Williams, 2002, p. 58). Indeed, it may be advantageous for a member of the group to mislead their peers and tell them an untruth. Williams's genealogical account of truthfulness, then, attempts to explain why, and under what conditions, epistemic cooperation is possible: "What would be needed for it to work?" (Williams, 2002, p. 59). The intrinsic value of truthfulness is thus explained by its being needed in order to make possible epistemic cooperation. It is in this sense that Williams's genealogy addresses itself to a "transcendental" question: a "how possible" question. How is the knowledge system, construed as a mutual assurance game in Amartya Sen's sense, so much as possible? (Sen, 1967; Thomas, n.d.).

In brief, Williams's explanation for the moral status and intrinsic value of truthfulness is that it is necessary for the existence and sustaining of epistemic cooperation. To make epistemic cooperation possible, Sincerity and Accuracy – the two virtues of truth, each a part of our conception of truthfulness (Williams, 2002, p. 94) – have come to be regarded as moral virtues grounded in corresponding ethical dispositions.[3] (Williams capitalizes these words as they are what he calls "technical terms" in the context of his overall argument). Williams analyzes the problem of epistemic cooperation as essentially a social coordination problem – which arises because of people's self-interested actions – to then solve it by appealing to shared moral values (Williams, 2002, p. 95).

This solution is not too surprising. It has been argued that problems of social coordination such as the prisoner's dilemma – whereby individuals would gain more by working together but would then also incur a greater risk, and where these individuals for that reason better opt for the choice that is in their self-interest even if it yields the sub-optimal result – can only be solved if the parties all share the same, strong ethical values and not only uphold these values themselves "no matter what", i.e. non-

---

3 See Williams, 2002 p. 111, for instance where Williams explicitly refers to "the ethical disposition of Sincerity".

instrumentally, but also have the certainty that the others equally uphold these values. However, the problem of social coordination differs in key respects from the prisoner's dilemma given that there are two equilibrium points in an epistemic cooperation game that differ not simply in the levels of payoff, but also in the extent to which parties are prepared to trust each other. Prisoners dilemmas as a class differ from mutual assurance games, and Williams assumes the latter applies to epistemic cooperation (Sen, 1967; Thomas, n.d.; Williams, 2002, p. 89 and footnote 6).

Williams's genealogy thereby provides an explanation of the possibility of epistemic cooperation. This explanation is his account of how truthfulness could have developed to attain a moral status, whereby the epistemic virtues of Sincerity and Accuracy have come to be regarded as valuable "in themselves". The moral status and moral enforceability of these epistemic virtues can plausibly explain why human interactions are normally truthful: there are social repercussions if one is found to breach the norm. Indeed, it seems to be true that truth and reliability are valued morally: not telling a lie is a deeply ingrained moral maxim. The social sanctioning involves one's status. Here the social statuses of honor and shame come in so Williams points out (Williams, 2002, p. 120). That is, according to Williams, it is through the motivations of honor and shame that truthfulness is upheld in society. These external sanctions are needed, one may think, before the ethos of truthfulness is internalized by the participants in the practice, and afterwards, too, for those who failed to internalize them (Williams, 2002, p. 44). Indeed, Williams notes that the motivations of honor and shame remain crucially important to date, to ensure truthfulness in our practices (Williams, 2002, p. 121).

Williams's more basic point, however, is that truth and trust are connected. Those who are truthful – that is sincere in their speech and/or serious about finding the truth – are regarded as trustworthy or reliable. Williams even points out that "truthfulness" and "trust" are linguistically related (Williams, 2002, pp. 93–94). Such a social status, of being trustworthy, plays a major role in ensuring that people can trust each other. Moreover, the fact that the values are publicly upheld will ensure that the expectations at both ends of a transaction are clear. Likewise, it ensures that any defaulting by a speaker S on meeting the expectations of trust by her audience A, such as would be involved in deception or unreliability, will have repercussions for S's social status as a moral and epistemic agent. That is, for everyone personal reputation and social status are at stake, and this is known to be so by everyone. This helps to secure that amongst all contributing parties trusting relationships can be sustained, while these relations are such that they harbor and require truthfulness. Hence, such that epistemic cooperation is possible. Thus, Williams provided an explanation for our shared attitude of valuing truth and thereby provided a reason to assume the general trustworthiness of others. Truthfulness serves to explain, via its connection to trust, how we could come to have any sustained form of epistemic cooperation.

Yet, Williams himself warned against the piece of "bad advice" that "if we know

absolutely nothing about someone except that we share a language with him, then we should start by trusting him" (Williams, 2002, p. 111). That is, we cannot deduce anything much from the general norms that govern assertion, even if we know that in general people will assert that p if and only if p is true, given the internal connection between belief and assertion, on the one hand, and the internal connection between belief and truth on the other. Individuals can still lie, because strategic considerations continue to play a role. That is, a breach of a moral norm is possible. Instead, says Williams, we had better ascertain – or at least have reason to believe – that the "normal circumstances of trust" do indeed obtain in this case (Williams, 2002, p. 88).

As we will see, this may be one reason why Paul Faulkner regards Williams as falling in the camp of "evidentialists" about testimony. That attribution is also made by Richard Moran (2005, p. 339), but for a different reason. Moran claims that Williams maintains that sincerity amounts to a speaker giving the hearer access to what they believe when they assert something; i.e. that there is no gap between what the speaker presents as true and what that speaker herself believes to be true, too (Moran, 2005, p. 327, p. 332). This is not an adequate account of sincerity, on Moran's view, as it does not give a role to the speaker's intentions. Williams's view of sincerity is not helpful for understanding testimony, he thinks, as a speech act of telling someone something is intentional and testimony is typically distinctively linguistic: it differs from any other form of giving evidence (Moran, 2005, 2006). So Moran argues that we need to see sincerity as about the intention of the speaker rather than as about access to her beliefs and/or congruence of what the speaker says with the content of her own beliefs.[4]

But as I argue, Williams's strong emphasis on the connections between truthfulness and trust speak against the attribution of evidentialism. Rather, Williams appears to have pointed to the fact that – in so far as they push us in different directions – we need to try and ensure that truth and trust are aligned. As I have emphasized, Williams's account of truthfulness is premised on the idea that epistemic cooperation is characteristic of what testimony is about. Then, truthfulness is required (and likewise trusting someone as a source) in so far as the parties are both trusting and cooperative. That is, when there is an internalized ethos on both sides that secure the circumstances of trust and hence

---

4 A difference between Moran and Williams is that the former focuses on the speech act of telling, while the latter is concerned, more generally, with sincerity in speech, whereby assertions are a broader category (Moran, 2006, 2005). Moran notes that sincerity is only relevant, according to Williams, when it comes to those assertions that seek to inform (Moran, 2005). I do not think there is an actual difference between Williams and Moran when it concerns sincerity in the case of tellings. As we have seen, for Williams, sincerity is a form of truthfulness and essentially comes to trustworthiness. This is a matter of one's intentions, indeed, in a case of presenting something as a truth, as in the act of telling. I think that my account of testimony as a generative source can explain why Williams focused on assertions: if Williams sought to explain the possibility of the social generation of knowledge, that explains why he thought the notion of sincerity pertained to assertions as well as other speech acts such as tellings. It also explains why Williams did not think speaking from knowledge was required for truthful speech. Linguistic contributions that are sincere help to generate knowledge, and hence knowledge need not exist prior to sincere speech acts. A case in point, which Williams refers to is that assertion can serve as a way for a subject to stabilise his or her mind and come to acquire self-knowledge and that 'It is the presence and need of others that help us construct even our factual beliefs.' (Williams, 2002, p. 194).

of cooperation. I will argue that this is not so much a matter of reciprocity, as Williams himself noted, but of epistemic justice. What epistemic justice involves, more precisely is something I come to in chapters VI–VIII.

For Williams, there is no simple rule that can provide an easy answer in specific cases (Williams, 2002, pp. 120–121). Reasonable trust is a matter of the relationship between the speaker and hearer, but not only of that. It is a matter of the governing norms and the context that relates the two parties (Williams, 2002). At any rate, according to Williams, the only way to act truthfully and trust responsibly is to act from and recognize in others the dispositions that belong to the virtues of Sincerity and Accuracy. Williams goes a long way to spelling out these dispositions, at least, in their schematic form. For now, it will have to suffice that on his view, Sincerity comes to not deceiving others in speech; he describes Sincerity as "trustworthiness in speech". Accuracy on the other hand, is related to the process by which one comes to have a belief in a way that secures its likelihood of truth. Accuracy is best seen as the virtue of resisting the temptation to believe that p is true because one would wish that to be so (Williams, 2002, p. 127).

3. Assumptions of Williams's Genealogy and its Plausibility

Alan Thomas has pointed out that it is *assumed* – without argument – in Williams's discussion that the epistemic cooperation in question is to be characterized as a mutual assurance game (see Williams, 2002, p. 89, p. 290 footnote 6; Thomas, n.d.). In turn, that means that the product it yields is something which individuals cannot achieve through simple trades but have to achieve as members of a collective (as Thomas puts it, achieving a "steep good" of knowledge through the "economies of scale" of mutual cooperation). This assumption can be inferred from the fact that only in the latter case does a problem of free riding arise; a problem that Williams invokes the virtues of truthfulness precisely in order to solve (Williams, 2002, p. 89; Thomas, n.d.). I believe Williams's view was that knowledge constitutively depends on the social practice in which it is embedded, and I also believe that if Williams did not make this explicit or provide enough backing for this claim, then Thomas's complementary argument does so (Thomas, n.d.).[5]

I will also provide my own argument in the next chapter on testimony, vindicating Williams's idea about the relation between truthfulness and trustworthiness. There I argue that Williams's ideas concerning assertion and sincerity make most sense within a

_____

5 One troubling point, however, is that Richard Moran has interpreted Williams's stance on testimony such that Williams defends a reductive position, which Moran calls "the Evidentialist View" (Moran, 2006, p. 277; Moran, 2005, p. 332) This would seem to entail a form of individualism. I will argue later that I view Williams's genealogy as aimed at explaining the conditions needed for knowledge being socially generated and sustained through testimony. This does not alleviate an individual from taking epistemic responsibility in giving or accepting testimony, but such responsibility is shared.

knowledge-generative account of testimony, whereby testimony is a fundamental source of knowledge, and indeed, as in Fricker's term, "a social epistemic practice": it is the pivotal form of epistemic cooperation. This also explains why Sincerity in speech is more central to Williams's focus than Accuracy (telling the true from the false). Of course, his account of Accuracy remains, so it appears, individualistic in terms of who qualifies as the one responsible for knowledge, but this is something we can see as true in so far as we are concerned with an individual's contribution to the whole as a member of the epistemic community. Indeed, it is that for which the virtue of Accuracy was looking to provide an ethical standard.

4. The Politics of Knowledge: Fricker's 'Epistemic Injustice'

Miranda Fricker explicitly conceives of testimony as a "social epistemic practice", whereby an epistemic practice is "a human practice through which knowledge is gained, or indeed lost" (Fricker, 2007, p. vi). While she does not argue for this kind of conception of testimony as a "social epistemic practice" it is a natural starting point given that her aim is showing that there are moral-epistemic wrongs and injustices. I believe that once it is shown that testimony is a social practice, and central to both the *generation* and preservation of knowledge, then we can start to substantiate the idea that there are moral-epistemic norms that pertain to testimony and that derive from the social nature of knowledge.[6] This is the task I set myself over the next two chapters.

One way of substantiating the substantive norms yielded by the social generation of knowledge through testimony is from the "bottom up". Miranda Fricker takes such an approach, beginning from actual practice and identifying certain ways in which we may epistemically wrong another person (Fricker, 2007, p. vii, p. 4).

Another way of doing so is by arguing that, in generating knowledge through epistemic cooperation, testimony meets the conditions of a mutual assurance game, and that therefore an internalized ethos is required for its success and sustainability. The conditions for a mutual assurance game are that individuals share a need for achieving a common purpose, which they cannot achieve on their own, but which they can achieve together (Williams, 2002, p. 89; Thomas n.d.). By taking knowledge as a product of a necessarily joint endeavor, and viewing testimony as part of that endeavor, Williams has assumed that testimony meets that condition (Williams, 2002, p. 89). However, as Thomas noted it is a position he has not argued for (Thomas, n.d.).

My own approach, in the next chapter, is to vindicate and substantiate the view that

---

6 Martin Kusch also holds the view that a social conception of testimony yields links between normative epistemology and ethics (Kusch, 2002, p. 52).

knowledge is social by arguing that testimony is best understood as an epistemic practice, and moreover, as the primary epistemic practice. That argument provides room for ethical and political norms related to knowledge sharing, that cannot be accommodated on the individualistic conception. So, I will argue that testimony has wrongly been conceived as only – at best – a vehicle for second-hand knowledge. In contrast, I argue that in generating knowledge through epistemic cooperation, testimony is rightly deemed an "epistemic practice", and that therefore, as Williams argued, an internalized ethos is required for its success and sustainability.

In the next chapter I will argue that the current conceptions of testimony and of knowledge are not at all congenial for trying to show that there is an ethics of knowledge sharing because of their individualistic bias. The predominant individualistic conception of knowledge and the static conception of testimony that derives from it only accommodate individual responsibilities on the part of those who wish to acquire knowledge on the basis of another person's testimony. Typically, these responsibilities are classed as epistemic ones. Even if such responsibilities are normative, they may often be rationalized as prudential rather than ethical. While being epistemically responsible may be laudable, it is currently viewed as ultimately based on rational self-interest and not on other-regarding reasons. I believe, however, that there is a third way – other than assuming that testimony is a knowledge generative practice or assuming that it is a mutual assurance game – to create room for an ethics of knowledge sharing. This is to engage with the current theories on testimony and challenge these from a meta-philosophical perspective.

Already in her earlier 1998 paper – nine years before her book– Miranda Fricker provided what I consider to be a top-down argument for epistemic injustice – whereby she started from the assumption that our epistemic practices are social (Fricker, 1998/2011). While I will discuss and extensively use Fricker's later work in Chapter VI, as I develop a positive theory of epistemic justice there, Fricker's earlier paper is of special interest in the context of this chapter. There, her argument is "top down", as opposed to "bottom up" (leaving aside, for the moment, that in the book she also provides a genealogy for the "corrective virtue" of epistemic justice, whereby she discusses Craig's genealogy).

By "bottom up" I mean that she provides examples drawn from fictional literature, which are used to yield the intuition that some wrongs and injustices are specifically epistemic and non-coincidentally so. This argument is bottom up in the sense that it prompts intuitions that support, directly, that some things are part of our lived experience. It is bottom up because we look at examples that could occur in actual practices, and take it from there. The top down approach, instead, would be to argue for the existence of a phenomenon such as epistemic wrongs and injustices by deriving or explaining their existence a as function of something that would give rise to it, i.e. the social epistemic practice.

A top–down genealogical argument is what has been provided, in fact, precisely by

Williams's genealogical argument and also received further support from Thomas's argument from an analogy. The bottom-up style of argument takes a different starting point, and seeks to derive truths from abstraction of cases, rather than the other way around. The top-down approach is one of deriving abstract truths and then testing their application. It will theorize first and look at the cases after, to see whether there are cases in the area where such cases were predicted, and whether or not the cases in reality affirm the theory.

There are two reasons for discussing Fricker's earlier paper here. First, the argument she provides for the existence of epistemic injustices, in the 1998 paper, is that it starts from the assumption that knowledge is socially produced. This assumption is less present in her later work, as she there makes the case for the existence of epistemic injustices bottom-up.

Second, the 1998 paper more readily suggests that credibility is seen as a good that may be distributed. This is also evidenced in the 2007 book, for the title of the second chapter is "Prejudice in the Credibility Economy". The image of an economy suggests that something is distributed, and that the item to be distributed is "credibility". In the 1998 paper, she also suggests that there is a special principle for doing so, namely the norm of credibility (Fricker, 1998/2011, p. 60). Her speculation that in our society the norm of credibility may get "rigged" by the socially powerful and/or shaped by the structures that replicate themselves to keep elites in power (Fricker, 1998/2011, p. 62), which explains how epistemic injustice can arise in a top-down fashion, suggests that credibility is something which may be distributed *unfairly* as well as its being assigned *improperly* (as in the case of a testimonial injustice in her 2007 book).

5. Epistemic Justice as the Alignment of Truth and Trust

At the same time as Bernard Williams's writing of *Truth and Truthfulness*, Alasdair MacIntyre discussed in his Tanner Lectures moral attitudes concerning lying (MacIntyre, 1995). While MacIntyre and Williams share the basic intuition – contra Kant – that lying is not necessarily worse than other forms of deceit, and even that lying is not always wrong, they differ in their approach of explaining why others have thought differently on this score. Williams believed, quite radically, that anyone who thinks truthfulness always requires honesty has his or her intuitions about truthfulness wrong and their disposition of Sincerity is "out of shape" (Williams, 2002, p. 115). MacIntyre, in contrast, thought we should see the moral requirement of whether or not lying is allowed – or indeed even called for – as one that pertains to one's outlook on what value is deemed to be at stake, truth or trust (MacIntyre, 1995, p. 338).

That is, McIntyre provides a meta-analysis in that he claims that the moral

philosophers who view lying as primarily an offense to trust allow for lying, while those who view it as primarily an offense to truth are opposed to lying in all circumstances (MacIntyre, 1995, p. 335). The first view, he contends, is the position on lying taken by utilitarians, who tend to value social trusting relations as a precondition of truthfulness. On their view, lying is a breach of trust that is on a par with other forms of deceit. The second position is that of deontologists, who tend to prohibit lying in all cases, though they can allow other forms of deceit that do not directly breach the rule that one must not lie (MacIntyre, 1995, pp. 335–336). What explains the deontological position is that for them any false assertion is viewed as inherently a moral wrong. The violation is one of a rule that embodies what it is to be rational and to treat others as such, as it is a constitutive requirement of the practice of assertion (MacIntyre, 1995, p. 316). By contrast, on the consequentialist view, the immorality of a lie is judged teleologically: it depends on the motivations and likely social consequences of that lie in the particular case. This does not yet mean that lying is easily allowed, but it does mean that there are justifications in exceptional circumstances. These justifications, so MacIntyre clarifies, stem from the underlying social relationships, because social relationships – and more generally "trust" – provide the rationale of the rule that one should not lie and of any justified deviation from it.

Both Williams and MacIntyre may appear to fall into the first camp of the utilitarians on MacIntyre's analysis, as they argue – as we have seen – that lying is morally allowed and that it may even be morally required in certain circumstances. Yet, both resist consequentialism in their own way. Williams resists it, as we have already seen, by arguing for the intrinsic value of truthfulness and by providing a genealogy that backs up this claim. In addition, presumably to show that he is not deviating from what truthfulness prescribes when he says that lying is sometimes allowed or even permitted, he maintains that truthfulness is a broader notion than an avoidance of lying, namely, it is "trustworthiness in speech" (Williams, 2002, p. 97). (He is also, of course, one of the twentieth century's most prominent critics of utilitarianism.) That allows him to argue that the moral distinction that has been made between lying and other forms of deceit, such as misleading another by way of one's use of Gricean conversational implicatures, does not justify the latter as a good way of avoiding the former, because both lying and misleading constitute breaches of trust (Williams, 2002, p. 97). It also allows him to make a case that lying is actually what truthfulness requires in some cases, namely when the social relations are very bad, that is "aggressive, extortionate, structured by threats" (Williams, 2002, p. 114). Here one could think of the murderer at the door who wants to know whether his victim is hiding inside your home. In such a circumstance, Williams contends along with Constant, a person does not have "a right to the truth" given his or her malevolent objective (Williams, 2002, p. 114).

MacIntyre, on the other hand, does not vindicate the utilitarian take on truth telling either. Rather, he seeks to integrate the two moral traditions vis-à-vis truth-telling and

lying, namely Kantian deontology and utilitarianism as defended by Mill, without siding with either one. Explaining moral philosophers' divergent attitudes to lying by taking a meta-perspective, he argues that we are pulled in both directions, because there is a tension (in the case presented) between the values of truth and trust, while we care about both of them. He then defends a non-consequentialist rule, which captures both Kantian elements and the utilitarian spirit that human relationships are at stake:

> It would be misleading to state it as though its form was "Never tell a lie, except when.. ." For this would suggest that we were first formulating a rule and only later, as a second thought, introducing an exception. But this is a mistake. The rule that we need is one designed to protect truthfulness in relationships, and the justified lies told to frustrate aggressors serve one and the same purpose and are justified in one and the same way as that part of the rule that enjoins truthfulness in relationships. (….) The rule is therefore better stated as "Uphold truthfulness in all your actions by being unqualifiedly truthful in all your relationships and by lying to aggressors only in order to protect those truthful relationships against aggressors, and even then only when lying is the least harm that can afford an effective defense against aggression." This rule is one to be followed, whatever the consequences, and it is a rule for all rational persons, as persons in relationships. (MacIntyre, 1995, p. 357)

Hence, MacIntyre defends a view that truthfulness requires one not to lie in most cases but that, because truthfulness is inherently linked to trusting relationships that both require it and that sustain it, there are possible exceptions where lying is allowed or even needed in order to protect trusting social relationships.

So MacIntyre's claim is that in some cases lying is allowed or needed in order to "uphold truthfulness". The act of upholding truthfulness which MacIntyre envisages in these cases is that of protecting trusting relations, which can only be protected by lying. At a first glance, I am sympathetic to this idea, which seems reasonable in the sense that it is akin to a distinction along the lines of what furthers truth more overall, e.g. in the long run as opposed to the short run. Yet, it seems to me dangerous to suppose that truthfulness has different meanings, and that truth-telling is required only in cases where one may reasonably assume that others are trustworthy (too). The better way, as I will argue further below, is to explain the deviation of truth-telling not as embodied by a broadened notion of truthfulness, or as a way of upholding it, but in terms of a principle that strives to achieve the right alignment of truthfulness and trust, i.e. by appeal to epistemic justice. In turn, epistemic justice demands truthfulness in most cases whereby the circumstances of normal trust may be assumed, and a just or rightful response in other cases where trust is lacking.

Thus, both Williams and MacIntyre appear to have argued that whether or not lying is allowed is rightly viewed as a matter of context – and that it is not so much one's ethical theory that decides on this, but respectively, one's understanding of the terms involved or what value should be attached to them when they conflict.

Williams is explicit about what this context is: that the circumstances of trust obtain. That trust would have to be the context fits with MacIntyre's ideas that (1) that lying is

only justified when trusting relationships are threatened, and (2) that moral philosophers of the utilitarian and deontological tradition tend to find the cases of disagreement in the context where the values of trust and truth are pulled apart, such as the case of the murderer at the door. But it also fits with Williams's own insistence that truthfulness is what trustworthiness requires in the epistemic context (Williams, 2002, pp. 93–94), and that in that context – the context of social epistemic practices, i.e. of the cooperation needed for (generating,) sustaining and disseminating knowledge, and pivotally the context of testimony – truthfulness and trust(worthiness) are to stand in a reciprocal relationship (Williams, 2002, p. 111, pp. 118–120).

In what follows, I argue that MacIntyre's meta-analysis can be seen as complementary to the work by Williams, even while it is also in some respects a corrective to it. It is complementary, as it vindicates Williams in his insistence that truthfulness and trust stand in a reciprocal relation in the epistemic context, and that they may be viewed, in most matters, as the two different sides of the same coin: whether to be truthful or whether to be trustworthy comes to the same thing, requires the same action, namely to not lie, in the context of (sustainable) epistemic cooperation. Yet, I think MacIntyre's work also helps us to correct for a misunderstanding that Williams introduced when he speaks of truthfulness in a way that incorporates not only demands of being honest and sincere, but *also* their opposite.

To start with this second point, there is a rather striking passage in *Truth and Truthfulness*, which I alluded to above already, about having one's disposition concerning truthfulness "out of shape" (Williams, 2002, p. 86). From that passage, it would appear that Williams believed it is a matter of truthfulness (or Sincerity) that in a case where it is judged that others not be deserving of truth, we may be allowed or even required to lie and deceive. That is, it appears that Williams wanted to build into the concept of truthfulness an awareness of and an appropriate response to untrustworthy behavior of others. People have to be *deserving* of truths (Williams, 2002, pp. 114–115).

But here I disagree. It is not helpful to see truthfulness, conceived of as "caring for the truth" as requiring that we sometimes lie. That is, truthfulness should not require the opposite response of speaking the truth in some circumstances, which response is deemed morally reprehensible and sanctioned in other circumstances. Rather, what requires us to act differently than truthfulness demands is *not* a matter of truthfulness: it is a matter of epistemic justice. To show why this is so, we can turn to MacIntyre again. For I think that MacIntyre is right in his point that the dictates of truth and trust can come apart, leading us in different directions depending on what we find to be primarily at stake (MacIntyre, 1995, p. 358). (Indeed, precisely this explains the dilemma people face when they need to defy their own nature to speak truthfully, and to assume others' good intentions as a matter of affective trust.)

MacIntyre's meta-analysis then helps us correct Williams's view on the role and demands of truthfulness, and to locate Williams intuitions better. The idea is this: we can

do justice to Williams's intuitions concerning truthfulness, by locating them as intuitions of *epistemic justice*, such that epistemic justice is the *broader* notion than truthfulness. Epistemic justice encompasses the requirement of truthfulness in circumstances of trust, but also the (required) responses to untrustworthiness in cases where an audience is not deserving of truth, because the circumstances of trust do not obtain.

As an example, one can think of how the socially powerful expect the socially powerless to explain how their experience of marginalization feels like "from the inside" – whereby they are expected to speak fearlessly, and open up, rendering themselves vulnerable to inevitable lack of misunderstanding, and – all too often – an unwillingness to see any fault with the situation, and hence to comments that add insult to injury. This is, as underscored by Berenstain, a form of "epistemic exploitation" (Berenstain, 2016). It is an unjust demand for knowledge, which can rightly be resisted, because explaining one's experience is draining and unrewarding and serves only to continue oppression, even if the request may appear a virtuous quest for knowledge (Berenstain, 2016, p. 571).

So, I would argue that truthfulness is to be narrowly conceived as honesty and openness, while the broader notion that Williams makes of it includes ideas of epistemic fairness – fairmindedness and open-mindedness – and is better captured by the idea of epistemic justice. All that we need to see, to understand this point better, is that epistemic justice is best understood as the proper alignment between truthfulness and trust. This is why epistemic justice can be viewed as a measure to define and restore the balance between truthfulness and trust in our social epistemic practices, and why we may need to side sometimes with "the moral" (that is, trusting relations) and sometimes with "the epistemic" (that is, upholding the criteria that are socially set for truth) even as both types of normative considerations, ideally, require us to do the same thing.

I believe that Williams's position is maintained and improved upon by this amendment – and moreover, I believe he would endorse it. First, as said, when reading Williams, and particularly his discussion of how categorical attitudes against lying reveal a misconception of the nature of truthfulness, it is striking that the way he uses "truthfulness" is such that it is natural to substitute "epistemic justice" in its place, e.g.:

> Some relations between people are, on the other hand, notably bad— aggressive, extortionate, structured by threats. In such a case, and that of the murderer at the door is an example, a lie may be a necessary form of defence. Here, it must surely seem to most people now that Constant was right, and that to "have a problem" about lying in such a situation (except perhaps in the sense that you may not be very good at it) is a sign that there is something wrong in your conceptions of what truthfulness requires. (Williams, 2002, p. 114)

So, I argue that the concept of truthfulness should be reserved for honesty (and openness, sincerity), while epistemic justice invites the considerations that allow us to deviate from being truthful (narrowly conceived in the sense of being open and honest or sincere) when others are not deserving of the truth or when social (moral-epistemic) reasons exist

that speak against informing people of some fact at all or simply at some particular occasion. So, on my view, withholding truths can only be explained as a matter of epistemic justice as opposed to being explained as a matter of truthfulness. Williams makes a similar point:

> The reason why the murderer or even the intrusive questioner does not deserve the truth is that he is no longer in a relation to us which is structured by the normal expectations of trustful exchange. Seen in this light, Sincerity will involve in a certain way the sense of justice, and there doubtless are some ideas of justice and injustice involved: it is a very natural thought, for instance, that it may be perfectly *fair* to deceive the importunate questioner, because he has put himself out of line. (Williams, 2002, pp. 119–120)

In the above passage, Williams himself alludes to the notion of justice, and clarifies that for him, this notion features in a conception of Sincerity and hence that *justice is part of his conception of truthfulness* (in speech).

Yet, what I suggest is that it should be the other way around. We have a clear idea of sincerity and of truthfulness, and it appears we should value that clarity and reserve the term for only those contexts where it rightly applies: the circumstance of trust. That Williams is unable to reserve the term for these circumstances is because he rightly sees that truthfulness – in the sense of being honest and open – is not required in some (exceptional) circumstances. However, since truthfulness (again, in the sense of honesty and openness) is not required *as a matter of fairness or of justice*, we should not consider it a matter of truthfulness, but a matter of epistemic justice.

That term, epistemic justice, appears to have been invented precisely to describe this case: to make clear that truthfulness is not always required because truthfulness and trustworthiness are to be reciprocated and hence conjoined: the circumstances of trust have not been met when they come apart. Indeed, MacIntyre too has underscored that once we see how truthfulness is tied to trusting social relationships, we can see how truth and trust are – ideally – values that are aligned, i.e. complementary to one another:

> What reflection upon Mill and Kant has led me toward is a conception of truthfulness as informing and required by rational human relationships, a conception that does seem to go some way toward integrating concerns about truth and concerns about trust. For to understand the rules prescribing unqualified truthfulness as governing relationships, rather than individuals apart from their relationships, is also to understand how the concern for truth and the concern for trust can become complementary. (MacIntyre, 1995, pp. 358–359)

Aside from this, as we have already noted, epistemic justice is conceptually preferable as the broader notion than truthfulness, since it speaks more clearly to a politics of knowing, which is not part of truthfulness because the latter is an ethical notion. In contrast, while epistemic justice may at first appear to be a political notion rather

than an ethical one it can be used as having both an ethical and a political dimension.[7] Williams acknowledges this political dimension, too, when he writes that:

> An individual's commitment to the virtues to truth may stand opposed to a political culture which destroys and pollutes the truth. (Here, as often, intrinsic values turn out to have their uses.) (Williams, 2002, p. 127)

For Williams, then, the fact that there are collaborative enterprises is a reason to insist that the ethics of individuals should be in order and that the virtues should be internalized so as to not give way too easily, being valued for their own sake – and yielding the rightful expectations that others abide by them, too. In contrast, others have instead sought to expand the idea of a virtue from the ethical to the political, by giving it a political use through applying it to institutions, e.g. Anderson and Fricker (Anderson, 2012, p. 170; Fricker, 2013).

6. Ethics versus Metaphysics

Might one be skeptical of the very idea that ethical norms are connected in any way to epistemic success, and this without reducing the notion of truth? Such skepticism has been developed by two epistemologists: Richard Rorty and Alvin Goldman.

In defending his view, Williams argued against Rorty's pragmatist view according to which truth is not the aim of inquiry, while agreement is all that we can strive for (Williams, 2002, p. 59; Rorty, 1998). These views by Rorty, which Williams did not share, were nonetheless relevant for Williams to distinguish his own account from, because Rorty offered an ethical view in response to the question of how truth is to be secured but did so by deeming truth as nothing over and above intersubjective agreement, and by reducing objectivity. Objectivity, on Rorty's view, reduces to the ethical notion of solidarity (Rorty, 1984, 1989).

On the surface, then, there may appear to be a similarity between Williams's approach to truth and Rorty's, in so far as they both are naturalist in outlook, while both emphasize the role of ethics. The difference, however, is that Rorty provided a purely ethical approach, and looks only at human attitudes that have to do with these ethical values. Williams, instead, maintains that the ethical notions are sustained by a notion of truth as intrinsically valuable and irreducible. Moreover, Williams's arguments, as he made clear, are not based purely on attending to our attitudes of truthfulness cast as ethical values, but also on the role they play in "belief, assertion and communication"

---

7 This is not only true when the notion if conceived of as a virtue, as does Fricker, but also when we view it as concerned with both the use and allocation of (epistemic) power. I will defend the latter position, later in this work.

(Williams, 2002, p. 61). This is to insist that there are other than purely ethical grounds that lie at the basis of our ethical dispositions, even if we arrive at our philosophical conclusions "from a naturalized outlook" (Williams, 2002, p. 60), i.e. through attending to our attitudes and practices. The contrast is clear, too, when it is noted why Williams seeks to argue for the intrinsic value of truth.

Williams's main motivation, in providing his genealogy of truth and truthfulness, is to show how these two notions can be "stabilized" (Williams, pp. 2–3). He maintains that all theories that seek to reduce, instrumentalize or eliminate the value of truth are mistaken, while these views are too easily dismissed by those who disagree with them (Williams, 2002, pp. 4–6). Countering the (meta-philosophical) view of pragmatists that truth-seeking or truth-preserving practices and activities are really only valuable to the extent that they are helpful to achieving human goals or useful to avoid danger, Williams argued that "to the extent that we lose a sense of the value of truth, we shall certainly lose something and may as well lose everything" (Williams, 2002, p. 7; see also Williams, 2002, p. 59).

Williams's approach is commendable in that it does not reduce the notions of truth and objectivity in any way. Indeed, he makes a point of not defining truth (Williams, 2002, p. 63). Instead, he provides a genealogical account of truthfulness that serves to illuminate the importance of the notion of truth, understood as independent from the will and from wishes (Williams, 2002, p. 140). The genealogy is to show why truth is valued intrinsically, while it is a notion that sustains many of our epistemic and communicative practices (Williams, 2002, p. 90).

In providing a genealogy, Williams's outlook is that of a naturalist philosopher, who seeks to understand the notion of truth through the importance it has in human life (Williams, 2002, p. 60). Hence he focuses on why we care for truth, and why we should care for it, which deep care is labelled "truthfulness" (Williams, 2002, p. 60). As he explains, on the one hand, postmodernist writing has made many skeptical of there being such a thing as "the truth", viewing it as a politicized notion, but on the other hand, this very skepticism is driven by an even deeper and undeniable sense that truth matters (Williams, 2002, p. 1). This attitude of caring for the truth, untainted by interests and biases, is the attitude of truthfulness, captured by "the unconditional will to truth" in Nietzsche's terminology (Williams, 2002, pp. 14–18; see Nietzsche, 2001, p. 344).

Williams's argument, then, starts from the idea that the method of a genealogy is uniquely able to square naturalism in philosophical understanding of key concepts and a conception of truth that encompasses the notion of objectivity (Williams, 2002, p. 1, p. 7). He maintains that a genealogy can provide an account of truthfulness that illuminates why truth is valued intrinsically, i.e. non-instrumentally, and why it should be so valued (Williams, 2002, p. 35, p. 90). The argument which can be made using a genealogy is that truth being intrinsically valued is necessary to maintain a wide variety of social practices, such as belief, assertion and communication (Williams, 2002, pp. 4–7, and Chapter 4 and

5, especially p. 61, pp. 66–67, pp. 74–75, p. 84, pp. 87–90). One reason why the genealogical method is uniquely suited to make this argument, is that it allows us to resist equating the notion with its function: the notion of truth is not defined. Accordingly, a functional definition of the notion of truth and hence an instrumentalization of truth is resisted.

This also means that the commonsensical conception of the nature of truth(s), in the sense of truth(s) being independent of the human mind and will, is not questioned. As Williams notes, it is not necessary to espouse a Platonic notion of truth in order to secure that conception (Williams, 2002, p. 61). He maintains that his genealogical account of the intrinsic value of truth is compatible with minimalism about the nature of truth (Williams, 2002, pp. 6–7). Minimalism is a metaphysical theory about truth as a disquotational linguistic device. A positive account of the value of truth is compatible with it, because the minimalist does not consider the value of truth – i.e. our caring about truthfulness – for our practices (Williams, 2002, p. 65; see Hornsby, 1997; Horwich, 1990; Quine, 1970/1986; Tarski, 1944; Wright, 1992).

Minimalists have no need to deny the normative claim that truth is intrinsically valuable, as nothing about the normative status of truth, i.e. its role, follows from minimalism (Williams, 2002, pp. 6–7). This may be though to be the case, if minimalists are also deflationary about truth's role in human life, but minimalists need not sign up to that sort of reductivism. Williams's argument is sustained, then, because even while "truth" is correctly viewed as a property of assertion in a strict sense, in a different and normative sense, in which it is often used, too, it is as a guiding ideal for the practices in which it features, such that it makes sense to speak about the "value of truth" and the "virtues of truth" (Williams, 2002, p. 94).

In the context of discussing opposition to his own views, Rorty made clear that objectivity and solidarity are, in his view, a question of "either–or": metaphysicians opt for the former, naturalizing epistemologists for the latter (Rorty, 1984). This latter group, so Rorty claims, tries to reduce objectivity to solidarity, i.e. to ethical attitudes. The former group resists such a move, and discards the need for intersubjectivity. It seems that Rorty's analysis applies to Williams's solution. While being a naturalist, Williams indeed sought to explain or illuminate truth via our attitudes of truthfulness. Moreover, the answer that is proposed to solve an epistemic problem – once that problem is cast as one of social cooperation – is ethical in nature.

Even so, I think this Rortyan critique can be dismissed as resting on a false dichotomy. Williams (2002) and Fricker (1998, 2007) resist any suggestion that objectivity can be reduced to intersubjectivity, even as they do insist on the social nature of knowledge and the need for ethical values to generate knowledge and sustain the proper – socially set – metaphysical criteria for it. Rorty's critique may affect, however, the views that have been defended by Misak and Longino, to the extent that they appear to commit themselves to a (hypothetical) reductive definition of, respectively, truth and knowledge (Misak, 1991/2004; Longino, 1990, 2002). In the case of Longino's view, I believe this

is not the case as she emphasizes epistemological pluralism and a partial localized conception of knowledge alongside a realist metaphysics (Longino, 2002, p. 93). It may be thought to affect the communitarian of account by Kusch, too, except that Kusch opts for a metaphysical solution, namely relativism (Kusch, 2002, pp. 174–175).

Throughout this thesis I have emphasized both that the concept of knowledge is indefinable, that it is essentially "commonable", but there is no way to reduce either truth or knowledge to idealized social norms, such as idealized rational acceptability. Those working in the neo-pragmatist tradition, such as Misak, may be drawn to such an identification, but on my interpretation neither Fricker, Williams, nor myself take this step. It would, after all, form part of the kind of reductive definition of knowledge that knowledge first theorists reject. It is also a position for the coherence of which the two genealogical accounts have offered the resources, as the genealogical method is equipped to give us "an explanation without reduction" of the intrinsic value of truth, and of the social nature of knowledge (Williams, 2002, p. 90).

Yet, it is important to recognize – as we have noted at the start of this chapter – that if knowledge is social in a strong sense of being produced socially and dependent on social processes, then there are continually real-life issues when it comes to securing objectivity. A social conception in fact helps to bring into focus the real-life issues, which we – as a community – face in our efforts to achieve knowledge. And, conversely, a view that tries to deny that knowledge is social in order to sustain the importance of individual understanding and hence of individual intellectual authority, so as to resist authoritarianism, "may lead us to neglect the ethical issues" (Welbourne, 1986, p. 65), which issues are of particular importance precisely so we understand that the threat of intellectual authoritarianism exists (Welbourne, 1986, p. 67).

One reason that Alvin Goldman has resisted a move away from individualism is that he sought to guard against relativism, according to Longino, because he assumes a dichotomy between sociality and rationality, which Longino argues is false (Longino, 2002, p. 42). Goldman's version of social epistemology is veritistic; i.e. "concerned with the production of knowledge (…) in the 'weak' sense of *true belief*" (Goldman, 1999, p. 5). Longino describes it as an individualistic, aggregative form of sociality (Longino, 2002, p. 43).

Reviewing Longino's work *The Fate of Knowledge*, Goldman (2002) complained that ethical norms do not help us to secure epistemic success. He believed that Longino had not made it plausible that her four criteria for successful social epistemic cooperation actually make it more likely for us to arrive at truths. [8] On the face of it, such criticism appears justified.

---

[8] These four are discussed in Chapter Chapter VIII section 5. They are: 1) Venues. 2. Uptake. 3. Public standards. 4. Tempered Equality (Longino, 2002, p. 129-131). It should be noted that the last criterion is an epistemic one, whereby abilities play a role. Longino explains that this is why this is why equality must be tempered according to the "role and purpose of this criterion" (Longino, 2002, p. 131). Its role is mainly to ensure that, even despite differences in political or social power, "all relevant perspectives are represented" (Longino, 2002, p. 131).

But a first thing to note is that Longino did not claim that her criteria would improve on the likelihood of truths. Instead, she has argued that these criteria determine the degree to which a scientific community is objective, as conditions that "are necessary to achieve the transformative dimension of critical discourse" (Longino, 1990, p. 76). The guiding thought, there, is that objectivity requires criticism from alternative viewpoints. So, what makes a social conception, coupled to epistemic norms, better than one which is individualistic and allows for conjoined efforts? And why it is that the specific chosen norms/structures improve, as opposed to distract from, the epistemic value of the outcome?

To tackle these questions, which are part of Goldman's critique – and that of other metaphysical realists who claim we need to be methodological individualists, prioritizing the individual above the community in order to refrain from any threat of relativizing the notion of truth – can be answered in a number of ways. A main insight is that there is no agreement on some fundamental questions about what the social nature of knowledge entails. But to insist that the view differs from his, and that he assumes a false dichotomy, is a weak reply. It would appear that his challenge of explaining how morality can secure epistemic success is a real one.

To address that challenge however, some rejoinders can be put forward. The first one is that the worry of relativizing truth is not appropriate, as I explained in reply to Rorty. But this is to say, again, that the dichotomies are false, and may not be the best way to structure these issues. The second rejoinder is to engage with the critique. That means accounting for the ethical and political norms that we would need to govern our epistemic conduct in order to arrive at knowledge; and, if possible, meeting the challenge of "proving" that they are the right ones. That answer is one I try and give, at the end of the thesis, but the chance that it will convince depends on one's leniency for accepting the genealogical account, and on agreement as to how the norms of epistemic justice are to be spelled out.

As a preliminary answer, and one that engages at the right level of abstraction, one can point out that *a naturalized epistemology* has some starting assumptions about the ways in which knowledge is validated. One of these assumptions is that our criteria are subject to revision, when we find out that they have not functioned well or that they cannot accommodate our experience. This point is owed to Quine and Ullian (1978), who argued that our web of belief is subject to revision and that we use beliefs to validate the truth of other beliefs in a coherentist way. While I maintain that we should focus on knowledge as the unit of our analysis and not on true beliefs, I hold, similarly, that what currently qualifies as knowledge is our standard for the validation of future knowledge claims, but that this does not block radical revisions when we have recalcitrant experiences (knowledge). Even within a knowledge first and coherentist framework there is scope for radical revision.

This serves to show that there is room, within the social generation of knowledge, to

revise the criteria for knowledge and for credibility. This does not yet provide an answer as to why the specific recommendations made by, e.g. Longino (1990, 1995, 2002), would work to increase and improve our knowledge. But it is to argue that the process of changing the practice, and any amendments to it, are subject to review. A practice which does not function well, or any better than the previous one, will simply not last. This is also a point Longino makes (Longino, 2002, p. 137).

Two points can be made here, in support of Longino's proposals. First, the conditions are to enable knowledge and are not framed as the sole criteria for the determination of knowledge. So, casting them as the sole criteria, as does Goldman in his review is a mistake (Goldman, 2002). For we bring our logic and observation to bear as part of all the individual judgments that are being made and pooled (Longino, 2002, p. 139).

Second, the conditions that are set are to ensure that all relevant sources are included. In so far as preventing (shared) bias is the aim of Longino's insistence on both diversity and on enabling the possibility of contestation, her criteria will help to ensure that the knowledge a community generates is tested from a wide variety of perspectives and that it can withstand further scrutiny. As such, it is "objective". Indeed, there is some proof in the form of the Hong and Page theorem that diversity ensures better epistemic outcomes (Hong and Page, 2012, p. 67, pp. 70–71).

The point I am making here is not that intersubjectivity is the sort of objectivity that we should seek; I have already explained why that is no part of my view. Longino's proposals for the proper structuring of scientific debate and inclusion of outsiders will serve to reduce the biases that are otherwise able to distort the outcomes – and they are based on an analysis of cases where such biases could come into play (Longino, 1990, 2002).

Aside from the epistemic reasons, there can be political reasons for inclusion. Such reasons have been given by Helen Landemore (2013), David Estlund (2008), Philip Kitcher (2003) and recently by Iris Marion Young (2000). It is part of my effort, in this thesis, and specifically in chapter VII and VIII to explain why and how political (as well as moral) considerations of the sort they adduce are relevant from a standpoint of epistemic justice, too. Their relevance, so I will explain, is due to the fact that epistemic justice is a subcategory of justice overall.

That I note this here as a separate point is in order to flag that the issue of democratization is one that has not always or primarily been associated with the best epistemic results or outcomes, but there is a possible broader perspective here that very much bears on epistemic outcomes which we may need to take into account, too. This is because and to the extent that epistemic justice cannot exist in isolation of our societal arrangements – while we need it if we are to achieve social epistemic success, on a realist conception of what such success is. It seems reasonable to conjecture (and something we may investigate) that epistemic justice requires other forms of justice, too. These other forms of justice can, conceivably, yield additional constraints on our epistemic interactions and societal epistemic arrangements in turn.

Conclusion

This chapter opened with the question of how it was so much as possible to explain how the social nature of knowledge without the conclusion that "anything goes" in inquiry? I have argued that it is. Moreover, I have argued that if knowledge is social (as this idea is developed by Craig and Williams), then it is required that people care for truth and objectivity. However, I have also explained that if knowledge is social, then because it is generated socially and has a social life, this means that we can further understand why it is dangerous to reduce knowledge to social agreement. This is why we should also refrain from giving any kind of instrumental/functional definition of knowledge, according to which knowledge is a social institution/phenomenon, i.e. a social status, for the reasons that knowledge first epistemologists have pointed out.

I have argued that there are both ethical and political aspects to knowledge. Even while there is no conceptual problem that stems from the social nature of knowledge, we do have to be aware that there are real life challenges that stem from it. These challenges are of an ethical and/or political nature, and accordingly, they require both awareness and a workable, ethical and/or political solution. I have argued that epistemic justice is the notion that we require.

I will define and explain the notion of epistemic justice more closely in Chapters VI, VII and VIII. However, over the next two chapters I turn to the issue of testimony. I will demonstrate that testimony has been misconceived so far, because – other than orthodoxy about testimony has it – testimony is to be conceived as a practice, and moreover as a practice that serves to generate knowledge (by our communally filtering claims to knowledge and identifying reliable sources as well as promulgating their claims as knowledge). So far, testimony has only been seen as a means of knowledge transmission, but that would mean it could at most serve to sustain and disseminate knowledge.

Chapter IV | Testimony and Epistemic Individualism

Introduction

In philosophical theorizing, philosophers have used the term "testimony" to indicate the communicative interaction(s) of an informative nature, whereby – putatively – knowledge is being shared (Coady, 1994, p. 7; E. Fricker, 2006). The main problem of testimony has been to explain how it is possible that knowledge be gained through "mere say-so" (Welbourne, 1986, p. 1). Clearly, that topic is relevant for my thesis, for the reason that understanding whether or not knowledge should be shared presupposes that we understand what knowledge sharing comes to. Moreover, it presupposes that such sharing is possible. The questions that I set out to answer in this chapter are accordingly: (1) what philosophical problem is posed by testimony, and (2) can this problem be (dis) solved satisfactorily?

I will first explain that current theories of testimony have taken for granted that testimony is at best a transmissive source of knowledge. This is the first indication that the puzzle of testimony is premised on an individualistic conception of knowledge; but as long as knowledge can be transmitted, there appears to be no problem. Next, I discuss how the problem of testimony has been analyzed, and how theorists have sought to solve it. Here I argue that the puzzle has, as one of its premises, the claim that an individual can only gain knowledge through testimony if he or she has an *internalist* justification for accepting it as knowledge. That is, *externalist* solutions to the puzzle, or solutions that stipulate a general entitlement whereby the hearer may appear to be epistemically off the hook too readily, tend be seen as "unsatisfactory". This is a point made by Faulkner, who explains that the challenge of the reductionist about testimony is inherently a question of an internalist nature.

Hence, those who regard the acquisition of knowledge by testimony as puzzling require that the person who gains knowledge be viewed as epistemically responsible for the knowledge gained. I endorse this requirement as an apt description of our practice: we do tend to think that the person who accepts a piece of testimony as knowledge should (moral-epistemically) not do so blindly. The credo which is intuitively correct, so I believe, when it comes to accepting knowledge as presented by someone else is: "trust, but verify". While that may appear contradictory, Michael Williams too maintains – building on work by Robert Brandom – that epistemic authority conforms to a "default

and query structure" (M. Williams, 2015, p. 179). [1]

However, I will also argue that the puzzle of testimony, conceived as one that endorses this internalist requirement on justification, cannot be solved satisfactorily if one assumes an individualistic conception of knowledge. To this end, I discuss the various solutions given but will focus on the account by Faulkner. Accordingly, this chapter serves to set the stage for my argument in the next chapter that the puzzle can be dissolved – and that a convincing error theory can be offered for why the puzzle arose – by accepting that knowledge is social, whereby testimony is a practice that serves to socially generate knowledge, and whereby the dissemination of knowledge claims is part of this function.

2. The Puzzle of Testimony

A first thing to note is that a social conception of knowledge, as defended by Welbourne, starts with the assumption that knowledge is shareable. As we saw, Welbourne maintains that knowledge is essentially commonable, i.e. that it can be had by virtue of mere say-so. At the same time, the stipulation that knowledge is commonable and that testimony is central to knowledge appears to require some explanation. Why or how is knowledge commonable – when others have disputed this? What can be said in favor of his account of testimony?

I would add to this first point that in Miranda Fricker's work on epistemic injustice testimony is conceived of as a social epistemic *practice* through which knowledge is gained or lost. Her definition of testimony, then, is somewhat broader than the conception of testimony held by most philosophers, which tends to be that of a one-off exchange. The idea of testimony as a practice involves the idea that there are many interactions each of which is part of the practice. While she explains that her conception of testimonial practice is broad, such as to also encompass for instance the questions that may be asked, Fricker does not, however, argue that testimony is to be conceived of as a social epistemic practice (Fricker, 2007, 2010; Hookway, 2010). This is my aim in this chapter and the next.

It appears that a social conception of knowledge is not in play when philosophers try and explain the possibility of gaining knowledge through testimony. For what is considered puzzling about testimony is that it involves epistemic dependence on others. Indeed, in so far as knowledge through testimony has been deemed possible, the suspicion

---

1 Michael Williams explains that he previously used Robert Brandom's idea of "default and challenge" to describe the structure epistemic authority conforms to, but that not all questions are directly challenging another's authority. The question, "How do you know?" may be asking merely for an explanation, rather than implying that one questions another's authority, which is a point made by Austin (M. Williams, 2015, p. 179; Austin, 1946/1979, p. 78).

has been that such knowledge has a special status, yielding the term "testimonial knowledge", whereby there is some hint that it is not on a par with knowledge gained through other sources.

The difference, as explained by Audi, is that, even while most philosophers regard testimony as a source, and as foundational for much of our knowledge, they also regard it as a source that is only transmissive of already existent knowledge, whereby that knowledge was based on other sources (Audi, 2013, p. 510). Testimony does not generate new knowledge then, even as the idea of a source may appear to suggest that testimony is generative of knowledge. One way to express this point is that knowledge is a source of second-hand knowledge (E. Fricker, 2006). However, the position that testimony is foundational while it is not generative appears somewhat contradictory, as Audi admits by asking: "If, however, testimony is transmissional rather than generative, how can it be a foundation for knowledge?" (Audi, 2013, p. 510)

To this question, Audi replies that testimony confers knowledge without there being a need for inference on the part of the receiver. That explains, so he believes, why testimony is a foundational for knowledge, as it can yield "basic knowledge" in the sense of yielding knowledge non-inferentially. Even so, he maintains that we should not cast testimony as a basic source of knowledge; because only a generative source is basic. Such a source provides a justification as well as knowledge. Testimony, on his view, provides us with knowledge, but not with the justification that the person who passed on that knowledge held:

> You are my source of information, but your knowledge is not the ground of mine or my evidence for p. It is not because you know that I do; it is because you tell me (given that you know). Your knowledge that p is required for successful transmission, but my knowledge is not based on your knowledge, if this entails more than its appropriately depending on it. (Audi, 2013, p. 510)

How is the fact that so much knowledge is gained through testimony compatible with the idea that testimony does not generate knowledge? For, the fact that testimony is transmissive only, and yet somehow *foundational*, would make it appear that we all come to have knowledge which stems from some very few originating sources. Moreover, the idea that knowledge is passed on non-inferentially would appear to leave many false claims of knowledge in a position where they have not been corrected.

In my view both these points can be answered. On the first point, one may argue that the pervasiveness of testimonial knowledge does not entail that knowledge lacks a base, even while there are perhaps fewer originating sources than we might tend to think. On second point, one may argue that false claims do not qualify as being knowledge, and hence that there is no need to correct for claims that pass as knowledge from an epistemic viewpoint. Yet, the question that looms here is: What room is there, on a non-inferential account, for a critical assessment of knowledge claims? That is, how is

there room for us to separate the true from the false and resist the claims that do not properly qualify as knowledge, when it is stipulated that knowledge is communicated non-inferentially?

It would appear that "testimonial knowledge" – as the kind of special species which that term intends to pick out – is particularly weak in its justification, if there is no way in which we can explain that we have acquired it via testimony, as opposed to merely false belief. For this reason alone we should seek to ensure that our account of testimony includes what Faulkner has sought to provide: an account of the role of the hearer, and specifically, an account of what rationally explains or justifies a hearer's "uptake", which Faulkner intends to mean acceptance of testimony that p which issues in a state of believing that p (Faulkner, 2011, p. 19). Note that that technical term is needed, as Faulkner explains, because acceptance may be differentiated from belief. Acceptance of some piece of testimony that p may not need to issue in believing that p, if one considers is false; or one may already have believed it on other grounds such that the testimony is not one's reason for believing it (Faulkner, 2011, p. 19).

Strikingly, Audi himself makes a point – which is also very commonly made by epistemologists – that a person who is on the receiving end of testimony should not be gullible (Audi, 2013, p. 525). Hearers should exercise their critical capacities, or, alternatively put, hearers should ensure that their trusting dispositions are in good shape (Fricker, 2006; Audi, 2013). While the emphasis on the hearer is quite correct, what is notable is that Audi would appear to rest content with mere reliability: a reliable sensitivity on the side of the hearer to trust those who are trustworthy. The problem, here, is that this assumes that the hearer can be a reliable tracker of trustworthiness. Positing that kind of ability, we may as well assume that the epistemic criteria for acceptance are met. Yet, it could well be that the hearer is not very reliable and, in that case, whatever he acquires by way of "testimonial knowledge" has only a fairly low status – and can certainly not be regarded as on a par with, or as a basis for, other types of knowledge which he acquired on his own.

As for Audi's account, it appears that splitting the notions of knowledge and justification – and specifically that knowledge entails justification – is something we should not do as Faulkner also says in connection with Audi's earlier work (Faulkner 2011, p. 13; Audi, 1997). Faulkner argues that in many cases the explanation of how knowledge is transmitted will also serve to explain how justification is transmitted – e.g. on an assurance theory of testimony, in a case where a speaker provides assurance and thereby takes epistemic responsibility, what the hearer ends up with is either knowledge or justified belief, depending on whether or not the speaker's justification was "knowledge supporting". As such, there appears to be no reason, Faulkner argues, to make a distinction between knowledge and justification theoretically fundamental in the way that Audi proposes.

Audi himself has noted, in a later paper, that his idea of justification may be an

internalist one (Audi, 2013, p. 511, footnote 5 and 6). While he rejects this, I believe that his take on justification being an internalist one may explain why he thinks knowledge can be transmitted while justification cannot be transmitted. For, as Faulkner notes, the justification that a hearer obtains by trusting a speaker (on the assurance view) is externalist in nature – both for knowledge and for justified belief. It depends, so Faulkner states, not on other things the hearer believes, but on the other things the speaker believes (Faulkner, 2011, p. 14). Quite notably then, what Audi appears to resist despite wanting to allow for a special category of testimonial knowledge, is that we should come to accept the idea that an externalist justification is applicable across the board as sufficing for knowledge – and particularly, for beliefs being justified, no matter whether they are based on testimony. But this does make testimonial knowledge different from other types of knowledge, and – akin to Faulkner's point that we should not make unnecessary distinctions – I believe that that is a conclusion that we should resist.

3. The Assumption of Internalism

The main reason why the puzzle of testimony arises is the idea that knowledge requires justification, while that justification requires an individual inquirer having access to reasons in the form of premises that can be cited in support of the knowledge in question. For, as we tend to view knowledge from within an individualistic conception, knowledge is justified true belief. For someone to qualify as knowing it should be possible for that person to give a proper justification for the belief on an internalist conception.

Of course, there is the view that justification may not need to be available to someone for that person to be justified or to know. Externalists about justification hold that being justified by some externally describable standard is enough for knowledge. Whether or not someone is justified is externally determined by the facts – and may be determined by someone else (the one who investigates the issue, say). For a theorist can rationally explain that a person who has formed a belief in a given way is justified in having the belief, such as to qualify as having the knowledge they claim to have (if the belief is true).

Analyzed from a distance, those who tend towards externalism about justification care mainly that justification likely yields knowledge, and that we can explain why knowledge is had in the many cases where it appears it is had; while those theorists who tend toward internalism about justification will tend to care more about the formation of beliefs, i.e. beliefs being supported by reasons, rather than whether the subject has knowledge or that all knowledge which we think should qualify as knowledge is recognized as such (M. Williams, 2015).

Indeed, the game of explaining that knowledge is acquired through testimony is

different on an internalist and on an externalist understanding of what this involves. On an internalist conception, a person should have reasons available to him or her; on an externalist conception, the person should be entitled to knowledge because there are reasons for their epistemically qualifying as knowers even while the reasons are not available to them. Internalism and externalism about justification are mutually exclusive. Even so, both have some appeal. In everyday life, in the cases where we demand reasons, we tend to think of knowledge as requiring an internalist justification. At the same time, we are satisfied, quite often, by learning empirical facts from people who have been verified as proper sources by others, i.e. who have a good track record and no apparent reason to lie and even, frequently, reasons against lying, both as a moral rule and as potentially damaging their reputation as an authoritative source and/or as a moral person.

We tend to think we can have knowledge for the asking and that we are entitled – externally justified – in accepting what others say, even if we lack access to their reasons. Tyler Burge, in his paper "Content Preservation", has given an account of testimony which appears especially appealing here, whereby we have an a priori entitlement in the form of an externalist warrant based on the general reliability of testimony (Burge, 1993). What Burge makes available to us, then, is the idea that we can have second-hand knowledge based on an external justification, even as it is possible to reflect upon our entitlement and then we may have an internalist justification – or at least an a priori justification – for our knowledge qualifying as such, after all.

However, as Williams, Moran and Faulkner have independently pointed out, the puzzle of testimony arises because we want more than a general entitlement: we want to know that and why knowledge can be gained in individual cases (Faulkner, 2011, p. 11; Moran, 2006, p. 275; Williams, 2002, p. 85). This requirement is an internalist one, one may think, but it is also, and mainly, a theoretical demand for a better insight into the factors that make for (gaining) knowledge through testimony. An entitlement need not actually exist in all circumstances, for Burge's view is that the entitlement which someone else had is being transmitted from one person to another. But there is no conferral of entitlement when there was no entitlement to start with. Deeming someone trustworthy is not a self-fulfilling prophesy on Burge's externalist view: no epistemic reason or entitlement can be transmitted unless the speaker already held it (Burge, 1993, p. 486).

In distinction from Burge, Williams and Moran insist that while testimony can be viewed as akin to other sources of knowledge such as perception – whereby given that we are more often justified than not, there is indeed a general entitlement and a justification for reliance on the source in some general sense – there is a crucial difference between testimony and our own onboard epistemic sources such as perception and reason (Burge, 1993, p. 466; Williams, 2002, p. 111; Moran, 2006, p. 281 and footnote 7). That difference is that in testimony we rely on others in a way that depends on the quality of another's

will. Since we willingly depend on another person, and may decide not to trust someone, both the giving of testimony and the accepting of it are voluntary acts. Faulkner usefully coins the word "uptake" for this voluntary acceptance by the hearer (Faulkner, 2011, p. 19). The idea of these acts being voluntary brings with it that we are, as individuals, morally and epistemically responsible or liable for our roles within a testimonial encounter (Moran, 2006, p. 283; Williams, 2002, p. 119).

Faulkner argues that the puzzle of testimony requires, then, an internalist interpretation of the justification that – on the epistemic individualist view – knowledge requires, because a mere fact of reliability in most cases does not suffice to explain why reliance on testimony is justified in a given case: "where reliance is uptake what is required is an epistemic reason or a reason for belief" (Faulkner, 2011, p. 41). Taking the interdependence of the two parties who partake in a testimonial exchange seriously, Faulkner casts the problem of testimony as a social coordination problem and specifically as "the problem of cooperation" (Faulkner, 2011. p. 7).

4. The Assurance View as a Second Personal Account of Testimony

The puzzle of testimony in epistemology is to explain how taking another person's word could suffice for having a justified belief. In his recent work, Richard Moran has "solved" this puzzle by arguing for a picture of testimony as inviting trust. The central notion of his account is that of telling (This account was also, separately, defended by Edward Hinchman (2005) and McMyler (2011).

On this picture of testimony, the speaker who asserts in a testimonial fashion "that p" undertakes epistemic responsibility for bringing her addressee to believe that p. The invitation by the speaker to trust her in this way "at her word" makes available an epistemic entitlement to the addressee for believing that p. If the belief that p is indeed true, a person A could come to have knowledge by testimony since A's epistemic entitlement for believing that p is based on (rightful) trust in the speaker S and the transmission of entitlement from speaker to hearer via distinctively second personal reasons (Moran, 2013; Darwall, 2004).

This account has four distinctive features: The first is that trust plays a key role. Specifically, it concerns trust in the speaker. In my interpretation, this is spelled out as trust in the speaker as a morally accountable epistemic agent. Both Hinchman and Moran stress the speaker's moral accountability for inviting her addressee to believe that p as she is responsible for the consequences of her purporting to have epistemic authority on the matter whether or not p. Thus, trust in the speaker is not merely trust in the speaker as a reliable source of information on p. Just as virtue ethicists draw a distinction between virtuous agents and virtuous acts, so the assurantist draws a

distinction between trusting the speaker and evaluating the reliability of what is said. The speaker's epistemic reliability is merely a presupposed background condition for the picture presented. The epistemic agent is – unless there are countervailing reasons – trusted to appraise herself as a reliable source on p if and only if and to the extent that she is a reliable source on p. This trust in the speaker is the basis of why she can be held accountable.

The second distinctive feature is that sincerity is a matter of the speaker's intent. Sincerity is assumed in the sense that the speaker may not intentionally mislead the other or betray the trust placed in her by her addressee. If the speaker happens to have a warranted, but false, belief, then she is still sincere as long as she had the right intention. She herself believes it, and thus had no intention to deceive her audience, or to betray the audience who would accept the belief at her word. But note that this need not mean that a speaker is only sincere in asserting "that p" if, in doing so, she accurately represents her own beliefs. If the speaker does not have the belief herself, but she presents it as a belief that is held to be true by certain authoritative others, the audience could be brought to believe it with entitlement. Notably then sincerity is not a matter of the speaker's adequately depicting her personal beliefs. Instead, sincerity lies in the speaker's intent to ensure that the addressee who believes her at her word will thereby only come to have knowledge.

The third aspect – which Hinchman makes explicit – is that on this account trust in the speaker yields an entitlement prior to the assessment of the speaker's reliability (Hinchman, 2005, p. 578). The speaker may be trusted when there are no apparent reasons not to trust the speaker. Trust is the default (as Williams might add in the "right" kind of conversational encounter). Trust may be assumed by default in each individual case, unless there are specific reasons to withhold trust. This is defended in a very strong sense: even if the speaker S has herself no (or no sufficient) epistemic entitlement to the belief that p, S's presenting the belief as true – without there being countervailing reasons for questioning her authority or for doubting that p on the basis of background information – creates an epistemic entitlement for her addressee to believe that p. This has the strange consequence that one could be entitled to believe p as asserted in testimony on the basis of sheer trust in the speaker, even when that trust is not in fact justified, as the speaker has not met the requirements for rightful epistemic trust.

The fourth dimension is that the context of testimony is typically an interaction between two parties: the speaker and her addressee(s). As such there is an accountability relation and directedness toward her intended audience by the speaker. In asserting, the speaker aims to bring about that her addressee(s) come to have a belief which she presents to them. Therefore, we are not talking of just any speech act nor indeed of mere assertion: telling is a special case.

The relation between the persons involved is relevant: the fact that the speaker addresses an audience means she is requesting the trust of these specific individuals.

Moran emphasizes the importance in his own view of the idea of a second-personal reason (Moran, 2013). In asserting in testimonial fashion the speaker will only take responsibility for the beliefs of those to whom she has imparted her information. In doing this she makes a special effort to be understood as well as trusted, by providing reasons for her beliefs, if she can, upon being asked to do so. Neither her responsibility nor her effort extend to bystanders who might pick up (part of) what she says, or who may give it a different interpretation from that intended by her. Bystanders therefore do not gain an epistemic entitlement, and cannot hold S accountable (Hinchman, 2005, pp. 566-570, p. 576; Moran, 2013). Instead they have merely gained some evidence when they have overheard S's contention that p. The relevant distinction can be drawn by examining the underlying intentions in communication and the second personal reason that relates speaker to hearer – but not speaker to bystander.

Moran contrasts his view of testimony with the view that he claims is dominant in the literature: he also attributes it to Bernard Williams (Moran, 2005, 339). He labels the two approaches, respectively, his own "assurance view" and the "evidential view". Moran's account of testimony is that it encompasses only a subclass of assertions, namely those assertions that are cases of S's telling A that p. In telling her audience that p, the speaker assures her audience that p. This contrasts with the account of testimony which focuses on "mere" assertions and which Moran calls "the evidential view" (Moran, 2006). On the evidential view, a speaker who asserts that p presents her utterance as evidence for her audience that the asserted belief is true (Moran, 2006; E. Fricker, 2006). Evidence is then assessed using the standard forms of inductive inference and second personal reasons do not feature as relevant on the evidential view.

In my view, the difference between assurantism and evidentialism can be usefully characterized in terms of which of the two values the views are each most concerned to uphold: trust or truth. In making this point, I am indebted to Alasdair MacIntyre's meta-analysis of how our major moral traditions view lying, discussed in Chapter III (MacIntyre, 1995). We saw that MacIntyre argued with respect to the two main moral traditions – utilitarianism and contractualism – that their stance on lying may be understood according to which of these two values they see lying as being an offense. According to MacIntyre, the difference in perspectives is that for utilitarians lying is an offense against trust, while deontologists regard it as an offense against truth (MacIntyre, 1995, p. 335). I here argue for a very similar meta-philosophical perspective with regard to the two approaches to testimony. The aim is to point to the underlying tension between truth and trust in epistemic cooperation, and to show that the demands of truthfulness and trust can conflict. They are both fundamental yet irreducible to one another so there is no possible "trade off" between them.

On the assurance view, assertions in the case of telling are conceived as an invitation to *trust*. The speaker is asking to be believed – on trust. On this view the assertion is seen not as evidence for the truth of what is asserted, but as a reason for the audience to

acquire the belief on the basis of trust in the speaker who stands in a distinctive, "second personal" relation to the assertee (and is not in the same relation to anyone else such as a bystander). Assertion is a speaker's way of acknowledging her epistemic responsibility in moving the hearer to acquire the belief she expressed. By using the mode of telling the utterance is explicitly given epistemic significance by the speaker.

Having purposely picked this mode of communication, the speaker is answerable to what she has asserted. She is raising the expectation that she is telling the truth. Trust in the speaker is distinctive of assurantism even to the extent that the speaker's epistemic reliability is not assessed, but is presupposed. Rather, the speaker needs to have sufficient self-awareness and act as a responsible moral agent when she takes an epistemic stance and purports to have epistemic authority. Sincerity is required, but it lies in the intention – not in the accuracy of the depiction of her own beliefs, or even the truth of these beliefs – though of course she will have to answer to criticism if the warrant that she needed to have for presenting the belief as true is lacking: she unjustifiably purported to have epistemic authority on the matter of whether or not p.

By contrast, on the evidential view assertions are taken to constitute evidence for the *truth* of what is asserted. The idea here may appear to be that people generally speak the truth when they assert, hence, an assertion to the effect that p is evidence for p. Here it should be noted that there is no clear epistemic role for the speaker's intention to convey truth. What matters is whether the speaker is truthful, in the sense of not misleading the other, and in the sense of being accurate and reliable as a source. Qua source of knowledge, evidentialism applies the standard inductive procedures for the assessment of evidence and that a speaker testifies in a particular way is "just more evidence" to be assessed by the standard canons of inductive inference.

On this evidential view, it remains the independent task of the audience to ascertain the truth of p. Indeed, on this view, it is rather hard to separate the verbal behavior of assertion from other (verbal) behavior. Though the speaker has presented the utterance as true, this speech act is seen as itself mere evidence. It is evidence in the same way that pointing to something which is lying around at the scene of a crime provides evidence. Likewise, an assertion that p is true is viewed merely as a way of asking attention for considering it: the linguistic analogue of pointing or directing attention. For evidentialism, responsibility on the part of the speaker – other than being sincere and accurate – is not assumed. Responsibility resides with the one who takes seriously (or not) that which is asserted. There is no shared, "second personal" reason between speaker and recipient and no shared responsibility. Epistemic responsibility devolves entirely on the hearer: any hearer (including bystanders).

Therefore, evidentialism can be characterized as primarily concerned with the truth of what is asserted. Assertions are conceived merely as an evidentially grounded means to the truth. A speaker who asserts p therewith presents evidence that p is true. The evidence however is not itself a reason to believe what is asserted. For believing that that

which is asserted is true, the hearer will either need a good rationale or independent verification. It now has two separable factors: inductive appraisal of the truth of p and inductive appraisal of the reliability of the assertor. In contrast with the Hinchman-Moran picture, the audience will need to assess and test for the reliability – as well as the sincerity – of the informant in every individual case. To not do so would be to give up on the idea that one's beliefs have to be justified, i.e. that one has to ensure that one's source is reliable. The canons of evidence are orthodox: the standard inductive appraisal of evidence.

One may wonder whether the assurance view and the evidential view are really incompatible. In one sense, the evidentialist can easily accommodate the claims of assurantism: the facts pointed to by the assurantist are simply added to the list of evidence as "more evidence". More generally, are both views conflicting ways of tackling the problem of testimony as Moran suggests, or do they simply focus on different aspects of the puzzle? Hinchman, for one, has suggested that both approaches work in their own terms. He claims that assertions that p normally provide mere evidence that p is true, but that those sincere assertions which are instances of telling are special because the speaker who infuses this mode of speaking therewith creates, and makes available, an epistemic entitlement to her addressee (Hinchman, 2005, p. 566-567). Assume for now that the views are compatible if applied to different cases. Do the "solutions" work?

My view is that neither work. There is a problem of bootstrapping for both assurantism and evidentialism. The evidential view has this problem as the need for evidence means that one in fact will have to assess the reliability of the informant, up to the point where one is assessing the truth of p, in order to be epistemically justified at all. Evidentialism attaches no role to trust. It substitutes, for the second personal normative commitment between teller and recipient, the fact that this commitment obtains: simply more evidence. This is a reason to be suspicious of the view, and indeed also a reason to try and modify it. The main problem however is that evidentialism will not be able to bridge the gap between truth and trust, because its sole emphasis is on truth. Therefore, I do not find Hinchman's proposed reconciliation between the two views plausible. Saying S is trusted only after we have assessed that S is reliable finds no role for trust: it occludes its special role by incorporating it into the category of inductive evidence.

It also amounts to bootstrapping, in this form: in a common case we cannot judge S to be reliable on the matter p without finding out whether p is true or not. Accordingly, evidentialism does not explain why we could justifiably take a belief that p at someone's word – or explains only a very small number of cases, the number of which depends on one's ability to trust in spite of having no full certainty about S's reliability. Evidential assessment simply bypasses trust. It plays no independent role in justifying knowledge acquired by testimony.

There appears only one way out for evidentialism which is to espouse an externalist view of epistemic justification. Externalism about epistemic justification is the view that

we do not need to understand or have access to our reason or justification for why we hold our belief or have accepted it as true, in order to be justified (Burge, 1993, p. 458; see also Lackey, 2008, pp. 9-10; Faulkner, 2011, p.10). That is, we need not be able to tell the difference between, but nonetheless are justified if and only if we have a good justification by some external criterion. We will need to show only that testimony is generally a reliable method of belief acquisition. As we saw, both Williams and Moran have argued that this is not satisfactory: the challenge posed by a sceptic about testimony is to show that trusting another is epistemically justified in an individual case (Williams, 2002, p. 85; Moran, 2006, p. 275). Otherwise our epistemic responsibility only relates to choosing a generally reliable method, and to watching out for any signs that the method might be flawed. Faulkner too views the reductive challenge as an internalist one of explaining how acceptance of testimony is "subjectively reasonable (Faulkner, 2011, p. 11).

I would argue, however, that assurantism has a very similar bootstrapping problem. Recall that on the Hinchman-Moran picture, the speaker explicitly takes on a burden of responsibility, and therewith makes available an epistemic entitlement for the hearer. The hearer who recognizes the invitation to trust the speaker at her word and who also accepts that invitation will come to have an epistemic entitlement to the belief in question as long as he does not have reasons – such as the speaker's track record or apparent insincerity – to withhold his trust. The hearer, however, *need not assess or test* the speaker's epistemic reliability on the matter at hand before taking her at her word. The entitlement, so Hinchman argues, has already become available: it has to be so, because on the picture of testimony as telling, which illocutionary act when consumed gives one a direct entitlement, trust is prior to assessment (Hinchman, 2005, p. 564, p. 578). Assurantism thus emphasizes that we can and do trust others, but does not give us an explanation when and why we can trust. It lacks any idea of epistemic standards and of trusting responsibly in an individual case. As such it does not alleviate the worry that simply trusting cannot be a basis for epistemic justification.

The bootstrapping problem for assurantism thus consists in the fact that the epistemic justification for believing that p on the basis of testimony is not linked to the likelihood of p being true. The justification is conceptually only tied to trust. Here, as I analyze it, there is so much emphasis on trust, that someone concerned with truth might wonder whether the trust is justified. It seems the hearer will need to assess that this is so, in order to ensure he does not trust a speaker who is not trustworthy. Yet there is no room on this picture for an assessment of trustworthiness. Trust, sheer trust, creates an entitlement, even if the source is not at all reliable.

Of course, this could be remedied by some background mechanisms in the general account that are argued to be preconditions of the practice and not represented within it. However, as Hinchman has rightly pointed out, this could not go as far as to demand an assessment of the reliability of the speaker who tells you "that p" (Hinchman, 2005, p. 581). It is indeed true that a speaker with a (readily) known bad track record should not

be trusted. And – if it is known – a speaker who has betrayed the trust in her on a given matter will not be trusted again. (At any rate, not regarding that area of expertise.) But there is no room on the assurance view to argue that trust needs to be earned or rightful for the hearer to receive an entitlement to believe that p. It need not be. Assurantism offers no solution to understanding our epistemic entitlement to beliefs we acquire through testimony, as precisely an assessment of reliability would be required in order to uphold the value of truth alongside that of trust.

The "pure" evidentialist view and the "pure" assurance view (assuming that they are compatible, as they are respectively concerned with cases of assertion and of telling) are thus inadequate as solutions to the puzzle of testimony because they refuse to accord to truth and trust an equally fundamental status. Focusing on one value to the detriment of the other is unsatisfactory. It shows that the values cannot be reduced to one another. They are each fundamental for the success of epistemic cooperation. To solve the puzzle of testimony it needs to be shown, instead, how they are aligned. In doing so, I believe that dropping the individualist assumption is key, since that opens the way to epistemic cooperation, and hence to a proper dynamic of truth and trust in which both parties have a shared epistemic responsibility. Unless we acknowledge the role played by shared responsibility in testimony, neither account has room for the idea of a speaker/hearer living up to the expectations that each party has of the other. The dynamic nature of testimony is not adequately captured as a result of this. Nor is there room for understanding the ethical dimension of testimonial exchanges. There is a case, then, for the individualist assumption to be dropped.

5. Faulkner's Critique of Second Personal Normativity

It has been argued by Paul Faulkner in a recent paper that while the reason that is given would need to be one to trust the speaker, such a second personal reason only exists on condition that the hearer has the affective disposition that he trusts the speaker (Faulkner, 2014, pp. 338-340). This means that the second personal reason is not the foundational type of normativity involved. Rather there are norms that govern a practice and those who engage in it expect others who engage in the practice to abide to the rules of that practice – while this can mean there are second personal reasons when two people who partake in the practice are interacting with one another as part of that practice.

So, I argue that a participant stance is possible even if the normativity in play is not one of a second-personal nature. It is enough that we view testimony as a practice. By understanding testimony as a practice (which extends beyond a one-off encounter as modelled on a court room case) we can come to see that there is in fact not only

dependence of one party on another, as the second personal account coupled to an individualistic conception of knowledge suggests, but interdependence. Another way to put this point is to say that the practice is dynamic: all those who partake in the practice are responsible for the epistemic state that they and the other parties involved end up with.

I concede that in Faulkner's own account too, trust is needed. It is needed, so he claims in order to explain the rationality of epistemic dependence. For he claims that norms of trust guide these exchanges. That these norms exist is something he proves by pointing to Grice's "Cooperative Principle", and its norms for effective interpretation (Faulkner, 2011, p. 187; see Grice, 1989, pp. 26–27). What I would argue, however, is that only a social conception of knowledge can both explain the existence of these norms and their moral-epistemic nature: these cases involve not mere dependence, but interdependence. The fact of interdependence is, so I argue, what explains why we expect others to correct us if we mistakenly believe something, and why we should tell them if they are mistaken, too. Interdependence is the antonym of (epistemic) individualism, as explained by Longino:

> A third form (of non-individualism) is neither eliminationist nor wholist, but socialist. It stresses, that is, the interdependence of cognitive agents and subjects. The difference between (as much as the commonalities among) such agents are what enables them together to generate and justify content. Their interrelations license the attribution to any one of them. The characterization of knowing subjects as interdependent yields individual subjects without individualism (the view that individuals know independently of their relations with others). (Longino, 2002, p. 91).

6. Faulkner's Solution

Faulkner explains that he regards the requirement of spelling out "the demand that acceptance (of testimony) be backed by reasons" as the requirement to explain that the acceptance of testimony is "(individually or subjectively) reasonable" (Faulkner, 2011, p. 10). This is to take the reductive challenge to testimony as being "internalist" in its nature (Faulkner, 2011, p. 10). He explains:

> (…) I will understand the reductive theory to be internalist. The reason for this is that I think the best motivation for the requirement that an audience have some capacity to discriminate true from false testimony comes from the argument from cooperation. And this argument institutes the demand that acceptance be backed by reasons. (Faulkner, 2011, p. 10)

In taking this stance, Faulkner is emphasizing that his theory seeks to combine the non-reductive claim that testimony is a self-standing source of knowledge and the "reductive" demand that acceptance be backed by reasons (Faulkner, 2011, pp. 10–11).

It may seem that doing so is inconsistent – if testimony is a self-standing source of justification, justifying the recipient's belief that p, then it is inconsistent to demand that supporting reasons are necessary for coming to believe that p (see Faulkner, 2011, p. 11). However, there is no need for an inconsistency if either the object of the supporting reasons is something else than the belief that p, or if the nature of these supporting "reasons" is going to be what non-reductionists typically assume, too, or assume as a background condition – namely that the speaker is trustworthy and/or that the hearer's trust in the speaker is his basis of accepting what she says. We will see that Faulkner does both, and that is why the inconsistency is avoided, and why his account looks attractive.

What I want to focus on here is that Faulkner appears to assume that all non-reductionists have been either externalists or defenders of a default entitlement, and often both.[2] This is what allows non-reductionists to maintain that there is no requirement that would require an inference which would necessitate supporting beliefs and to deny that an assessment of the speaker or of what is said is necessary for reasonable uptake. The problem, however, is that these positions appear to leave the hearer off the hook entirely in any individual case where he acquires knowledge though testimony (Williams, 2002, 85; Moran, 2006, p. 275; Faulkner, 2011, p. 11).

Faulkner references Coady (1994, Ch. 9), Burge (1993, p. 467) and McDowell (1994, p. 438) as opting for an externalist and generalized solution: of arguing for a default entitlement which ensures the hearer does not need to be able to give an argument for why he accepts a piece of testimony in a given case (Faulkner, 2011, Sections 4.2–4.4, and p. 122, p. 130). If there is no default entitlement, then the hearer requires a justification or entitlement in each case. That entitlement could be construed in an internalist way, in the sense that it is an individually or subjectively available reason, or in an externalist way, where it is a mere entitlement grounded on a reliable, truth conducive, method. In the latter case, one could postulate that the hearer needs to have the right, reliably trained, sensitivity. In the case of a sensitivity, it may be argued that the hearer is not epistemically off the hook, because their capacities matter for whether or not they have acquired knowledge (E. Fricker, 1994; Hinchman, 2005, p. 178-179; Faulkner, 2011, p. 39). Yet, Faulkner says the reductivist wants a solution to what he calls "the problem of cooperation" (Faulkner, 2011, p. 7).

This problem of cooperation, in its testimonial variant, is the challenge of justifying the fact that a person accepting a piece of testimony provided by another person is reasonable or rational in this acceptance (Faulkner, 2011, p. 6). What it motivates is the requirement that "(t)he acceptance of testimony must be backed by reasons if it is to

---

2 While it would seem Faulkner includes the assurance view – as defended by Moran, 2006 – among the non-reductive positions (Faulkner, 2011, p. 9), Faulkner clarifies in chapter 6 of his book that he regards the assurance theory as defended by Hinchman as a trust theory (p. 167) – even if he also thinks that the assurance theory fails in its scope, because it is limited to *tellings* (Faulkner p. 168) At the same time he believes the assurance theory can be part of an 'adequate epistemological theory of testimony', as he believes it can 'provide a solution to the problem of cooperation.' (Faulkner, 2011, p. 109).

be rational" (Faulkner, 2011, p. 6). That a reason for trusting is required is motivated by an argument intended to show that trusting others for the truth is problematic, i.e. that there is reason to doubt that others are trustworthy. The challenge arises since (and in so far as) the audience and the speaker have divergent interests: if it is in the audience's interest to only accept the testimony as a basis for belief when the speaker is a trustworthy source, while it is in the speaker's interest to be trusted even if she is not a trustworthy source (Faulkner, 2011, p. 6).

While there could be externalist answers to the problem of cooperation, its impetus is internalist, as Faulkner makes clear (Faulkner, 2011, p. 10). The appeal of the cooperation problem – raised by the possibility that others can deceive you – is that it requires that the hearer actively discriminates between a trustworthy and non-trustworthy source, i.e. with a cognitive awareness of doing so. On an internalist view, someone requires a justification for accepting a belief on the basis of testimony in the sense of their being able to cite an argument for their acceptance (Faulkner, 2011, p. 10). On the internalist reductive view, a proper answer to the cooperation problem should be given, one whereby the hearer should ascertain that their trust is based on the speaker deserving it. Accordingly, the way things are with the speaker, epistemically, is something of which the hearer has to have a grasp. There are, on this view, no basic testimonial reasons: a further reason must be cited to ground the reason to believe the person who testifies: "testimonial justification and warrant are ultimately inductive in nature" (Faulkner, 2011, p. 10).

Faulkner writes that in his view "acceptance (must) be (individually or subjectively) reasonable" (Faulkner, 2011, p. 10). This means that Faulkner opposes theories that posit a specific kind of entitlement, namely ones that leave the hearer's cognitive efforts wholly off the hook. In denying a general entitlement, Faulkner's position is reductive (Faulkner, 2011, p. 11). It adheres to the requirement that the hearer should exercise reasonable trust in each case:

> If acceptance issues in belief – if we take a bit of testimony to inform us of some fact – then our reason for accepting that bit of testimony must be a reason for belief. (Faulkner, 2011, p. 3)

At the same time, Faulkner clarifies that he finds the reductive position overly restrictive in its account of *what qualifies as a reason* for accepting testimony (Faulkner, 2011, p. 11). He maintains that an attitude of trust, if and when it exists, can qualify as a reason for accepting testimony, too (Faulkner, 2011, p. 12). This means Faulkner seeks to defend a hybrid position, which is partly non-reductive, and partly reductive. Reductive because reasons for accepting testimony are required. Non-reductive in that testimony is a source of (second hand) knowledge, because the attitude of trusting the speaker can constitute a reason for acceptance –as opposed to one's needing to having reasons for accepting that the content is true (Faulkner, 2011, pp. 11-12).

In a similar vein, Hinchman has defended a hybrid view whereby he seeks to combine internalism and externalism about justification, in order to explain how knowledge can be gained on trust, while there is a requirement on the hearer, too, to exercise reasonable trust (Hinchman, 2005, p. 579). Hinchman stresses that this requires the cognitive self-government of a person (Hinchman, 2005, p. 584). He thinks a strong externalism in the form of a pure reliabilism can be rejected, since the hearer's role of trusting the speaker can be conceived as a form of cognitive self-management, even if the hearer is not required to assess the speaker or the content:

> Trust is not a 'mechanism' of belief formation; it does not merely cause beliefs but manifests the subject's cognitive self-governance'. (Hinchman, 2005, p. 584)

I share the basic intuitions by Faulkner and Hinchman: that testimony is to be a source of knowledge whereby we can gain knowledge on trust, and that reasonable uptake is required on the part of someone who receives knowledge. Yet, I think that the hybrid accounts proposed by them try and combine reductive and non-reductive elements in a way that is not fully satisfactory. Their accounts are pluralist, which can appear ad hoc, because the results yielded in terms of why and how a transmission is successful are explanatory rather than justificatory and this internalism that drives the cooperation problem requires a justification.

Before discussing Faulkner's hybrid trust-based theory of testimony in the next section, I will first present my argument that there is room for an alternative. I maintain that the requirement of reasonable uptake can be accommodated in a non-reductive account of testimony – or rather, in a generative account of testimonial practice – when the individualistic understanding of knowledge is discarded, so as to enable us to see how shared epistemic responsibility for epistemic cooperation is possible in testimonial encounters.

My argument is that the two main current accounts of testimony – evidentialism and assurantism – each accord full responsibility to only one of the parties to a testimonial exchange, and that both accounts are found wanting as a result. I argue that because both accounts are attractive and yet also incomplete, we need to understand from a meta-philosophical perspective on the nature of the dispute between them. Both sides in this deadlocked dispute are excessively individualistic in their focus, particularly when it comes to epistemic responsibility.

My MacIntyre-inspired analysis showed that evidentialists strive for truth in testimonial practice at the expense of social relations and trust, while assurantists do the opposite. The key for successful epistemic cooperation, however, is that trusting others should yield truths, while truthfulness should inspire and yield trust. So, while trust should not be granted cheaply, it should also not be refrained from as a matter of principle (Audi, 2013, pp. 529–530). Hence, in terms of what must be presupposed by the

idea of epistemic cooperation, there is an undue imbalance in both views.

My interpretation of this result is that the solution is that both these views are too individualistic in that they hold that knowledge always ultimately originates with a single knower. Both accounts end up giving the full responsibility for knowledge acquisition to only one party of a testimonial exchange, leaving the other party entirely off the hook. The solution, I suggest, is to take epistemic cooperation much more seriously, so that a proper dynamic of truth and trust can be accommodated for. This means that we must start from a more strongly social conception of (testimonial) knowledge than is possible from within an individualist approach. I will begin, however, with a necessary excursus on exactly what I mean when I claim that knowledge is "strongly social".

I argue that what a social conception of knowledge makes available is an understanding of the social aspects of testimony that have tended to be ignored. Recently, some advances have been made, in the understanding of the social aspects of testimony, by those who emphasize that in testimonial encounters there is an important role for trust. Assurantists – just as Welbourne did – have emphasized that trust mediates knowledge. On their view of testimony, a speaker extends to the hearer an invitation to trust. Some theorists have argued, in addition, that such an invitation, or assurance, provides a reason for the hearer to depend on the speaker. That is, they seek to apply the idea of second-personal normativity to testimonial encounters (Moran, 2013, p. 119; McMyler, 2011). Indeed, while Welbourne also emphasized the role of trusting the speaker, so Kusch has pointed out:

> I wonder whether Welbourne's 'believing someone' is not best captured with concepts like 'trust' or 'participatory stance'. Trust in a person's word is more than simple reliance; trust amongst people has a moral quality. (Kusch, 2002, p. 62)

Welbourne's view, too, is appealing because it shows there is a role for understanding how speaker and hearer are related to one another which makes them participants in a practice, and which helps us to explain how we may understand the normativity involved as that of interdependence. We can view the speaker and hearer, theoretically, as morally dependent on each other each taking "a participant stance" towards the other. I would argue that there is a similar presupposition in Craig's genealogy given that he maintained that a special characteristic of humans as sources is that we are working to achieve knowledge as cooperative members of the species.

However, while the idea of a participant stance is indeed relevant, in so far as testimony is to be viewed as a practice, a question that may be raised is whether second personal normativity is the best route to a solution here; I agree with Faulkner that it is not, but my proposal is, as I will argue in the next chapter that testimony is to be understood as a generative source of knowledge. Indeed, the role of testimony in generating knowledge serves to substantiate the claim that knowledge is social in a strong sense.

7. Assessing Faulkner's Hybrid Account

Faulkner's own account of testimony, which does focus on explaining the puzzle of testimony in its own right, is not fully internalist. Even while he does seek to explain the rationality of a hearer's uptake, he does not provide justificatory reasons for it (Faulkner, 2011, p. 15). Faulkner's claim is that while we choose affectively to trust, and hence we chose to rely on others, there is a self-fulfilling prophesy at work in our doing so (Faulkner, 2011, pp. 149-150, p. 182). For the fact that we morally trust another person to speak the truth makes them more likely epistemically reliable, as they are morally expected to live up to our expectation to be truthful and know that we depend on them to meet that expectation. Faulkner, 2011, p. 151). To not be truthful, then, comes to a betrayal of our trust. As such, Faulkner can "rationally explain" why testimony tends to be reliable not only in general, but in the specific settings whereby we actually choose to depend on others. Moreover, the claim is that our reasons can be deemed to be epistemic, because it is "objectively more likely" that a person who is being depended upon for the truth will speak it, provided that the situation is structured by the norms of trust (Faulkner, 2011, p. 159).

There is a caveat here, however, which is that we should exercise reasonable trust, i.e. trust wisely, even if the reasons can be moral as opposed to epistemic in their nature. To be reasonable means that epistemic reasons for distrust cannot be wholly discounted: a person cannot trust come what may (Faulkner, 2011, p. 152). Faulkner maintains that, in a number of cases, we can take an evidentialist stance to testimony and in some cases we have to (because no assurance is given, as not all acts of testimony are acts of telling). This also shows, so he contends, that a pure assurance theory is limited in scope, and would fail as a theory of testimony for that reason (Faulkner, 2011, p. 167).

Instead, he proposes, as we saw, a hybrid trust theory of testimony, whereby trust is basic in the sense that it can explain why testimony transmits knowledge in both the cases where we are given assurance, and in cases where assurance is lacking. In cases of "mere assertions" we can still trust the speaker if we are able have an attitude of affective trust towards the speaker (Faulkner, 2011, p. 177). That we can have an attitude of affective trust is because we have internalized the norms of trust (Faulkner, 2011, p. 171). When that attitude is not available to us, because in the given case trust is not reasonable, the circumstances not being those of trust, we can fall back on the evidential stance, and take what a speaker says at face value (Faulkner, 2011, p. 188).

Faulkner claims that we can take an evidential stance because we have come to see testimony as generally reliable; generally, people are motivated to speak the truth, because the norms of trust are in place (Faulkner, 2011, p. 190). I must admit that I do not see how this works, for we would be taking the evidential stance in cases where we doubt – or at least, lack a propensity to trust - the very source in question. Perhaps the

idea is that we nonetheless are trained to find testimony on the whole a useful means of gaining knowledge, and therefore still take it to be evidence.

My main objection to the account given concerns the fact that the account given by Faulkner cannot account for the norms that are invoked, themselves. Another objection concerns the hybrid nature of his account. What Faulkner in fact provides, is a reason for supposing testimony is reliable if and when we choose to depend on it, if we do so wisely.

Faulkner is correct to say, so I believe, that his account is largely externalist, in so far as he provides a rational story about the reliability of testimony. That story is not available to everyone who exercises reasonable trust: someone who chooses to trust affectively does not do so based on epistemic reasons, but does so out of trust, even if that trust is informed and shaped by experience or prudence or common sense. Moreover, there is not really a way of telling whether one's exercising reasonable trust is wise except in hindsight. That is, whether or not it will pay off epistemically cannot really be said in any individual case unless one uses epistemic reasons for trusting, after all.

Faulkner argues that on his account there is room for uptake being backed by epistemic reasons, since the attitude of trust can be shown to make it objectively more probable that the trustee will speak the truth, if the norms of trust apply (Faulkner, 2011, p. 156). Yet, in other circumstances where trust is inappropriate, it will be necessary to adopt an "evidentialist stance", in a case of a known pathological liar for instance. The problem with Faulkner's account, then, is starting to become clearer: which are the cases where we should trust and which are the cases where we should adopt the evidentialist stance? I contend that a principled basis is required for the pluralist account to be satisfactory – as opposed to an ad hoc, case by case, justification.

Faulkner does not offer such a principled basis. He maintains that there can be no deeper level of analysis; i.e. that the attitude of trust is what we should go on to determine when trusting is wise (Faulkner, 2011, p. 160). But this would imply that his account is no more satisfactory than that of assurantists. For, if it really is a self-fulfilling prophesy, then this makes it sound as if trusting is always wise. It does not give us a means of discriminating those cases when we should refrain from trusting.

From my perspective, Faulkner's analysis, in being externalist and pluralistic, falls one step short of being a fully satisfactory solution. The idea of explaining testimony as a source of (second hand) knowledge requires that we can understand how trusting wisely makes a difference. That we trust wisely is to be read off not only from the epistemic consequences of our affective trust, and willingness to extend it, but is to be a function of our background knowledge and assessment of the credibility of the source. Moral or affective trust – even if it may play a role in determining to whom we give credibility, or in how much risk we should be willing to incur to be lied to – does not suffice for characterizing what constitutes epistemically reasonable uptake. At least, it does not do so from the perspective of the internalist whom Faulkner wants to accommodate, as he admits.

What of the externalist who wants to know why trust is rational in individual cases? There is some reassurance to the thought that if one trusts, then one tends to trust wisely, if only because it motivates others to act truthfully. However, there is also an element of wishful thinking here which is being informed by this line of reasoning. That type of thinking defies the vigilance called for by Williams – to ascertain that the conditions of trust obtain – or else appears to require us to always take an evidentialist stance by default.

So far, my explanation of the puzzle of testimony has considered those authors who continue to work from within an individualistic conception of knowledge. What I will argue next is that the relation between truthfulness and trustworthiness needs to be reconsidered. Faulkner is right to insist that there are norms of trust in play. However, he does not offer the resources to explain why these norms exist, nor why they have the dual moral-epistemic nature that accrues to them, and which helps us explain reasonable trust in a much more helpful way. So, I argue that the separation of affective trust and epistemic trust is not a "divorce" given that reasonable trust should remain epistemic to some extent. But that fact follows from the insight that the moral and epistemic ought to be conjoined for epistemic cooperation to work, in order for knowledge to be generated through testimony.

Conclusion

In this chapter I have discussed how the problem of testimony has been analyzed and how theorists have sought to solve it. On my analysis, the problem arises because the individualistic conception of knowledge, according to which knowledge is the proprietary and unique responsibility of an individual, does not allow us to explain why epistemic responsibility for knowledge can be shared. Nor, relatedly, does it explain why there is an apparent conflict in the attitudes that we must display when receiving putative knowledge: why there should be both a default to trust and yet also an impetus on us to verify that what is asserted is true.

In the next chapter I will explain what testimony is, consistent with the (dis)solution of this puzzle, and whether or not it involves moral-epistemic norms. There I aim to show both that testimony is a knowledge-generative practice, and that this helps us understand the nature of knowledge better, while making good on the claim that testimony is foundational for knowledge.

Chapter V | Testimony as a Source of Knowledge

Introduction

In my view taking Welbourne's view concerning the social nature of knowledge seriously would have a transformative effect on how philosophers ought to think about testimony. This is a rather radical claim, which not even Welbourne himself made – even if it appears to be a corollary of his view that testimony is central to knowledge and that it is required for knowledge that it passes as knowledge in a community.[1] (Welbourne, 1985; Kusch, 2002, p. 60-62). For the latter claim makes sense only if passing as knowledge is a social status that accrues to a claim through social validation, which I believe is precisely why testimony is to be conceived as an epistemic practice, and why testimony so-conceived is central to the generation of knowledge.

So, I will argue in what follows – a point that has also been noted by Kusch (2002, p. 62) – that the natural conclusion of Welbourne's insistence on testimony as central to knowledge and on knowledge being commonable while belief is not, is to regard testimony as generative of knowledge, as opposed to regarding it as merely a vehicle for its transmission.[2]

2. What is a Generative Account of Testimony?

A generative account of testimony is an account according to which testimony is a practice that plays a central role in generating knowledge and not merely its transmission. It

---

1 Instead, Welbourne is leaving his work at the less committal level of drawing conceptual between knowledge and testimony, which he claims serve to illuminate both. For instance, in his later work, he writes: "My idea, then, is this: the concept of knowledge enters our repertoire of concepts on the back of testimony, and it is a concept of something that owes its value ultimately to the value of testimony in our lives as a propagator of (mostly) true beliefs." (Welbourne, 2001/2014, p. 125) While this statement could have a reading that shows how testimony is generative of knowledge, it remains perfectly compatible with the view that testimony is merely a way of transmitting knowledge which has already been created.

2 Kusch writes:
"Finally, pushing Welbourne's analysis to its natural conclusion, and thus further than he seems willing to go, we can formulate the resulting position as follows: 'knowledge' is a social status like money, and thus it only exists in so far as there are items upon which we are willing to impose the status. And the collective imposition of the status happens in and through local transmissions, first of knowledge-claims, and later of accredited knowledge-items." (Kusch, 2002, p. 62).

distinguishes itself from the traditional view that testimony is not a basic source of knowledge (Audi, 2013, p. 150; E. Fricker, 2006). This traditional view, according to which testimony is not involved in the generation of knowledge, I have called a "transmissive" account of testimony.

The transmissive account and the generative account are exclusive in the sense that the transmissive account denies the central claim of the generative account: the claim that testimony is a source of new knowledge, i.e. properly seen as a source of first-hand knowledge, rather than merely a vehicle for knowledge transmission. At the same time, the generative account of testimony does not deny that testimony can transmit existent knowledge as well as generate new knowledge. The central point of contention between the two types of account concerns the claim of the transmissive account that testimony can transmit only already existent knowledge, whereas the generative account holds that testimony can also produce new knowledge.

Even so, the difference between these two accounts would prove elusive, if it were the case that testimony hardly ever generates new knowledge – or does so only accidentally. This is relevant, as some theorists accept that knowledge could be gained from someone who unknowingly speaks truthfully (Hinchman, 2005, p. 586). Assurantists may hold that in such a case, the hearer can gain knowledge on the basis of trust (even if it is misplaced as the speaker intended to deceive him), and reductivists may hold that the hearer is able to gain knowledge on the basis of the speaker's testimony if he ascertains its truth, despite her own misgivings. So, for the difference between the transmissive and generative accounts of testimony to be relevant, the contrast between them is to mark a difference in our understanding of the principal role and characteristic nature of testimony.

I claim that we can only understand certain features of testimonial practice if we see it as generative of new knowledge. It is because generating knowledge and defeating false claims to knowledge is central to testimony that it can be characterized as an epistemic practice, i.e. as a practice through which knowledge is gained and lost (Fricker, 2007, p. vii). This understanding of testimony as an epistemic practice accommodates the epistemic roles played and responsibilities held by the parties who partake in testimonial exchanges.

In order to yield new knowledge, testimony has to ensure that what passes as knowledge does so for the right reasons: that the social criteria which have been put in place to increase the likelihood of truth are applied correctly when statements are socially accepted as knowledge. While social acceptance, if it accords to the proper socially set criteria, should increase the likelihood of truth, because the criteria are such as to secure this, there is no guarantee that what passes as knowledge is true.

Knowledge claims remain defeasible in the testimonial case as in any other; the practice as a whole is fallibilist. If testimony is foundational for knowledge, then social acceptance in the sense of "passing as knowledge" is a necessary condition for a claim to

be knowledge, but not a sufficient one. The reason is that this is a necessary requirement is that in that case knowledge is essentially a social status: a marker of trustworthiness or reliability. At the same time, the reason why passing as knowledge, hence being assigned the status, is not sufficient for knowledge (even if there is a social consensus) is that knowledge presupposes truth. Knowledge is factive. It could be that all members of society are mistaken in applying the criteria in the particular case or indeed, if we would wish to go as far as to accept the option of global skepticism, the criteria may themselves be wrong.

At any rate, the status of knowledge – even if it is appropriately given to a claim – may subsequently be defeated. The defeasibility of particular claims is built into the social conception. It could be that we learn new facts which forfeit the justification that we had obtained for counting some claim as sufficiently supported. Social acceptance embeds an idea of increased likelihood of truth, because it assumes that individuals have brought to bear epistemic reasons, i.e. that their acceptance has been epistemically responsible. So, a proposition's social acceptance as knowledge – reflected in its passing as knowledge, its social status of knowledge – is the collectivized version of an individual's justification.[3] Accordingly, it is a mistake to think that the social acceptance is all there is to something being knowledge. Rather, what we are stating is the truism that something passes as knowledge if it is socially accepted as knowledge. In addition, we are making a claim that it is necessary for something to pass as knowledge, in order for it to be knowledge. (The aim, in neither case, is a reductive definition of the concept of knowledge.)

There is indeed a worry that the view that knowledge is a social status may be unable to differentiate between proper and improper ascriptions of that status. That worry derives from the fact that the very idea of "status" reflects that there is an inherent fact of authoritativeness. The social conception of knowledge may appear troubling precisely because we are so used to thinking that, given a choice of what is to be preferred, epistemic authority should reside with the individual and not with a collective. However, the question of what we would prefer is not relevant. The question is: what conception of knowledge best explains the nature of our social epistemic practices, such as testimony, given that testimony is so central to it? That question is central because the focus should lie with those practices that can explain why we have the concept of knowledge in the first place (Welbourne, 1986, p. 75).

In my opinion, the best elucidation of the concept of knowledge is one which – as Welbourne has put it – explains why moral and political questions arise in the realm of

---

3 But note that this does not mean I condone a view of social epistemology that makes individual contributions key to the achievement of knowledge: the process of knowledge generation is interactive, and in that it is irreducibly a collective process (Longino, 2002, p. 122). This is also the view by Longino contrasts with the views by Goldman and Kitcher which remain epistemologically individualist (Longing, 2002, p.42). Even so, there is room for individual autonomy: individuals can make a difference, even as they are doing so as members of an epistemic community. This is possible because they can enjoy a degree of autonomy sufficient for questioning others, while they may also be part of more than one epistemic community (see Longino, 2002, p. 154-155).

testimony. These questions, so I argue, principally concern the proper use and allocation of epistemic authority. These questions can be explained by the fact that knowledge is a status that is attached to statements as the outcome of social processes. For it is because we can see that the outcomes of such processes may be tainted by undue social influences, while these outcomes are nonetheless presented as authoritative, that we understand the urgency of proper epistemic conduct. This is the urgency of attending to moral and political aspects of epistemological practice. So, there is an inherent connection between the account of testimony I develop here and the account of epistemic justice I present in the next chapter.

Interestingly, those who have an individualistic conception of knowledge have no role for testimony as involved in the generation of knowledge. The individuals, at either end of the exchange, are each to "do their bit". There is no sense that these individuals cooperate to gain an improved epistemic position: each furthering knowledge. The success of testimony on the transmissive model is seen as purely communicative success. On this view, testimony succeeds if the speaker has knowledge and if the hearer gains that knowledge. The testimonial exchange fails if the hearer does not gain knowledge from the speaker. If the speaker does not have knowledge to start with, then, in principle, the hearer cannot gain knowledge from the speaker.

"In principle", because the hearer can gain knowledge by a kind of fluke, if the speaker lacks knowledge but happens to have spoken truthfully and the hearer happens to have believed her while, in addition, the hearer is considered justified in doing so (Hinchman, 2005, p. 586) – here we are assuming the justified true belief account, as it is that account of knowledge which is typically appealed to in the transmissive model. So, on the transmissive model of testimony, it appears that testimony fails if the hearer does not end up with second-hand knowledge. Yet, it is worth considering the intuition, which there appears to be no room for on the transmissive model of testimony, according to which in some cases the best outcome of a testimonial "exchange" is that the hearer improves the epistemic position of the speaker, fairly and rightly contesting that what she presents to him is knowledge. Similarly, we can think of a case where a speaker who is not confident that what she says is true may end up with knowledge by way of sharing her thoughts with a hearer who validates them as true.[4]

I argue that this illustrates the point that the generative account can give us a conception of testimony according to which testimonial success is much more sensitive to its dynamic and cooperative nature. The interaction between parties and the cooperation which is assumed by each of them providing feedback to the other is wholly absent in a purely transmissive model, where individuals remain related atomistically:

---

4 In the first case I just gave, of contestation, knowledge may be lost (in which case the speaker's epistemic position may be deemed worse than at the start). In the second kind of case, whereby a person affirms the views of another, knowledge has been gained. Indeed, in this last case, by combining their separate insights, knowledge may have been generated by these two people, depending on whether or not it pre-existed and whether or not it its true.

knowledge is had by one of them, and taken by the other, or not. Whether or not they engage in a cooperative interaction and improve one another's epistemic situation, bearing responsibility for each other's epistemic position and that of other members of their community (who may learn or be epistemically influenced indirectly through the outcome of their deliberation later), is not a function of "testimonial success".

Ultimately, we want to explain why there are norms of truthfulness and trustworthiness that play an important role in testimony, and why there is to be room for both epistemic autonomy and epistemic authority when it comes to knowledge sharing and knowledge acquisition. My proposed answers are that testimony is central to the generation of knowledge, which follows from the strongly social nature of knowledge as "essentially commonable". This understanding fits with Craig's insight that knowledge involves a social status. The business of knowledge is social because we all flag the trustworthy sources together, conferring the status upon those who should be recognized as epistemic authorities, and filtering truths from falsehoods. This understanding of knowledge as a socially sustained and accorded status and the role of testimony as central to it requires that we view testimony as the dynamic, generative practice it is, and that we redefine what a successful testimonial encounter comes to. Once this is clearer, we can see why there are norms that govern our testimonial practice – which norms are in the background of both evidentialist and assurantist accounts – and the nature of these norms.

While I thus argue, mainly, that we should move away from a conception whereby testimony is never generative, I will also argue that the generative conception allows that knowledge can be lost, as well as gained, through social interaction. Accordingly, and in a way opposed to the traditional conception of knowledge, testimonial success is not, or at least not only, the successful communication of knowledge from one person to another. Instead, testimonial success is when a community achieves an improved epistemic position of its members. This can be by way of increasing the number of things that are known, or by way of decreasing the number of false beliefs. So, one crucial part of testimonial practice, and of its being successful in the generation and sustaining of knowledge, is that members of a community proactively inform and correct one another.

This allows us to move away from a conception of knowledge and testimony where there is no need to look after each other's epistemic interest, because each person is to cater solely for their own interest. On the generative conception of testimony, whereby it is key to the strong sociality of knowledge, the epistemic interests each of us have are shared and our ways of securing them are aligned. The reason is that knowledge does not belong to individuals, as it would be if it is their personal property which arises on the basis of the individual having his or her own proprietary warrant for it. Knowledge is a status which is accorded on the basis of interactive judgement using shared criteria (Longino, 2002, p. 148). It comes about only after social processes of validation (Longino, 2002, p. 122).

Of course, individuals can have knowledge, in the sense that they fulfil whatever criteria are to be met for being accorded epistemic authority. They may claim to have knowledge, and it may appear to be a personal achievement that they have it. But epistemic authority is a socially accorded status: the criteria which a person must meet are social criteria. The reason an individual cannot generate knowledge on his or her own is that some putative item of knowledge does not qualify as knowledge proper if it lacks the social status. The social status is the requirement that it "pass as knowledge" in the community. This will mean, at least, that the item is not challenged by those who have might or epistemic authority. It is very common in current philosophy – despite the rise of social epistemology – to think that individuals are the proper unit of epistemology. This view, however, cannot be correct for two reasons.

The first is that one's responsibility, even while it should not be denied that an individual has it, is already other-regarding from the outset. The reason is that knowledge is a social status. This is not superficially so but deeply so: I can claim to know something, but I do not qualify as knowing it unless I can justify it to others – as Mill claimed (Mill, 1989, pp. 22–23).[5] (While I also noted, earlier, that Mill seems to have intended to explain the nature of knowledge via his account of inter-subjective rational warrant in a reductive way – whereas I do not share that aim.) The fact that I am accountable when I claim that something is knowledge, given that others may rely on my judgment, also means that I am responsible. But that responsibility, which I have, is not individualistic in the sense of a concern that is just for my good. Instead, I am responsible to others, because I take part in the social practice of testimony when I claim something is knowledge, a social practice which gives rise to new knowledge.

Secondly, the responsibility which I have is not (typically) one which renders me ultimately responsible: the others involved in the practice of testimony each share the responsibility for whatever knowledge arises, too. Given that knowledge is a social status, the responsibility for accepting that something is knowledge is shared among the people in the community. No one individual is ultimately responsible for a specific item's status – unless we have granted that person the privilege of being an esoteric, who has privileged access to knowledge that is not available through any other route. But even then, the rest of us are there and have (tacitly) accorded that authority to this person enabling them to uncover knowledge for the rest of us – for good or for bad. Even as I may originate an idea, or a claim to knowledge, it can only become knowledge once it is accepted by others as knowledge.

Note that I am not claiming that individuals have only a small amount of responsibility or that they have limited responsibility or that they should rely on others more. It could be that all of us have very great responsibility even if we share it. For I have not said that

---

5 Longino holds a similar view, but makes the important nuance that what qualifies as a good justification depends on the context. The standards for being a knower need not be high (Longino, 2002, p. 151).

we may delegate or think less of what responsibilities we take ourselves to have. What I have said though is that the idea that knowledge belongs to an individual as a self-made property, whereby the individual bears ultimate responsibility for its being knowledge is mistaken. Of course, an individual should care that what she believes is true, and that her justification is sound, and that she does not reject knowledge or accept falsehood as knowledge. But those responsibilities, so I argue, are compatible with a social conception of knowledge. Indeed, they are even more important because of the social conception of knowledge.

Of course, some beliefs may be false but helpful, for instance because they are useful approximations of the truth. The guiding thought and main concern, however, is to signal and eradicate those falsehoods that can come to pass as knowledge, or even already pass as knowledge, that are unhelpful because and in so far as they stand in the way of uncovering truths that are socially relevant. So, I argue that the role of testimony is not only that of sharing that which we regard as true for good reasons but also that of filtering out the falsehoods that may have been shared and accepted (at least by a portion of the community) as if they are true while in fact there are no good grounds or even serious indications that they are false.

3. What Speaks in Favor of a Generative Account of Testimony?

A generative account of testimony can account for testimony as a practice, which is not merely two-personal but involves all members of a community. It can make sense of why feedback from different people is valued and why it is expected that people come forward with the knowledge they have: because doing so is important for generating knowledge. It is only through others' affirmation that knowledge claims can be socially accepted and (when widely accepted) can come to "pass as knowledge", which is required for them to be knowledge at least as a necessary condition.

A generative account can account for epistemic responsibility being had by all the parties who partake in a particular testimonial exchange or, more widely, in the testimonial practice. Since knowledge is socially generated, no single individual is ultimately responsible for the result of a knowledge exchange. Instead, we have a dynamic account of testimony whereby all parties are responsible for the epistemic position each one ends up with. This does not negate one's individual responsibility. Rather, it means that our responsibility is also other-regarding.

A generative account of testimony is key to a social conception of knowledge such as Craig's. At the heart of this conception is the idea that we serve as resources for one another. This means that we rely on each other's testimony. Notably, however, we do not simply rely on others: we work together to single out those sources which are most likely

reliable. This is why we have shared standards for a "good informant", which is the proto-concept of our concept of knowledge in Craig's genealogy. Identifying good informants is a joint, cooperative enterprise which is key to our being able to rely on others for the truth. The genealogy which Craig provides shows, essentially, that knowledge is a social status, backed by the joint endeavors of filtering the good from the bad sources. It may be conjectured that in cases where knowledge becomes detached from the originating human source, the mechanism that keeps the epistemic title of being "knowledge" a good one is the social practice of testimony. In testimony, knowledge claims continue to be put to the test, by the receiver. What is passed on by the speaker and accepted by the hearer literally "passes as knowledge". So testimonial practice helps to detach knowledge from its source, to "objectify" knowledge (Faulkner, 2011, pp. 192–193).[6]

But while this happens, the social processes of validation remain in place. Each individual can contest that what passes as knowledge is knowledge, and as a result, knowledge claims can be defeated. Responsibility is shared because the speaker has responsibility in passing on something as knowledge and the hearer has responsibility in accepting it. Both the speaker and the hearer then have responsibility, each, for the epistemic status of the beliefs they, and putative others, end up with.

Their acts of communicating putative knowledge and of accepting it as knowledge are two instances of other-regarding actions. They are other-regarding actions because there is a social consequence to accepting or rejecting some statement as knowledge, given that others can trust one as a source. Of course, these others too have responsibility, but as the claim gets detached from the source and gains more and more acceptance, the epistemic title to the claim will readily accompany it. Once a claim clearly passes as knowledge, unquestioned acceptance is going to prevail and will be legitimate. Prefixes of the form: "we know that", "it is widely accepted that", "as all of us have come to see by now", and weaker ones, such as "many people say that", "it is often said that", "it is hardly possible to dispute that", signal that we better accept what is said as it is not worth the effort to try and contradict it – unless, of course, we suspect that the speaker is a confidence trickster.

Since it may be hard to show that something which has passed as knowledge is in fact false, testimonial practice will work best if it can keep the source close to the information for as long as possible, and if the individuals who receive what they believe is new or unusual knowledge verify the source. It will also help is all of us remain somewhat skeptically inclined. That is, it may be best if knowledge claims don't come to pass as knowledge too easily – even if we do, psychologically, tend to accept what others tell us.

An upshot of my portrayal of Craig's genealogical account, and especially the idea that

---

6 I reference a point made by Faulkner, which again shows his work has some similar lines of thought as mine. The main point of difference is that Faulkner saw Williams's genealogy as aiming to secure a reductive theory of testimony, and that he is not focusing on knowledge, but on testimony in his own work, overlooking why there are epistemic norms – as opposed to merely social norms of trusting (e.g., Faulkner, 2011, pp. 199-200).

a test of the source is built into it and is what makes it social, is that Williams's genealogy of truthfulness is perhaps not strictly required by Craig's genealogy of knowledge. The reason is that Craig has an inbuilt test for knowledge, which is that the source must be considered reliable by the members of the community that assigns credibility/authority to it. The shared conception of a good informant, which takes as its standard that the informant should satisfy the demands of even the most demanding inquirer in order to qualify as a "good informant", would seem to rule out that the source is ill-intentioned or untrustworthy.

Of course, Williams intended to give an account of truthfulness in order to show that the notion of truth is one we care about deeply, and in linking his genealogy to Craig's account of knowledge, he also showed that proper moral-epistemic behavior is a precondition for a social conception of knowledge, whereby knowledge is socially generated through a reliance on testimony. Interesting, also, is that William's speaks of "pooling information" (Williams, 2002, p 57), while that term is not used by Craig: Craig's participants "share" information with one another, apparently in individual encounters. Faulkner notes that in a society more complex than Craig's described State of Nature, the source of putative knowledge, which is "pooled", may be anonymous (Faulkner, 2011, p. 191). This could explain why the problems of deceit and freeriding which Williams thought would stand in the way of successful epistemic reliance on others, are, after all, serious ones.

4. The Generative Account vs. Faulkner's Hybrid Trust Theory

In what follows, I will compare Faulkner's hybrid trust theory with the generative account. While I will argue next that the generative account fares better, it is important to note first of all that Faulkner's account remains within the transmissive paradigm. This being so explains why the norms of trust that Faulkner invokes in his explanation are seen by him as the norms that enable effective communication – as opposed to these norms serving to further the validation and dissemination of knowledge.

My first point is that the contrast between the generative account and the transmissive account is substantive because the roles they ascribe to testimony also determine the criteria for testimonial success. In the case of the former, the conception of testimony as a source that is generative of knowledge means that what "testimonial success" consists in is this: all those engaged in testimonial practice are to end up in an improved epistemic position. This means they either end up with more (actual) knowledge or with less false claims that passed as knowledge. Success on the transmissive model, in contrast, is given by the successful acquisition of knowledge by one party – knowledge that the first party already had. While bringing this about may require cooperation, as

Grice insisted, there is no clear co-dependence in the form of a shared end result, and, additionally, the success is clearly a communicative one – rather than an epistemic one (Grice, 1989, p. 30).

I will argue that this generative account offers a solution that is better than Faulkner's "trust theory of testimony". That theory tries to overcome the shortcomings of the traditional reductive and non-reductive accounts by combining aspects of each, being a self-described "hybrid" (Faulkner, 2011, p. 11 and Chapter 7, esp. p. 172; Faulkner, 2000). While I acknowledge that the starting point of both our arguments is Williams's position in *Truth and Truthfulness*, my view differs from Faulkner's in that I see testimony as generative of knowledge as opposed to only transmissive. In fact, I argue that Faulkner has misattributed the reductivist position about testimony to Williams, and that Williams's point is miscast when read that way (see Faulkner, 2011, p. 176).

A main reason why the generative account is preferable to Faulkner's hybrid theory is that it can provide a rationale that grounds normative appraisal of our epistemic conduct. This includes our reliance on others as sources. This rationale is that of co-dependence, and defies the idea that the cooperation problem can arise in the way that is suggested by reductionists and non-reductionists alike, who each provide only an incomplete solution to how and why someone can gain knowledge via or through testimony.

In comparison, while staying within the transmissive model, Faulkner's hybrid alternative only provides an ad hoc solution for our reliance on others. It points to our existent attitudes of trust or mistrust in a given case in order rationally to explain why we generally rightly rely on others if and when we do. In contrast, the generative account can point to the fact of epistemic cooperation as providing a rationale for – or against – the reasonableness of these very attitudes. The generative account makes this explanation available because epistemic cooperation is presupposed in testimony's role in knowledge generation and contestation. By "contestation" I mean cases when claims that are presented as knowledge are deemed to lack what being accorded the status of knowledge requires.

Williams argued that what is necessary for Craig's conception in our much larger community is that we, as a community, have mutually recognized internalized virtues of truthfulness. His genealogy of the virtues of truthfulness starts from the idea that they are required for the success of the practice of testimony, given that we may have practical reasons to be motivated by desires other than truth-telling. Faulkner, however, thinks that the virtues that Williams provides are not intrinsically motivating; specifically, that the virtue of Sincerity does not provide a motivation to act sincerely (Faulkner, 2011, section 7.2, esp. p. 179).

I believe that Faulkner does not correctly identify the rationale for Williams's position. The idea of the intrinsic valuing of these virtues is equated by Williams to these virtues being internalized by all members of a practice: i.e. that their being intrinsically valued actually means that people are disposed to act in accordance with them, such

that the practice has come to be structured by rightful expectations of truthfulness. In this way, it is a given that it is public knowledge that the virtues are part of the public ethos. My reason for this interpretation of Faulkner is that he insists, on the one hand, that the "intelligibility" of the virtues from the inside is not enough for them to have motivational force (Faulkner, 2011, p. 179). Yet on the other hand, he also states that a reductive theorist of testimony could have a solution to the problem of cooperation if the audience knows that the speaker prefers acting sincerely (Faulkner, 2011, p. 179). As for this last point, Faulkner insightfully remarks that precisely this knowledge of another's preference is lacking in the cases that give rise to skepticism about the possibility of trusting others so as to gain their knowledge through testimony, so that it is not the solution for which we should look (Faulkner, 2011, p. 179).

It seems to me that Faulkner's misreading of Williams is given by his wish to argue, as forcefully as possible, that predictive trust does not suffice. It is because he thinks Williams defended a reductive account of testimony that Faulkner insists that the reductive answer would involve "predictive" trust (Faulkner, 2011, p. 177). Yet according to Faulkner it is only affective trust and its account of trustworthiness as a thicker notion than mere reliability which can satisfactorily explain what grounds reasonable uptake (Faulkner, 2011, p. 177). He also claims that this affective trust is really what connects truthfulness (in the form of Sincerity) and trust in the way Williams envisaged. while phenomenologically "affective trust" – as opposed to "predictive trust" – is what we experience in our attitudes to lying (Faulkner, 2011, pp. 177–178). This is "shame or guilt" on the part of the offenders, "the reactive attitude of resentment in the "wronged'", and "a punitive attitude of disapproval or anger in the third parties" (Faulkner, 2011, p. 179; see Sripida and Stitch, 2006)

Summarily, Faulkner introduces an idea which he believed to be lacking in Williams's work, but which was actually also present there. That idea is that there are norms of trust that are socially sanctioned when breached (Faulkner, 2011, p. 179). Williams also invoked the motivations of honor and shame when illustrating how in our culture the internalization of truthfulness is brought about (Williams, 2002, pp. 120–121).[7]

What Faulkner appears to have overlooked, however, is that the intrinsic motivation of the virtues of truthfulness is assumed to work because there is a conception of knowledge, which Williams does not argue for, which makes truthfulness everyone's proper concern. This is Craig's social conception of knowledge, whereby testimony plays such a pivotal role. On Craig's conception, I argue, testimony is a generative

---

7 A difference is that on Williams's genealogy the norms would be internalized, and would come to be publicly recognized as part of the epistemic ethos. It appears it is this point which Faulkner tries to resist, because it enables a reductive solution, as explained. Yet, this reductive solution is one which Faulkner himself also suggests, when he presents his trust theory as a hybrid of the reductive and non-reductive account, in the form of providing an evidential explanation of knowledge gained through testimony, which explanation works in cases where one cannot extend affective trust but only because there are social norms of trust in play in testimonial practice, according to Faulkner (Faulkner, 2011, p. 191).

source – despite possible appearances to the contrary. While one individual may indeed give rise to a knowledge claim, and hence may be the originator of some proposition, the business of flagging that source as "trustworthy" (compared to other sources), and hence the attribution of the status of knowledge, is a communal affair.

In my spelling out of this conception, then, my focus is on actual claims to knowledge, and how they come into existence by being socially accepted as knowledge. The community plays a generative role of individual instances of knowledge by according to some people authority with respect to what is most likely true, and/or by directly according some statements the status of knowledge – when the source has come to be detached from the originating source of information. Indeed, Craig makes clear that his "state of nature story" assumes from the outset that knowledge is a social phenomenon. This is because he assumes that the concept has a social role and meets a social need. That social role is the role of flagging good informants.

Moreover, the business of flagging – of deciding who and what is trustworthy, and what information is trustworthy – is an ongoing process. The reason why I have emphasized the role of the community in Craig's account is to underscore that a corollary of his social conception of knowledge is to regard testimony as a social practice, and notably as a practice that is generative of individual bits of knowledge. It is this social conception of knowledge which is assumed by Williams. Moreover, it is because of the social generation of knowledge that Williams needed to worry about a complex society's ability to keep a notion of truth and of objectivity alive.

In saying that we need external sanctions, Faulkner does not acknowledge that knowledge, on the conception deployed, is a steep good (Thomas, n.d.). A steep good is one that is delivered in its entirety, or not at all – like a bridge over a river. (Two thirds of a bridge is not a bridge.) As such, it is a good which we can only have because we work together. Hence, the social nature of knowledge ensures that all of us have a reason to care that testimonial practice is not contaminated. All of us have reason to care for truthfulness, and for the proper assignment of epistemic credit to creditworthy sources. This is why the virtues which Bernard Williams posited do not, in the end, stand in need of external sanctions. Incidentally, it is also why, in large and complex societies, the social influences that can threaten to influence epistemic outcomes need to be kept in check. As I argued in chapter III, both Bernard Williams and Miranda Fricker start from the assumption that knowledge is socially generated, when they worry about the ways in which individual human beings or social power structures can wield undue influence and distort the proper functioning of testimony as a knowledge-generating practice.

5. Williams's Genealogy of Truthfulness (As Needed for Knowledge)

The difference between my position and Faulkner's becomes clear when looking at the way Faulkner understands Williams's genealogy of the virtues of truth. It is clear that Faulkner does not regard Williams's genealogy as building on and assuming Craig's genealogy of knowledge. Accordingly, he does not make a point that Williams tries to explain how knowledge is socially generated, and, specifically, how successful epistemic cooperation is to be explained. Instead, Faulkner canvasses Williams's aim more narrowly as follows:

> What Williams aims to do in *Truth and Truthfulness*, amongst many other things, is to provide a philosophical explanation of how testimonial knowledge is possible (….) This amounts to a philosophical justification of the presumption of testimonial reliability. (Faulkner, 2011, p. 172)

Faulkner claims that Williams seeks to provide such an explanation or justification of testimonial knowledge because, so he claims, Williams has a reductive understanding of testimony as a source of knowledge (Faulkner, 2011, p. 172). In making this claim, Faulkner points to Williams's regarding it as "a piece of bad advice" (Williams, 2002, p. 111) to trust a speaker in the absence of having ascertained that there are reasons for doing so, and cites Williams:

> (I)t may be said that a hearer never has a reason for believing that P which lies just in the fact that a given speaker has told him that P. He has to believe also that the speaker (on such matters, and so on) is a reliable informant. (Williams, 2002, pp. 77–78)

Interestingly, Faulkner does not cite Williams's point that a hearer should ascertain whether or not the "normal circumstances of trust obtain" (Williams, 2002, pp. 110–111). Also, while it may be a "piece of bad advice" to trust a speaker without ascertaining that the right conditions for such trust obtain, Williams has not denied that the hearer could end up with knowledge from the speaker on the basis of mere trust, provided that these (right) conditions are in place. And there is also a rather striking passage where Williams conjectures that being truthful is only required when someone is deserving of it, hinting at a sense of justice being involved.[8]

Within the set-up of the debate between reductivists and non-reductivists, who share the assumption that testimony is at best a second-hand source of knowledge, it is perhaps understandable that Williams is classified by Faulkner as a reductivist, given Williams's emphasis on there being a rather active role on the part of its recipient. However, Williams's own emphasis on the conditions of trust obtaining should casts doubt on that reading. Indeed,

---

8 This sense of justice, so I argue, is in fact what we ought to label "epistemic justice" as it is precisely the alignment of truthfulness (or authority) and trust (or due credibility) which has been central to Miranda Fricker's work (even as she focuses on cases of epistemic injustice).

Faulkner himself finds it hard to reconcile his and Moran's reading of Williams as a reductivist with Williams's non-skeptical, trust-based solution to the "cooperation problem" (Faulkner, 2011, p. 176, p. 178; Moran, 2005, p. 339). As Faulkner notes, Williams's invocation of the importance of the intrinsic valuation of Sincerity as a solution to the cooperation problem does not sit well with Williams supposed reductionism (Faulkner, 2011, p. 176). Here I argue we should side with the former, not the latter, as Faulkner does (Faulkner, 2011, p. 178). This is also the better interpretation, given that appealing to the intrinsic valuing of Sincerity as a form of trustworthiness is Williams's "official solution to the problem of cooperation",", as Faulkner notes (Faulkner, 2011, p. 178). This is right, and so Faulkner's interpretation of Williams has to be amended.

6. Consequences of a Generative Account of Testimony

I understand testimony as generative of knowledge. The central idea here is that individuals co-create knowledge through social and joint validation of knowledge claims. My account of how knowledge is generated is similar to that provided by Kusch (2002, Ch. 5). Kusch builds on a theory of social institutions developed by Barry Barnes, David Bloor, and John Searle (Kusch, 2002, p. 63). He explains knowledge as something that is shaped by a community's performative speech acts. The idea is that new knowledge is generated by the community through the acceptance of claims as knowledge. Such acceptance may be explicit or implicit. Kusch makes clear that the way this is done is by the commitments people take on and the entitlements that are validated. Very often the communal performative beliefs are not explicit, and not a single act, but dispersed as a belief that is evidenced by people's speech and the inferences made. On Kusch's view, as on mine, knowledge is a social status. He explains this means that knowledge is a bundle of commitments and entitlements. These can only exist in a social setting.[9]

An attractive aspect of Kusch's view is that it gives a one-size-fits-all account for both testimonial knowledge and for empirical knowledge, whereby testimony – described as the general and generative practice is it – is central. The reason is that speech acts are central to the account. The account is attractive, too, because it gives a very clear and direct sense to the claim that knowledge is a social status. On his view, knowledge comes into being through communal agreement. The view is also attractive because it makes clear that such agreement is neither hypothetical nor a matter of casting a vote at some convention: it is the cumulative result – the emergent pattern – of the many acts by the individual members of a community. Even while agreement is what is sought for, the result will also reflect the ordering – i.e. social

---

9 Kusch argues that normativity – the difference between "is right" and "seems right" requires social relations (Kusch, 2002, p. 174, Ch. 14) He argues that various epistemological theories cannot explain the rationality of empirical belief, because they fail to be able to situate normativity.

statuses – accorded to the people, as some people's word may be given a higher importance or prominence than that of others.

While Kusch's view is attractive, what is lacking in it is a focus on the standards for knowledge, which are to be socially determined, and on the epistemic efforts that individual members of the community make to abide by these standards: presenting and assenting to claims as knowledge only when they are to be presented and accepted according to the social standards. My proposal to remedy that problem is that we need another version of the way in which testimony is generative of actual knowledge claims, which is more akin to the idea of social verification. That is, we need to amend Kusch's account where it appears to rest content with mere "agreement": we need to make clear that there is an idea of validating claims by using our best social standards – which also, not quite incidentally, explains why a multitude of different perspectives is sought.

On the view of social verificationism, which is associated with the JTB- account, and defended by Alvin Goldman (1999), knowledge can be social in the sense that the creation of warrant can be a social affair. Collectively, we create a body of warrant – which Faulkner terms "an extended body of warrant" (Faulkner, 2011, p. 22). On my view, however, warrant is not required for knowledge, and testimony is not in the business of creating warrant: testimony is in the business of creating knowledge. On my view, justification is not the right level of analysis, because it does not fit with ordinary practice, and delivers us the misleading idea that testimony yields warrant. It belongs to the individualistic conception of knowledge whereby we should either accept full externalism about justification; or alternatively, place epistemic responsibility with the individual who acquires the warrant and remains individually responsible.

I am not arguing that there is no role for justification. Of course, there is a role for it, which is that people who assent to a claim or present it as knowledge, need to be able to justify doing so, answering queries and challenges in cases where reasonable doubts are raised. Contrary to the view that we are collectively creating warrant in testimony, I maintain that testimony is about socially ascribing the status of knowledge to claims that appear to be acceptable. We do bring our many perspectives and divergent background knowledge to bear when we engage in testimonial exchanges, as we decide to believe or disbelieve a source and determine what claims are to be accepted and spread. The status of knowledge, then, is given on the basis of how a new claim fits with the existent claims, while all claims accrue their status of being knowledge on a defeasible basis.

Of course, we may – and often do – defer the judgement on the score of whether some claim is to qualify as knowledge to those who are socially deemed more capable to deciding on the fit [10] – even as, in such a case, we still have a duty to ensure that our trust in those

---

10 And the fact we often do this explains why testimony tends to be cast, wrongly, as merely about transmission. This is wrong, however, as it fails to grasp the implication of accepting testimony, which is that the claim becomes more entrenched while it is spread and its status as knowledge becomes more secure (i.e. harder to challenge, and becoming part of the body of claims which constitute the background by which other and newer claims are judged).

authorities is rightful. So, we may need to check, for instance, whether their credibility is not assumed too readily, and ask them to explain to us why their expert judgement is what it is: to provide the justification (which so many people think is required for knowledge). There is then a sense, like Faulkner too emphasized, in which we seek to have explanatory reasons for why some claim is to be accepted as knowledge rather than justificatory ones. But that sense is not because the justification that a claim is knowledge is to be left to some other individual, but because the status of knowledge is a function of a social practice which is complex and ramified.

It is the social practice that sets the terms for deciding on credibility, and it is the social practice, too, that determines the standards by which knowledge claims are judged. So, there is no one individual justification for why a particular claim is deemed to be knowledge; there are many incomplete reasons. A claim is not knowledge because S says so, or because S has a good justification for it. The claim is knowledge because it has come to be accepted as such, and in being so accepted, it has proved to be "stable", i.e. to fit with the existent body of knowledge – all validated claims – and to withstand social scrutiny and reasonable challenges. The claim that something is knowledge is in fact always a stake: it anticipates that the claim is able to survive future challenges, even while it is spread further and further: that it really is commonable.

The attraction of my alternative view is that it appears to be more compatible with a realism about truth than Kusch's view. It argues for a communitarian epistemology on the basis that the group is, supposedly, more likely right than an individual, and/or a better guide to truth. Kusch's view, by contrast, appears to be best suited to the idea that both knowledge and truth are relative to a community (Kusch, 2002, p. 174). The alternative, instead, can insist that while knowledge is what passes as knowledge, passing as knowledge is not sufficient for claims to be true, and truth too is required for knowledge. This allows us to say that some communities are better than others at getting at the truth: there can be an external standpoint even if none of us can occupy the position because of our attachment to one standpoint and to the viewpoint of some community or other.

A further difference with Kusch's account, then, is that I seek to spell out what "good" standards for the validation of a claim are, in a way that is not relative to a (particular) community. So, I maintain that we aspire to standards of assessment that put an external constraint on what should qualify as knowledge. In saying this I oppose Kusch's relativism about truth, which he believes follows from relativism about knowledge (Kusch, 2002, p. 174). I maintain that we can be non-relativistic about truth, even while the standards for what passes as knowledge are communally agreed upon as our best way of arriving at "real" knowledge. In order to have this view, I maintain that knowledge – our possession of truths – is fallible: claims can come to pass as knowledge even while they are not deserving of that status.

A mixed blessing is that on this kind of fallibilist view, the possibility of global

skepticism remains open. This is troublesome in the sense that it may appear that I am deviating from a commonsensical notion of knowledge in order to accommodate a realist notion of truth – whereby we assume that this is necessary because knowledge entails truth, which would mean truth is a condition for something to actually be knowledge. If so, it would appear that I trade on equating two rather different conceptions of knowledge. One whereby "actual knowledge" belongs to a special metaphysical realm; and one which is "our everyday concept" which we assumed to be attainable, and is really "fallibilist knowledge", i.e. only knowledge in a "weak" form. It would mean that all our items of knowledge are really only fallibilistic knowledge claims.

I do not regard this idea as problematic; nevertheless, there is an alternative. This alternative – presented by Cheryl Misak, who defends a form of Pierce's pragmatism – is to say that a realist notion of truth is possible and compatible with minimalism about truth (Misak, 1991/2004, pp. 37, p. 129). Misak maintains that Pierce successfully shows that truth is neither a metaphysical property that transcends us, i.e. that is out of reach in some Platonic realm, nor to be equated to our best current opinion, which would make it dependent on our perspective:

> I shall argue that he succeeds in establishing a position which avoids taking truth to be something that transcends all perspectives and avoids taking it to be something that is relative to different perspectives. (Misak, 1991/2004, pp. 1–2)

Rather, so she insists, Pierce argued that if some hypothesis is true, then what follows from its being true (its practical consequence) is that it cannot be improved upon (Misak, 1991/2004, p. 42). This is how she argues we should understand Pierce's claim that truth is the hypothetical agreement of a would-be community at the end of inquiry. On her view, far from conceiving truth as defined by epistemic success, Pierce was a fallibilist: so, what we consider true might be overturned later (Misak, 1991/2004, p. 51).

The main point on which Pierce insisted, then, according to Misak, is that of there being a link between truth and inquiry (Misak, 1991/2004, p. 41). The practical relevance of truth, its role, is that it is to guide inquiry, if truth as the aim of inquiry (Misak, 1991/2004, p. 41). Supporting this point, she insists that the main concern of pragmatism is to avoid the outcome that inquiry be blocked prematurely. This way of viewing truth, then, makes us care for the methods by which we assess putative truths. These methods are to be such as to uncover reality. To care for truth, is to try and ensure that all evidence and arguments have been taken into account.

I am committed to a view, like Misak's, that sees truth as an external and methodological constraint on what qualifies as knowledge. Other than Misak, I take knowledge to be the end of inquiry, rather than truth. For knowledge, if properly ascribed, indicates possession of truth. At the same time, the pressing question – which Kusch for one would want to see answered – is: How can reality constrain our judgements? (Kusch, 2002, p. 97). For, as we saw, the fact that knowledge and truth are both a matter of our

(social) interpretations and normative judgements have been the basis of his relativist stance. Can we argue that reality is what it is, mind-independently, while it nonetheless constrains what truth and knowledge are?

Misak believes that there is room to do so at least when it comes to truth, and invokes Pierce's pragmatist stance towards both reality and on truth. In doing so, she is explaining objectivity in the sense of mind-independence:

> The pragmatic view of reality is that reality is the 'object' of true beliefs—it is what true beliefs are about. Reality is what beliefs in the final opinion would fix on (Misak, 1991/2004, p. 131)

> He makes the same point about truth. What would be ascertained to be true would be so ascertained whatever anyone here and now thinks. A hypothesis may be believed, then doubted, and then believed again, but this does not alter whether it would be believed at the end of a prolonged inquiry. Independently of whatever any 'definite' group of inquirers may think about the truth-value of *H*, *H* either would be or would not be a member of the final opinion (*CP* 5. 565, 1901). Thus the truth-value of *H* is an objective matter—it does not depend on what anyone happens to think. (Misak, 1991/2004, pp. 131–132)

This argument, however, may not appear to help much, in so far as Kusch denies the crucial premise, which is that reality is what constrains our judgements. Yet, a question is why he should think so. Here it appears pragmatists have a point, that reality is what is independent of our thoughts. In a similar vein, Bernard Williams too argued that our survival depends on our adjusting correctly to our environment. It seems somewhat strange to deny that reality should constrain our judgment, even if Kusch may insist that there are different possible interpretations. While we appear to reach a deadlock here, there is an aspect of the approach that means that Kusch's conclusions about truth can be resisted. For, while Kusch seeks to define truth, a point of convergence between Pierce and Williams is that they both stop short of defining – and hence of reducing – it.

Misak is clear that she thinks it is a mistake to cast Pierce as doing what he saw verificationists do, which is to *define* truth in terms of human epistemic success (Misak, 1999/2004, pp. 42–43).[11] Instead, truth is something which we should aspire to because it furthers the epistemic goal of stable beliefs (Misak, 1991/2004, p. 138). While truth on her account is instrumental, its instrumentality concerns the methods that we should use and our assessment of these methods: truth is to be our standard (Misak, 1991/2004, p. 82). Truth as an aspiration is not an empty standard, however. It is a standard that serves us to be real: to take into account all evidence and arguments and all the perspectives that can be brought to bear on the formation of beliefs/claims and their verification. All sources of information and of doubt that are potentially relevant for overturning the belief need to be included (Misak, 1991/2004, p. 63).

If this approach, of viewing truth as instrumental to knowledge, helps us to rescue

---

11 I should note here that I resist the same thing for knowledge.

truth and the nature of objectivity, which I claim it does, then a key to it is that truth is not to be reductively defined in instrumental terms, as Williams argued. Truth, as I conceive it, is a constraint on knowledge which is to serve the point that there is a match between what we believe and reality – the world as it is. Differing from Misak, I deny that truth – in the sense of stable beliefs – is the purpose or end of inquiry; knowledge is the end or purpose of this practice (see Misak, 2004, p. 40). What I do believe, however, is that knowledge claims which stand uncontested for longer are more likely to hit the mark of being true, as do a number of other authors (Goldberg, 2011; Goldman, 1999; Kornblith, 1999; even while we should remain vigilant; Quine and Ullian, 1978, p. 39; Longino, 2002, p. 152).

Therefore, to the extent that our vision of truth is similar, my dispute with Misak appears to be semantic: I substitute the terms "claim" where she uses "belief" and where many belief-theorists about knowledge use "belief",", too, instead of the term "claim" I do the same. My substitution of these terms is explicable by the social account of knowledge I have argued for and by the fact that I substitute "knowledge" for the end of inquiry, where she uses "truth" (in the sense of stable belief).

However, other than Misak I also do not endorse any further *definition* of knowledge as that would be incompatible with my "knowledge first" commitments. In addition, Craig appears right that equating knowledge with stable belief would run counter to the intuitions that beliefs which are shaped by upbringing and indoctrination are knowledge – though, of course, in the sense of surviving contestation, I think this kind of intuition is already answered (Craig, 1990, p. 7). If so, then the idea that true beliefs are stable ones makes sense, even if we do not yet arrive at a definition of knowledge that way. I remain committed to the knowledge first epistemologist's claim that truth is indefinable, just as I am equally committed to Welbourne's claim that this does not obscure any understanding the social role of knowledge but in fact helps us to appreciate it.

On my view, truth is a criterion for the assessment of knowledge claims, and is what serves to achieve the match or our assertions with reality. Knowledge claims that are "true" should be uncontested socially *while we make use of social standards for truth that are constrained by reality and serve to ensure that our assertions match it* – by including and accommodating (or reasonably rejecting) arguments, evidence, and the perspective of all other human beings.

It follows from this understanding of what it takes for knowledge claims to be true is that there ought not be any arbitrary exclusion of others as sources of knowledge, or of particular evidence and arguments, because doing so undermines the possibility of knowledge being what it ought to be: stable and commonable. It has been denied, already by others, that truths can be derived from knowledge, or even that we can derive what standards are those that yield truths. However, I do believe we can evaluate our standards as better or worse than other standards, and that we do so by looking at how

easily claims to knowledge which have come to be accepted are defeated: we can see how many false beliefs are had, and evaluate the success and "truth-conduciveness" of our knowledge system according to that fact.

What I take from Misak, then, is the thought that if our standards of assessment (i.e. are criteria for truth) are successful, in the sense of linking us to reality, then they yield knowledge. This is not what Misak herself says. Her point is one about truth: that hypothetical agreement at the end of inquiry is neither a definition nor a criterion of truth, but what we would expect concerning a statement if it is indeed true (Misak, 1991/2004, p. 42). On her rendering, truth is linked to inquiry in that it describes what a we should expect of a true belief, while the role of aspiring for truth is that it constrains what methods we can use to validate our beliefs (Misak, 1991/2004, p. 41).

An important point about truth, which helps us to distinguish it from reality (which it may be equated to if we think of what a concern for truth is about), is that because truth is a methodological constraint on knowledge, *it can be used to assess our standards for assessing knowledge claims*. This is an important point because it can explain the sense in which "truths" are contested. When this happens, there is a debate about what, methodologically, will deliver knowledge. Contestation that something is true, is not a mere denial that something is knowledge, but is a challenge to its passing as knowledge: to its being viewed as properly assessed and to the standards employed being used to assess future instances. As such, the idea of truth as a constraint on knowledge claims to properly pass as knowledge – i.e. on the right criteria being used for knowledge assessments – is distinct from the idea we have of "reality" (the way the world is).

Also relevant, vis-a-vis the point just made about truth, is that the fact that (a) knowledge (claim) is fallible does not mean we cannot evaluate that some methods are better than others for arriving at knowledge. This we can do, for instance, by considering how many knowledge claims are debunked – especially when these claims were well-entrenched as knowledge. Also, we have space for evaluating the system, by evaluating how easily our claims are shared and/or defeated after being shared – as well as considering how easily our standards allow a claim to pass as knowledge, and how much care people take to ensure that the standards are followed properly, and how readily people resist false claims to knowledge and have access to platforms to protest false claims.[12]

Validating knowledge claims for their truth – and ideas or hypotheses for their plausibility – is a reciprocation of the trust placed in one, as a receiver, and as a responsible member of the community. Something becomes knowledge when it passes as such: when it is

---

12 Note that Pierce, according to Misak, rightly worried that stability is not sufficient to ensure the methods would yield knowledge as opposed to make-belief (Misak, 1991/2004, p. 56-58). Likewise, claims being debunked is not in itself a bad thing and we may deny that a mark of a good epistemic system is that the claims are not easily debunked. At the same time, an epistemic system in which claims are readily debunked may not inspire the needed trust in the methods used, and would appear to be fragile from that viewpoint, signalling a lack of care – unless there are explanations for why so many claims turn out false that inspire trust in the new methods. Some level of stability appears required for trustworthiness, then, and for ensuring that the cooperative spirit, or the ethos of epistemic justice, is not undermined.

deemed that there is sufficient evidence for it, that all relevant perspectives have been raised and considered, and when no convincing counterargument has been made.

How is this point reached, and how can we know that this point is reached? Is it even true that there is ever an end-point to inquiry? To answer the last question first. There is a point where the demands of the community have been satisfied, and where the efforts of its members may be better employed verifying claims with lesser likelihood. So, there is a point where we are satisfied that something is knowledge, even if there is strictly speaking never an endpoint to inquiry, because we realize that we are fallible and that any knowledge claim, may have to be retracted. This is a point emphasized by Quine and Ullian, when they insisted that much of what is regarded as "common knowledge", in the sense that "it is regarded by large numbers of persons as true and by almost no one to be false" will be repudiated, a claim which they base on "the intellectual history of the species" (Quine and Ullian, 1978, pp. 39-41).

What their metaphor of "a web of beliefs" beautifully illustrates, is that there is ample room for an ongoing and dynamic creation of knowledge even if and when we use our background knowledge and beliefs in the process of assessing new ideas and claims to knowledge (Quine and Ullian, 1978; see also Quine, 1951, p. 39). In the process of doing so, we may find that some of the claims that we earlier accepted as knowledge need revision: some of our putative "background knowledge" is not knowledge after all.

Quine's image of beliefs that lie at the core of the web also pictures nicely the point that some beliefs are so entrenched in our worldview that it does take more to question those beliefs – that we readily assume form our shared beliefs and common background – than other beliefs (Quine and Ullian, 1978; Quine, 1951, p. 40–41). On a picture whereby this translates to the social nature of knowledge, we say that these claims to knowledge have much more social support, social importance and validation, and are therefore much harder to question than others. Challenges to what is accepted as knowledge are hard; claims that have been accepted as knowledge, and that pass as such, have acquired some solidity and are harder to question than claims that do not have this status as yet. They form our starting point and they also define our social being.

It may be responsible epistemic conduct in some cases, to take a claim on trust because it is already epistemically "proofed" by many others – or by people who are in a better position to validate the truth of the claim in question. We tend to take it that responsible behavior, including epistemically responsible behavior, is also behavior that recognizes that we need to defer to others on the score of the status of some knowledge claims, i.e. that we defer to the reigning social judgement that sufficient proof has been given. When we are asked to join a community of knowledge, we are asked to trust that others have made the efforts required – including those of socially validating the claims made. In the right circumstances (namely, those of trust), such responsible behavior yields to authority for their verdict.

In my view, there are three important things that determine the right amount

of trust, or deference, in a given case. The first concerns whether or not the person who submits that something is knowledge is regarded as having greater competence regarding the matter at hand. Second, whether the thing asserted fits with the accepted knowledge, to date. Third it is important to understand what is at stake in accepting it as knowledge; are there detrimental effects, and possible insights or other sources which would be relevant which are being overlooked? As such, there is a back and forth between the types of demands of individual responsibility placed on the members of an epistemic community.

Conclusion

I have argued that testimony is a generative source of knowledge. Testimonial practice works to generate knowledge by its involving a broader community effort. The social act of testifying to some truth and the social scrutiny that follows by filtering what is passed on is a mechanism that increases the likelihood that only true statements come to pass as knowledge. That is, I believe that knowledge is a good that requires social cooperation, and particularly, epistemic cooperation. The idea, which is also part of Kusch's account of testimony, is as follows: while a claim to knowledge comes into being when it is stated its (defeasible) status as a piece of knowledge is begotten through being socially accepted as such, having been subjected to social scrutiny and contestation, whereby people assess whether or not the claim is true using our best criteria for truth.

Knowledge, then, is a social status that can be retracted when it turns out that there is reason to doubt the truth of what has been passed as knowledge. So, passing as knowledge is not sufficient for knowledge, but – pace Welbourne, Williams and Fricker – knowledge exists because there is a social practice of testimony that requires and sustains it, and because we cooperatively strive for truth(s).

Chapter VI | A Theory of Epistemic Justice Centered on Epistemic Power

Introduction

In this chapter I provide the outlines for a theory of epistemic justice as essentially a hybrid: as both an ethical and a political notion. In order to do so, I will develop an interpretation of Miranda Fricker's *Epistemic Injustice: Power and the Ethics of Knowing* (Fricker, 2007). Fricker was the first to develop the idea that our everyday social epistemic practices include both epistemic wrongs and epistemic injustices. She arrived at this conclusion "bottom up", by attending to cases that arise in our epistemic practices. Notably, in her account of the practice of testimony, she drew on examples from fictional works, e.g. to illustrate when a person is not believed due to a racial bias on the part of the hearer she chose the character of Tom Robinson in Harper Lee's *To Kill A Mockingbird*.

While I am much indebted to Fricker's work, the aim of this chapter is to move beyond it. First, I aim to provide a theory of epistemic justice that can explain which cases are *duly* regarded as epistemic wrongs and as epistemic injustices. Second, I aim to provide a theory that can explain both the "discriminatory" cases and the "distributive" cases. This is necessary, so I argue, for the reason that Fricker recognized the latter type of case only in later work. In *Epistemic Injustice* she at first cast doubt as to whether the second type of case raised a "distinctively epistemic" issue (Fricker, 2007, p. 1). Third, I will show that epistemic justice is not merely a useful slogan, but that it can be operationalized as a specific subcategory of justice.

I argue that a natural understanding, which can achieve all of the above aims, is that we ought to understand epistemic justice as centrally concerned with regulating epistemic power. This understanding is natural in that it corresponds to Fricker's initial idea of the inverse notion of epistemic injustice. It supports the intuitive idea that justice is centrally concerned with power in its various forms. On the conception I envisage, the broadest notion of justice – "justice overall" – regulates the use and allocation of all the different forms of social power: physical, economic, political, legal, epistemic, psychological, sexual, etc. Epistemic justice, then, fits within this framework as a subcategory of justice concerned with *one* form of social power, epistemic power.

Adopting this framework allows me to explain how epistemic justice has its own specific form of normativity. I will be working with a pluralist normative framework of justice that encompasses both the aspects of recognition and the aspects of distribution

that have, so far, been separate projects in political philosophy. Part of this pluralism is the acknowledgment of the fact that there may be reasons of justice that outstrip epistemic justice. That, in turn, implies that there is room for a limited set of justified reasons for not acting as epistemic justice requires in particular cases.

My approach to epistemic justice is different from Fricker's in two ways. First, I work in so-called "ideal theory", striving to provide an ideal of epistemic justice, positively conceived, while Fricker's work on epistemic injustice is in "non-ideal theory" (Fricker, 2007, Preface, pp. vii–viii, p. 2; see also: Sen 2009, Ch. 4; Simmons, 2010, p. 35). Second, I deviate from her virtue-ethical framework. While I am sympathetic to the virtue ethical approach, it is my aim to provide a theory that can supply us with articulated norms and contestable criteria.[1] In addition, my view goes beyond the agent appraisal that is central to virtue ethics, and the act appraisal that is supplementary to it, in order to focus on the social conditions that enable these narrower forms of judgement.

2. Fricker's Cases of Epistemic Injustice

Fricker seeks to establish, through giving examples, that there is a distinctive class of "moral bads" – wrongs and injustices– that are "specifically epistemic" and neglected by philosophy. They are discovered when we attend to our "social epistemic practices" (Fricker, 2007, p. vii). As she characterizes this class of wrongs and injustices, they are theorized as "consisting most fundamentally in a wrong done to someone specifically in their capacity as a knower" (Fricker, 2007, p. 1). To elucidate what is meant by this, and to make her case for the existence of epistemic injustice and the pervasiveness of such injustices, Fricker provides two "primary cases".

The first primary case is one of *testimonial injustice*. This is the injustice of not receiving due credence as a speaker because of a prejudice on the part of the hearer which causes the hearer to give the speaker lesser credence than she deserves. The example which Fricker gives is derived from *Kill a Mockingbird*, and concerns Tom Robinson, an African-American man who is not believed by an all-white jury when testifying in court, and who is sentenced as a result of this, even while he is innocent.

In this example, as Fricker makes clear, the epistemic failing of the jury is that they give the speaker less credence than he epistemically deserves. The moral wrongness of such an act is that the person in question is not treated as he epistemically should have been treated, namely as a (putative) knower, i.e. as someone whose testimony is to be taken seriously.

Testimonial injustice, then, describes an epistemic failing on the part of the hearer

---

1 This is not to deny that a normative theory can also be cashed out in terms of virtues.

which of course also disadvantages the hearer, who does not obtain the knowledge he or she could have gained, but it is an epistemic wrong because it morally wrongs the speaker who is not taken seriously to the extent that is due. According to Fricker, given how a person's rational capacity is related to the conception of one's moral worth, such an undue denial of someone's epistemic status as a knower can ultimately harm that person in his or her human dignity (Fricker 2007, p. 54).

The second case that she considers is exemplary of the class of *hermeneutical injustices.* This is the injustice of not being able to make sense of one's own experience because the necessary hermeneutical resources are lacking as a result of social inequality. The example Fricker gives is of a woman who cannot understand her experience of sexual harassment because the society in which she lives lacks that concept, which in turn is due to women not having a role in shaping the collective hermeneutical resources of that society. To clarify how this is an epistemic injustice, Fricker writes:

> I seek to explain this sort of injustice as stemming from a gap in collective hermeneutical resources – a gap, that is, in our shared tools of social interpretation – where it is no accident that the cognitive disadvantage created by this gap impinges unequally on different social groups. Rather, the unequal disadvantage derives from the fact that members of the group that are most disadvantaged by the gap are, in some degree, hermeneutically marginalized – that is, they participate unequally in the practices through which social meanings are generated. (Fricker, 2007, p. 6)

As Fricker presents her case of hermeneutical injustice the injustice is due to an asymmetry in social power. Some members of a community are less able than are other members of the same community to achieve self-understanding and to render her experience intelligible to others.

This disadvantage is not merely bad luck or unfortunate, but unjust. It is not an accident that a society does not have the conceptual resources to cater for those who are socially oppressed or marginalized, but is a consequence of that oppression. Fricker argues that it is the unequal and undeserved disadvantage which is unjust. This injustice is epistemic, because the disadvantage is an epistemic one, namely, an inability to render one's experience intelligible to oneself. Gaining self-understanding and self-knowledge is a cognitive activity. The same goes for making one's own experiences socially understood to others.

The two cases illustrate what Fricker means by injustices of a distinctively epistemic kind. In testimonial injustice, one is wronged in one's capacity as a knower. In hermeneutical injustice one is disadvantaged in one's capacity as a subject of knowledge – namely, a subject of self-knowledge. So, the two cases show in what way wrongs and injustices can be specifically epistemic: they affect a person either as a knower or as a subject of knowledge (and of social understanding).

3. Social Power on Fricker's Conception

Fricker notes that social inequality can serve to explain why the epistemic injustices arise, in both cases. Social inequality is relevant because we engage in epistemic practices as socially situated epistemic agents, that is, as agents who stand in certain relations of dependence to one another, whereby some of us have more social power than others (Fricker, 2007, p. 3). Fricker gives a working definition of social power as, roughly, one's ability to influence the social world: that is, "a capacity we have as social agents to influence how things go in the social world" (Fricker, 2007, p. 9). But given that such power is often coercive, even if it need not require direct agential action of one agent controlling another, she further characterizes social power as:

> (A) practically socially situated capacity to control other's actions, where this capacity may be exercised (actively or passively) by particular social agents, or alternatively, it may operate purely structurally.
> (Fricker, 2007, p. 13)[2]

An example of an active reliance on identity power is when a man patronizes a woman, which he gets away with in virtue of his social identity power. This power, in turn, is grounded in a collective social imaginary according to which a male has more authority than a female. This is active reliance because the male implicitly uses the stereotype to assert his authority over others (Fricker, 2007, p. 14).

There can also be *passive* operations of such power, for example, when a person with less relative power is caused to not even make an effort to try and be heard. This effort is pre-empted by the fact that their identity-based power is too weak to have a (fair) chance of success. The structural operation of social power happens in a case where a group of people is being controlled by a certain social structure or conception without there being any agent who actively or passively coerces them, but whereby this control always preserves the power involved. Fricker gives as an example Foucault's "delinquents" in his *Discipline and Punish*, and the example of a dominated social group of which the members informally disenfranchise themselves because they consider that "voting is not for them" (Fricker, 2007, p. 13; Foucault, 1979).

So, while in principle *any* failure to give another person due credence, for instance because one is simply not paying enough attention to the testifier, can be seen as a testimonial wrong, the more serious cases of testimonial injustice track the person throughout his or her life because they are based on that person's identity. They

---

2 Fricker's conception of power has been influenced by both Foucault (1979) and Thomas Wartenberg (1990). I will discuss how I have adapted her conception and how I conceive of power and its proper place in political philosophy and particularly as related to my pluralistic conception of justice later, when I focus on epistemic power and the positive conception of epistemic justice.

have the potential, cumulative, effect of eroding someone's confidence, and hence of incapacitating them as a knower.

These more serious cases, so Fricker maintains, revolve essentially around "identity power". Fricker defines "identity power" as "a form of social power which is directly dependent upon shared social-imaginative conceptions of social identities of those implicated in the particular operation of power" (Fricker, 2007, p. 4; see also p. 14). This is to say that there are social conceptions that, whether we want to accept them or not, are shared between us and others in our social environment and influence how others see us and how we view others vis-à-vis ourselves. Fricker notes, further, that the specific collective conceptions that she believes play a role in the central cases of testimonial injustice "amount to stereotypes" and that such stereotypes "may or may not be distorting ones" (Fricker, 2007, p. 15).

Thus, given that testimonial injustice involves negative identity prejudice, while hermeneutical injustice is the result of social marginalization, both the two primary cases as Fricker presents them (in their most serious form) involve reference to one's membership of a given social group, as defined by the collective social imagination. So, both types of epistemic justice are discriminatory epistemic injustices, as they relate to social identity. There is a real sense then in which, as Fricker emphasized, the more serious forms of epistemic injustices are explained by the workings of social (identity) power (Fricker, 2007, pp. 155–156).

4. Epistemic Justice: Discriminatory and Distributive

At the outset of her book, Fricker notes that her conception of epistemic in/justice may not be the most intuitive sense of that term that would occur to the reader. For one naturally assumes that justice is about distribution, such that epistemic injustice would be primarily focused on the distribution of educational advantage. Fricker comments as follows:

> Given how we normally think about justice in philosophy, the idea of epistemic injustice might first and foremost prompt thoughts about distributive unfairness in respect of epistemic goods such as information or education. In such cases we picture social agents who have an interest in various goods, some of them epistemic, and question whether everyone is getting their fair share. When epistemic injustice takes this form, there is nothing very distinctively epistemic about it, for it seems largely incidental that the good in question can be characterized as an epistemic good.
> (Fricker, 2007, p. 1)

The reason she gives for setting aside this class of cases is that she wants to argue that epistemic injustice concerns cases where there is a specific moral-epistemic type of

wronging in a way which is not captured by existent frameworks of distributive (in) justice. Epistemic injustice is a *sui generis* category*,* one that has neither received attention from epistemologists nor from ethical and political theorists, because it belongs to an area of inquiry that falls between these two defined fields as they are practiced (Fricker, 2007, p. 2).

A sociological explanation of this intellectual neglect would be that those who want to address the issue tend to be those who are socially marginalized and hence not the ones who enjoy the status and power to introduce the term in academic and societal discourse. Traditional epistemologists approached their subject without a focus on the social situatedness of inquirers although with the publication of Fricker's book the situation is changing (Fricker, 2007, p. vii).

In later work, in response to constructive criticism by Coady, Fricker has conceded that epistemic injustice does intersect with, and have implications for, distributive issues (Fricker, 2013, p. 1318; Coady, 2010, p. 101, p. 109). She now regards "epistemic injustice" as an umbrella notion, which can cover both the discriminatory form of epistemic injustice and distributive cases (Fricker, 2013, p. 1318). She explains that both types fall under her overall characterization, namely that of "wrongs done to someone specifically in their capacity as a knower" (Fricker, 2007, p. 1). I concur with this judgement. On my view, both knowledge and credibility may be viewed as "epistemic goods" and their distribution is importantly a matter of epistemic justice, as they are a means to epistemic power. I conceive of epistemic power as a form of "power-to", namely as an ability to influence what comes to be regarded as knowledge and who qualifies as authoritative, and will argue that the use and allocation of epistemic power is the central concern of epistemic justice.

My concern about Fricker's concession within her framework, however, is that it threatens to undermine the claim that epistemic justice is a self-standing category of justice with its own dual moral-epistemic normativity. One reaction might be that this claim can be abandoned without loss. However, I believe there is a problem even if the general idea of "epistemic injustice" can still be used for strategic purposes. We have good reason to care about a sui generis notion of epistemic injustice. Only this can explain the intuitive pull of the cases to which Fricker has drawn attention. If there is no substance to the independent life of the concept, no specific type of normativity that it pinpoints, then it is easy to see how people will come to simply dismiss any invocation of the notion as either nonsensical or overblown.[3] Epistemic justice is weakened, as a category of interest, if it simply re-categorizes cases that are already treated as wrongings or injustices, but merely labelled differently. My own view is that it does identify issues

---

3 Also note that reduction allows that philosophers can continue to do "pure" epistemology or "pure" ethics. On reductive readings, we can leave the social out of the epistemic, happily continuing *a-social* epistemological questions, and we can regard epistemic issues within the moral realm as restricted to those that morality itself poses, e.g. about moral responsibility and blame.

that were not identified by orthodox understandings of the nature of justice. The neglect of epistemic injustice was, ironically, its own form of hermeneutic injustice: the denial to people of a concept that they need to understand the wrongs done to them.

A sui generis concept of epistemic justice ought, in my view, to form part of the turning away from the individualistic biases of orthodox epistemology and the restriction of the theory of justice solely to distributive issues. In these cases, theoretical tidiness and unity are not helpful. On its traditional understanding, epistemic authority unquestioningly attaches to those who "have" truth and objectivity on their side. Moral authority attaches to those who use the correct moral criteria, where those criteria are a matter of fact or reasonable consensus. What these conceptions of both types of authority appear to deny is that both of these theoretical outcomes may have been influenced by power. But if this authority is, at least in part, socially determined, then this possibility cannot be ruled out ab initio. An accurate result would depend on the proper inclusion and weighting of all the available perspectives (Longino, 1995, p. 384; Douglas, 2009, p. 128; Hong and Page, 2012, p. 67, pp. 70-71).

One might surmise that the case for a sui generis class of epistemic injustices would have been made if one could demonstrate that there are cases of discrimination that are wholly epistemically grounded. However, even that point has been contested. For it is not clear that discriminatory epistemic injustice is sui generis and a reductive challenge has been put forward to account for such cases.

Reductionists argue that the moral and political intersecting with the realm of the epistemic does not mean there is such a thing as epistemic injustice. The examples all involve a much broader concept, namely, social domination. Conversely, sceptics can maintain that, despite the dire consequences, the failure of a testimonial injustice is really an epistemic failure and not one which is in any way moral. Nevertheless, I will argue that there is a sui generis category of epistemic injustice that can cover both discriminatory and distributive cases of injustice. To further this goal, I will develop a positive conception of epistemic injustice that shows how and why it covers both classes of example by relating both to my more fundamental idea, namely, epistemic power.

If the category of epistemic injustice is to avoid the charge of being ad hoc, then my view is that we need to locate it within a broader framework that describes the just form of an epistemic community. That justice, in turn, is related to the idea of the proper distribution of epistemic power. This framework unifies my previous explanations of the strong sociality of knowledge and the generative role of testimony. If we lack this rationale, then the concept will not only seem ad hoc, but will capture both too little and too many wrongs and injustices. So, my version of Fricker's central conception of an epistemic injustice – "a wrong done to someone specifically in their capacity as a knower"– would read as follows: "a violation of the norms that regulate the proper use and allocation of epistemic power" (Fricker, 2007, p. 1).

The advantage of my approach is that it can accommodate both the ethics of knowledge

sharing via the notion of epistemic wrongs and its politics since the notion of an epistemic justice indeed is political as well as ethical– a point that Fricker acknowledged (Fricker, 2007, p. 8). The appeal to epistemic power can explain why domination is a central concern even if – on Fricker's characterization of the umbrella concept of epistemic injustice – domination is not central to spelling out the notion of epistemic injustice.[4] In addition, it can include a large class of cases which Fricker has omitted, and which some critics have rightly tried to include: distributive cases of epistemic justice, which are cases where inequalities in standing and access are not justifiable by recourse to some epistemic considerations but are instead clear violations of the norms of epistemic justice.

Moreover, this accommodation is not unnatural: distributive as well as discriminatory issues can rightly be classed as epistemic injustices, either because an epistemic failing violates the proper inclusion and weighing of all available sources or because there is an inherent epistemic injustice in the fact of the epistemically improper allocation of epistemic power which makes some people undeservedly epistemically worse off than others.

5. Skepticism regarding 'Epistemic Injustice'

Fricker's pioneering work has met with some critical challenges. The first type of criticism relates to the identification of cases. Fricker excluded distributive matters of how education and information are to be distributed. But it seems this is counter-intuitive because information and education will be relevant for realizing and expanding one's capacity of being a knower. [5]

Accordingly, an improper distribution of these epistemic goods can wrong a person in his or her capacity as a knower or as a seeker of (self-) knowledge. It seems, then, that Fricker's definition of an epistemic injustice as most fundamentally "a wrong done to someone in their capacity as a knower" already covers such cases too. The reason is that any social inequality that unduly affects one as a knower or subject of knowledge is an epistemic injustice so long as it results in an unfair epistemic disadvantage.

The second type of criticism relates to the dual moral-epistemic nature of the injustices and wrongs that Fricker describes. Some have doubted that there is such a dual nature

---

4 Fricker notes in later work on its relation to political freedom that epistemic justice is a "compound constitutive condition of non-domination" (Fricker, 2013, p. 1324).

5 This is so unless the idea of capacity is understood exclusively in terms of having a social status. The point is that the social status is part of what gives one epistemic influence, and hence, both the status and other things that can empower one, are concerns of epistemic justice. Credibility may be viewed as a status, when it is enjoyed, but also as something that is distributed. But aside from credibility, knowledge will impact who is able to exercise epistemic influence, for knowledge provides both motivation and content that one needs to exercise influence.

because they doubt that an "epistemic wrong" and an "epistemic injustice" are special species of justice. This requires that the moral and epistemic are linked. Instead, such critics think that all cases reduce *either* to an epistemic failing with some moral valence or consequence *or* reduce to an already existent moral failing or social injustice that surfaces in the epistemic domain. These are wrongs with an only incidental epistemic relevance. This claim is precisely that which Fricker gave as a reason for rejecting the distributive injustices related to epistemic goods as falling under the category of epistemic injustice. She considered the wrongfulness or injustice of such cases to be only contingently epistemic, as opposed to being wrongs and injustices that are "distinctively epistemic" in kind.

There are clearly cases in which the moral and the epistemic are connected, but only in an incidental way that does not warrant a new conception of moral/epistemic harm. For example, suppose a dictator orders a brain surgeon to operate on a person and the result is a lowered epistemic capacity for the victim. Even if the result is that this victim's capacities as a knower are diminished, or disappear, it does not seem that the moral wrongfulness is at all epistemic. Rather, the injustice is political and physical, and only epistemic in some secondary sense (if at all). If we accept that other forms which are only contingently related to the injustice are epistemic too – because of the type of harm that is effected say – we will have to accept cases that no longer match our intuitions of what an epistemic injustice centrally is.

It is illuminating in this context to consider Laura Beeby's criticisms of the category of hermeneutic injustice (Beeby, 2011). In the case of sexual harassment that Fricker presents as paradigmatic of hermeneutical injustice, Beeby objects that the injustice does not appear to be epistemic, since both are equally disadvantaged epistemically by the lack of the concept (and social understanding) of harassment (Beeby, 2011, pp. 582–583). It is grounded, rather, in the unequal relation between the woman and the man in this example.

Beeby is right to claim, so it seems to me, that the male protagonist was epistemically no better off in this respect (of understanding his own actions) than the woman that he harasses. To find the woman a victim of hermeneutical injustice then, and to a greater extent than the man who is in an equal epistemic predicament with respect to understanding his experience, it appears that we are introducing the effects of his actions into the case. After all, we do not usually judge that the harasser is equally to be considered a victim. Yet the man too could not understand his act as one of harassment. And as Beeby points out, he may well have wanted to know that he was causing some form of harm or distress. This raises the question why there is an epistemic injustice in this case, as opposed to a social injustice.[6]

---

6 Even worse, the explanation that Fricker is proposing may be exculpatory of the man's actions; after all, he lacked the social resources to be able to conceptualise his action as the wrong that it is. How, then, can we hold him accountable for it in the specifically moral sense of blaming him for not having a reason he ought to have had? "Ought" implies "can"; yet, in this case, he cannot access the very concepts needed to understand the wrongness of his action.

Anticipating this type of objection, Fricker had claimed that it is especially in the interest of the woman to understand this particular experience (Fricker, 2007, 151). However, it seems that if we assume that the man is well-intentioned, and that is not willfully complacent in exploiting his identity power, then we have no reason to say that the man is not equally disadvantaged as the woman, epistemically. So, Beeby's point that the injustice is social domination rather than epistemic injustice seems to have force.

One response to this critique is to argue that the injustice in question is derived from a specific type of marginalization that is already an epistemic injustice. It is already an epistemic injustice that some members of a given group cannot partake equally in the processes through which collective meanings are shaped. In fact, I would myself consider this the primary form of hermeneutical injustice. The fact that a woman cannot make sense of her own experience is an injustice that derives from her being hermeneutically marginalized and such exclusion is epistemically unjust.

Yet, I realize that this is not a reply that really helps, as it runs into a regress. For, it could be said – quite rightly – that that first hermeneutical marginalization would have as its the root cause social injustice, namely the social oppression of women. And then it could be said that that case is not an epistemic injustice after all. For then it is not a specifically epistemic injustice if it is the social oppression really is that makes for the injustice. I do not agree with that objection, but I grant that pointing to a prior type of epistemic injustice is not enough to answer Beeby's objection.

Instead, I argue that the reason why the case which Fricker described is an epistemic injustice – and indeed the fact of hermeneutical marginalization too – is because it basically consists in having an unfair epistemic position: an epistemic position that is different even while there is no good reason that can justify it. The way to differentiate between the man and the woman in this example is by noting that the man does not occupy the same *epistemically* marginalized position with respect to achieving self-understanding, which comes to being able to understand particularly important experiences and rendering them intelligible to others.

What is right, then, in Fricker's defense of her claim that the woman suffers an epistemic injustice as a result of the fact that the concept of sexual harassment was lacking is that the woman is unable to understand her own feelings while the man, who acted out of ignorance, does not have that same problem. He has, so we suppose, no issues with his self-understanding as a result of the lack of the concept of sexual harassment. Even if he may be wrong to not view himself as a perpetrator, there is no nagging self-doubt that affects his self-trust. Hence, a comparative disadvantage in terms of understanding one's own experiences is distinctive of the epistemic injustice in this case.

Yet, a comparative disadvantage is not enough for it to qualify as an injustice. It could be a matter of bad luck. The reason it is an injustice, is because there is a difference in the epistemic position of these two people (with respect to their self-knowledge)

that cannot be justified epistemically, by the amount of effort they put in to achieve self-understanding for instance. I think this is why Fricker repeatedly emphasizes that hermeneutical injustice presupposes hermeneutical marginalization. Indeed, as I will argue, epistemic justice centers round differences of *epistemic power*.

6. Epistemic Wrongs and Epistemic Injustices

Before I move on to give my positive conception of epistemic justice as centering on epistemic power, it will be important to clarify the distinction between epistemic wrongs on the one hand, and epistemic injustices on the other. Why this is so important will become apparent later, but let me say here already that attempts to reduce epistemic justice are more readily successful when we have not clarified the difference between these two categories. So, in this section I explain the difference between epistemic wrongs and injustices, even as I also argue that they are nonetheless rightly regarded by Fricker as having an important commonality: they are, so I argue, rightly viewed as having in common that they are instances of epistemic injustice, i.e. violations of the same type of moral-epistemic normativity.

To understand the difference between a wrong and an injustice it is helpful to focus on an important difference between the cases of testimonial injustice and hermeneutical injustice. There are two quite different ways in which epistemic injustice can come about. The first case, of testimonial injustice, stems directly from an act of one or more agents. Notably, the wronging is of an epistemic nature: the hearer fails to give the amount of credence which is epistemically due.

The second case, on the other hand – at least the way it is presented by Fricker – does not come about as the result of an epistemic act of one or more others. Instead, it is the result of social inequalities – and subsequent impoverishment of the social conceptual repertoire – that means that some people are cognitively at an unfair disadvantage.

In the second case then, the cause of the injustice is not epistemic, even if the outcome is an injustice. Here the injustice is epistemic not because of the type of act that caused it, then, but because of the type of disadvantage that was caused. As I have spelled it out, this disadvantage is unfair in so far as it is undeserved. It is undeserved because we expect people to be in a similar epistemic position unless there is a *relevant difference* that can justify why they are not in that same position. It is an epistemic injustice then, because it is an epistemic disadvantage that has *no epistemic reason* that could justify it.

Overall, what the two cases show is that we should draw a distinction between wrongs and injustices. Wrongs require an agential act or failure, injustices do not. That act failure is itself epistemic as it is a misuse or abuse of epistemic power. Epistemic injustices, on the other hand, concern the unjust (unfair) allocation of epistemic power. This can come

about through structural reasons, such as segregated schooling that perpetuates social inequality in the epistemic domain, or simply as a result of many acts each of which are individually morally permissible, as Elisabeth Anderson has argued (Anderson, 2012, pp. 164–165, p. 171). Because a misuse of epistemic power can cause a misallocation of epistemic power, a wrong may be classed as an injustice. This is to look at its outcome. Yet, there is only an epistemic injustice if there is an undue epistemic disadvantage which leaves one vulnerable to *epistemic* domination, which is a point missed if we focus on the more generic form of *(social) domination*, and the upshot of the generic focus is that we can no longer single out epistemic injustice in a meaningful way.

7. What Makes Wronging Epistemic?

From what I have argued thus far, it appears that epistemic injustices could arise, all too easily, from just any kind of failing. This makes it hard to say what makes them qualify as specifically *epistemic*, as opposed to wrongs to someone in their capacity as a knower whereby the failing is an epistemic one. In the case of the distribution of such "goods" as information and education, the only epistemic aspect appears to reside in the type of good distributed. As we saw, Fricker initially sought to draw a distinction between discriminatory and distributive epistemic injustice precisely in order to avoid the conclusion that epistemic injustices are only contingently epistemic.

This approach, however, has yielded the strange paradox that testimonial injustice is presented as surfacing only when one is giving knowledge, and not when one is on the receiving end. For, notably, all cases of receiving knowledge appear to be ones in which, according to Fricker, we simply think of epistemic goods and wish these to be distributed fairly. If that is true, then we could appeal to the existent framework of distributive justice, as there is nothing specifically epistemic about epistemic goods, or their just distribution, except for the nature of these goods (Fricker, 2007, p. 1; cf. Coady, 2010, p. 109).

While I agree with Fricker that there cannot be a merely contingent qualification of an epistemic injustice as epistemic, my point is different: I argue that we *can* keep epistemic failings central to epistemic injustices in the case of testimonial wrongs, regardless of whether it is a speaker or a hearer who is wronged. There is, so I argue, no reason to bracket those cases where the wrong affects a putative knowledge receiver from those which concern a knowledge giver. In the below, I shall illustrate this with my own example, which is in fact a historically true one, followed by the case of "preemptive" epistemic injustice that Christopher Hookway provided.

The example concerns famous ex-footballer Johan Cruiff, who, as a Supervisory Board Member of football club Ajax, had a vote on whom to hire as a new CEO for Ajax. Yet,

as the only one of the board members he is not told that hiring Louis van Gaal is on the agenda for the next meeting. As a result of this, he does not make it to the meeting, given that the invitation for the meeting (without its agenda being sent along with it) came so late, that he could not attend. At the meeting, the panel decides to hire Van Gaal. Cruiff finds out and doesn't approve, but is told that the decision is made according to the Statutes, and that if he had wanted to veto it, he should have been there.

Here the epistemic injustice consists in the fact that Cruiff is not told something which all the others know, and which they know he should have an interest to know. As such, he is not given a chance to convince the board otherwise, nor to cast his verdict on whether such a decision is wise – and indeed, he is disempowered politically (and, apparently, legally as well). The point here though, is to illustrate that the case starts with an epistemic wrong, and that the specific abuse of epistemic power in question is the wrong of not being given the knowledge which the others had, or at least not on time. It could be objected that this knowledge was knowledge he was entitled to receive, such that in fact there was a special obligation in this case to inform him. However, I think that the case still shows that not sharing knowledge with someone can constitute a wrong, and that that wrong is epistemic, because it is an abuse of one's epistemic power.

What this case confirms, too, is that an epistemic wrong can happen regardless of whether or not someone is socially powerful in rank. Of course, we could redefine what it means to be "socially powerful", as epistemic power plays its part too in that, but it seems clear here that the type of discrimination Cruiff endures has nothing to do with his social identity or with past social injustice suffered on his part, and instead came about because of his ability to block a hiring procedure, which ability proves his social (political) power. While it is possible to argue that Cruiff suffered an epistemic injustice as a consequence of the epistemic wrong, being unable to influence the vote of others or to argue his case, the main disempowerment that he suffered was of a political and legal sort, and so I would cast it as a political and possibly as a legal injustice – if it should have to be called an injustice as well as a wrong. What I argue though is that this case shows that the diminished (legal) decision power was effected illicitly, through an abuse of epistemic power.

It also actually makes very little sense to talk of the injustice of assigning someone a lowered credence only in a case where someone is giving testimony, i.e. aiming to share knowledge by speaking. Someone who asks a question and is not answered because they are not deemed worthy of an answer, even while in fact they would be perfectly capable of understanding the answer, which has been labelled a "pre-emptive epistemic injustice" is obviously wronged in this same respect of not receiving due credence as the person who gives testimony but is not being believed to the extent which is due (Hookway, 2010, p. 157). For the person is blocked from partaking in the epistemic practice. This has a bearing on the success of that practice, too, given that knowledge is generated by claims being validated and disseminated.

An example is a case where a secretary is treated as less worthy of information about organizational changes that will be affecting her job than is an employee who is her manager. The knowledge not being shared, due to a disregard for her capacity to provide useful input about the way her work may still be done, renders her as I would put it "epistemically powerless", having been excluded from being able to contribute her knowledge. Its consequence is first frustration, and after not being regarded as having a valued opinion for longer, a lack of motivation as well as a belief that she is indeed epistemically far less capable: her confidence is undermined.

In such cases where credibility is not denied but not even in question, it could even be that the same identity prejudicial mechanism is at work, such that the person is not given due credence because the sole cause of not being informed is the prejudice on the side of the knowledge giver which causes her to give you a lowered credence for various reasons. This could be, for example, because she judges you as not being worthy of her time, being deemed unable to understand the answer or simply deemed unable to reciprocate in the relevant way by contributing knowledge to what the person in the know already knows and could share with you.

Thus, my argument is that the case of testimonial injustice, as Fricker has described it, covers only some of the cases of injustice in the interpersonal epistemic practice of knowledge sharing. The reason why this is even more problematic is that a social conception of knowledge whereby testimony takes center stage ought to do justice to the reciprocity which is part and parcel of knowledge sharing.[7] If one espouses a dynamic conception of testimonial practice as generative of knowledge, whereby the moral-epistemic norms of reciprocal trustworthiness and rightful trust regulate communicative exchange, then the thought comes very naturally that in order to be able to contribute knowledge, one has to be empowered to do so.

The idea that there is a dichotomy between distributive and discriminatory cases of (testimonial) epistemic injustice, however, not only problematizes our understanding of testimonial injustices – whereby we lose sight of the plight of those who wish to receive knowledge. It also problematizes the identification of epistemic injustices more generally.[8] What constitutes an epistemic injustice becomes problematic if we have to look either to the sphere of discrimination, or to the sphere of distributive justice. In either case, the approach would appear ad hoc: we start trying to determine whether any epistemic good (such as credence) is denied on a discriminatory basis, or whether something of epistemic worth is unjustly distributed.

Doing so, we are hampered by lack of a standard by which we can judge whether the

---

7 Bernard Williams points to the requirement of there being such reciprocity in the practice of knowledge sharing, very notably in his reference to Constant, who held the truth is to be told only to those deserving it (Williams, 2002, p. 114).

8 This is so because current identifications of epistemic injustice proceed by analogy to testimonial injustice, or again by analogy to hermeneutical injustice, as these have been presented as the prime cases of epistemic injustice.

denial is *epistemically* fair or just, and equally a standard by which we can approve of the distribution of *epistemic* goods, or the granting or denial of access to epistemic goods. And that lack makes understanding how distributive and discriminatory cases are both instances of the same type of injustice harder, as we lack a standard that applies to both types of cases, too. As such we fail to understand what makes both testimonial injustice and distributive injustice an *injustice* or *wrong* to start with. What we need is a positive account of the moral-epistemic norm of justice that is being violated.

8. Epistemic Power as Central to Epistemic Justice

Epistemic justice, on my conception, is a positive notion in its own right. So, it is not reducible to the republican notion of justice as non-domination simply applied to the epistemic realm, as James Bohman (2012) casts it. Instead, epistemic justice as I defend it centrally revolves around the notion of epistemic power. Epistemic power is a subspecies of social power. If social power is "one's ability to influence the social world and to control/dominate others",[9] then epistemic power is one's ability to do so *epistemically*.

Epistemic power is, on my account, one's ability to exert epistemic influence, that is, one's ability to influence what comes to be regarded as knowledge and who qualifies as authoritative. It is the ability to influence what others believe, think and know. As it is a "power-to" it may be used to enable or disable others from exerting such influence, i.e. epistemically empowering or disempowering others. Some examples of exerting epistemic influence are: sharing one's knowledge that P, convincing someone that P, voicing an opinion that P, (publicly) questioning someone's belief that P, giving a source credence or refusing to do so – and all these at a relevant time.

My account assumes that an agent who exerts influence has the requisite standing or credentials to do so, at least such that he or she is believed. So, authority and creditworthiness are often prerequisites for exerting epistemic influence. Note, however, that the definition of "epistemic power" alludes to an "ability" which is broader than merely having epistemic authority. Even if having authority is an important aspect of having the ability to exert epistemic influence it is neither necessary nor sufficient.

Less than authority may be required, as one's reputation or status may not always

---

9 This deviates somewhat from Fricker's definition of social power (Fricker, 2007, p. 13). Notably, however, while Fricker regards the control of other's actions as central to her notion of social power (Fricker, 2007, p. 4), she writes at the outset that she begins from "the strongly intuitive idea that social power is a capacity we have as social agents to influence how things go in the social world." (Fricker, 2007, p. 9) This basic notion is what I will stick to, as it appears to me that something is lost if we focus solely on the idea that power can be used to control others' actions: it can also be used to make a difference in the social world. My conception of power is thus closer to Amy Allen's, who has argued, against the account given by Thomas Wartenberg (1990) that 'power-to' is a modality of power which is basic. Her reason, which I endorse, is that "power over" casts all modalities of power too readily as entailing domination, whereas there should be room, too, for (self)empowerment, and for resistance and solidarity, which requires the modality of "power-to" to be basic (Allen, 1999).

enter into one's being believed. Consider, for instance, a case whereby what one says immediately strikes everyone as the only possible true or good answer. The utterance causes a "eureka" response on the part of everyone who hears it even in spite of prior beliefs about, or disregard for, the person who said it. Here antecedent authority or credibility is not necessary for being believed and for exercising epistemic influence. On the other hand, even while authority or credence is often needed, it is even then not sufficient as (often) more than authority is required. That is, one's ability to exert epistemic influence can crucially depend on such things as: what other knowledge or beliefs someone has; and on whether or not one's message reaches the right audience; and the time at which one has the knowledge.

One may think that epistemic power is also had by those who have good connections, and as such enjoy easy access to knowledge. Indeed, people with a wealth of so-called "social capital", that is, good social connections, will have a fair amount of epistemic power. At the same time, to influence others and to have the knowledge and skills on which to base one's own decisions and social success, being connected to numerous influential people is not in itself sufficient as a ground for epistemic power – even if one may be perceived as epistemically powerful, and hence have *undue* or *excess* credibility. Indeed, then, being well-connected may increase one's status, and so it may help with being regarded as a knower or trustworthy recipient of knowledge and credibility. However, we seek (and ultimately should always seek) to assign epistemic power based on the qualities and knowledge that are had by the person, who tries to sway us, not their connections.

It may be thought that the first part of my account concerns the ability to exert epistemic influence in a way that focuses on the knowledge giver or opinion maker, rather than the recipient. However, it is complemented by, the second part of the account that focuses precisely on the side of the audience or recipient. Allocating epistemic authority or credence and their opposite, i.e. refusing to grant epistemic authority or credence or even discrediting someone, are also exercises of epistemic power. Even if one does not (necessarily) exert epistemic influence directly in such a case, one does influence the epistemic outcomes indirectly, by enhancing or decreasing another's ability to exert epistemic influence.[10] A hearer then has epistemic power to believe or not believe a speaker, and to publicly commend or discredit the speaker. The latter can be done either by denouncing a speaker's reputation, or by making use of one's own authority on a matter P, when the speaker attempts to exert epistemic influence.

However, there is another reason why the account encompasses hearers as well as speakers, and it relates to the due allocation of epistemic trust. The reason is that many cases of exerting *direct* epistemic influence are also cases of exerting *indirect* epistemic

---

10 "Epistemic outcomes" I define as what knowledge or beliefs people have, or their degrees of confidence in specific beliefs being true. Of course, one could argue that epistemic outcomes also include people's beliefs about specific others being epistemic experts in which case there is a direct influence.

influence. Consider the case of contradicting S's claim that P: doing so comes to publicly discrediting S as a reliable source or epistemic authority on P (and, potentially also as an authority regarding closely related matters). Also, the single act of sharing knowledge is a way of exerting both direct and indirect influence: one both influences another's beliefs and enables the other to share with yet others what he has come to know.

Notably, the decision to entrust someone with knowledge could be separated from the act of knowledge sharing; whether or not someone is deemed worthy of a given piece of knowledge, and/or trustworthy as a disseminator of it, is then seen as a prior decision. The judgment that is made concerns the allocation of epistemic creditworthiness. That this judgment is an exercise of epistemic power may be clear when we consider a case where someone is not deemed worthy of receiving knowledge. It shows that someone with knowledge has the (putative) power to exclude another from knowing that P as well. There can be good and bad uses of such power, but the point is that in the very act of deciding that someone is not deserving of trust, not worthy of receiving a (putative) piece of knowledge, one is exercising epistemic power too, and hence not only when giving or receiving knowledge. Not entrusting someone, epistemically, can be an epistemic injustice.

It should be noted that in both cases of exercising epistemic power directly and indirectly, one is risking as well as potentially enlarging one's own epistemic credentials, i.e. one's own epistemic reputation. One's reputation is an important measure of how much trust one deserves, and one's risking it is a measure both of how trustworthy one is and of how trustworthy another needs to be in order to be entrusted. The dynamic of trust in epistemic cooperation is one reason why justice in the epistemic realm is so important: relying on another's word is risky, but so is giving someone assurance, as this can backfire too.

Epistemic power in to be used with care and yet, more often than not it means inclusion and giving someone the benefit of the doubt. This is not a naive stance that the truth will prevail, in the long run, a view that Mill attacked in *On Liberty* (Mill, 1989, p. 31). For we cannot say these things without us ensuring that it does. Rather, it is to say that *because* truths can be distorted, and because our lived experiences shape what we know – as individuals – and shapes what knowledge is had within communities, that we must ensure all members can contribute in filtering out untruths and distortions and in completing the dimensions that were out of view.

One important question is whether epistemic power is always a matter of domination. I think it is very important that we understand that in many contexts, there is an implicit assumption that someone with power is someone who is dominant. Calling someone powerful means that they can dominate, if they chose to. At the same time having defined epistemic power as an ability – even if it is the ability to influence others, and hence to exert a form of social control – means that everyone could have it simultaneously, and that it is not necessary to dominate others. Indeed, using one's epistemic power well

means that one does not use it to dominate others: it should be used to exert influence where that influence is required.

Moreover, in principle it is possible for everyone to have equal amounts of epistemic power, in an ideal world where we all have the same abilities and all share what we know with others (in so far as that knowledge is relevant for them). More likely, in the real world there is no equal division of knowledge and credibility. Yet, we can ensure that everyone has some epistemic power and that knowledge and epistemic status are not abused, not exploited. This may require certain political changes, but it requires at any rate a shared ethos: an ethos of epistemic justice.

9. An Ideal-Theoretical Pluralist Framework

I have argued that by relating the concept of epistemic justice to epistemic power the former can avoid the charge of being an "ad hoc" re-labelling of cases that are either purely ethical or purely epistemic. In addition, drawing out this connection allows one to unify both "distributive epistemic injustices" as well as the "discriminatory" cases under a single concept. However, that does not imply that the concept of epistemic justice, while distinctive, cannot be further explained. It can be further explained as regulating the proper use and distribution of epistemic power.

This unitary explanation of epistemic justice can be formulated positively so as to provide a positive standard, an ideal, for epistemic justice. It provides a standard for evaluating both the morality of epistemic conduct and the epistemic justice of social structures. Epistemic justice is constituted by its own substantive norms that spell out what "proper use and proper allocation" of epistemic power come to. It is not merely a strategic label for redressing some certain kinds of injustice and wrongs by re-labelling them in a rhetorically persuasive way.

I will develop this argument further by showing that the notion of epistemic justice fits into a larger theoretical framework of justice and that the notion of epistemic power has its precedents. This larger framework is pluralistic in the sense that one can conceptually distinguish different kinds of justice such as economic justice, epistemic justice, political justice and legal justice. Examples of pluralist theories of justice are the theories of Michael Walzer and Iris Marion Young, but neither has been extended to epistemic justice (Walzer, 1983; Young, 1990).

Both theories are internally pluralist, which seemed to John Rawls to be a problem with any theory of this general form: they cannot give an account of the just, all things considered. I share Rawls's view that we ought to aim at an ideal conception of "justice overall", which will be achieved when all the various kinds of justice, which may conflict with one another, are optimized. While there is ample room for exploring how "justice

overall" may be achieved, starting from a pluralist ideal conception of justice that specifies the kinds of powers in play and justice with respect to each of these, will be the most helpful starting point.

A point that is important for understanding my approach is Nancy Fraser's argument that there is no contradiction between a theory being pluralist and a theory being ideal-theoretical (Fraser, 2007, p. 319). Moreover, there is no contradiction between striving for a comprehensive theory – i.e. a notion of justice overall, all things considered – and engaging with this from a perspective of critical theory, as Fraser has also argued (Fraser, 2007, pp. 319–320). A critical theoretical conception of justice is indeed the perspective that is most readily associated with a conception of justice that focuses on power. Internal pluralism is the result of the various forms of power that it regulates.

This is not to endorse Fraser's own ideal theory about justice. Fraser focuses on the structural arrangements needed for "participatory parity", which is an ideal according to which all can participate as equals in social life (Fraser, 2005, p. 42). For her, the concept of "participatory parity" as an ideal serves to unite three dimensions of justice: 1) recognition, 2) distribution, and 3) representation (Fraser, 2005, p. 42). These three dimensions of justice regulate what is needed to achieve participatory equality, respectively:

1. in terms of recognizing *social statuses*,
2. by distributing *rights and resources*, and
3. deciding on criteria for membership of the community that is bound by justice, i.e. questions of *representation*.

As she sees it, the three dimensions also correlate to different kinds of domains, namely: 1) recognition of status to the cultural; 2) distribution of rights and resources to the economic, and 3) membership of a community to the political (Fraser, 2005, pp. 42–43).

I do not share Fraser's outlook, as I view the dimensions she focuses on – recognition, distribution and representation – as perspectives on achieving justice. I say this, even as I acknowledge that these dimensions may depend on the kinds of power they regulate. The diverse forms of justice, as I envisage them, are picked out only indirectly. My direct focus is on the substantive forms of justice: economic justice, political justice, legal justice, epistemic justice. My account of the subcategories is thus more akin to that of Walzer's theory according to which there are different (distributive) *spheres* of justice (Walzer, 1983, p. 20). It is also similar to Iris Marion Young's view who made different forms of power central to identifying forms of oppression and domination (Young, 1990, pp. 32–33). The latter are the two conditions that define injustice on her view (Young, 1990, p. 37).

Yet, while Walzer's theory and Young's both focus on the different kinds of power, Walzer's purely distributive framing is somewhat misleading, while Young does not

aspire to provide an ideal theoretical account. I share Fraser's view that a theory of justice aims to find unity and comprehensiveness. Fraser notes that doing this from a critical theoretical perspective requires that comprehensive theorizing should be "infused …. with the spirit of fallibilism and openness to new evidence" (Fraser, 2007, p. 320). This is especially important given the fact that the concepts of knowledge and justice are both such that they may be appropriated by the socially powerful, since they are "essentially contested concepts" (see Gallie, 1955, p. 180).

Conclusion

In this chapter I have both introduced Miranda Fricker's pioneering work on epistemic injustice and set out the outlines of my own view. I have examined some of the main lines of criticism of Fricker's work and suggested responses to them. I have also set out a view that builds on Fricker's work and thereby goes beyond it. It places the idea of epistemic power at the center of an account that can unite both wrongs and injustices. At the same time, the idea of the use and allocation of epistemic power grounds the normative force of epistemic injustice as a sui generis category, in terms of its regulation being necessary for the social generation of knowledge, on the conception of knowledge which I have defended to this point. In particular, it casts light on my account of testimony as a generative source of knowledge. In the next chapter I will further develop this connection between epistemic justice and knowledge sharing.

Chapter VII | Epistemic Justice: Ethical Norms

Introduction

In the previous chapter I argued that epistemic justice is a subcategory of justice centering on epistemic power. What I seek to do in this chapter and the next, is to spell out the notion of epistemic justice further, delineating the proper use and allocation of epistemic power. In this chapter, I aim to render understood what norms epistemic justice entails for our individual conduct as participants in social epistemic practice. By focusing on ethical norms for conduct, my focus is mainly on the use of epistemic power. In the next chapter, I discuss how epistemic justice can be used for evaluating the social structures that shape our epistemic practices. There the focus is mainly on the allocation of epistemic power. Accordingly, in this chapter and the next I argue, respectively, that epistemic justice:

1. provides the rationale for moral-epistemic norms that relate to our (social) epistemic conduct, and
2. provides a way of articulating principles that can be used to evaluate the societal structures that shape our epistemic practices, and in particular the criteria we use for knowledge and credibility.

There is some overlap between the two chapters, while the first chapter aims primarily to delineate the proper use of epistemic power and the second focuses on delineating the proper allocation of epistemic power. That this overlap is unavoidable is because tending to the proper use of power presupposes an idea of its proper allocation: properly allocating epistemic power is one way of properly using epistemic power. Conversely, power being properly allocated helps ensure that power is used well, even if it is no guarantee of that. A proper allocation of epistemic power is such that the best use – or at least, the least abuse – will be made of epistemic power.

Despite this overlap, the use and allocation of epistemic power do each need separate attention. One reason is that it could be that even while people each use their epistemic power properly, the cumulative result, in terms of the allocation of a power, is one which is not just (Anderson, 2012, p. 164). Hence, the question of how epistemic power is allocated is related to, but still separate from, the question of how epistemic power

is used. A second reason is that, even as it may be tempting to provide a definition of the proper allocation of power in terms of epistemic power being correctly used, we should realize that epistemic power can be abused by those who have it, even if the allocation was reasonable. One way to see this is that epistemic power granted on the basis of one area of expertise is sometimes used to claim such power in another area; e.g. a professor in economic theory who, as a result of his title gets credibility, may comment on television on political issues in her society. If so, there is a good chance she is treated as having authority on such issues, even though she is not an expert in the field in question.

A third reason that allocation and use are to be viewed as separate issues that should not be defined in terms one another, is that the actual allocation of epistemic power is not only a function of it being conferred through epistemic acts, but also a function of the already existent social structures that tend to seek to replicate themselves (Fricker, 1998). This is important because the allocation of epistemic power reflects who can participate in epistemic practices and thus *who can shape the criteria* for knowledge and authority that serve to allocate such power. Those who have been allocated epistemic power have the power to include or exclude others.

Aside from attending to the cumulative effects of the use of epistemic power, then, in order to achieve a proper allocation of epistemic power we should pay attention to the structural features – i.e. criteria for knowledge and authority – of an epistemic practice and the wider societal context in which our epistemic practices are shaped; particularly, how other forms of power than epistemic power are distributed and used. These topics will be explored in the next and final chapter.

In this chapter, I first argue that epistemic justice is the rationale for the moral-epistemic norms that are relatively uncontroversial, such as the norms that Williams provided for accuracy and sincerity, and the norm of credibility which is key to Fricker's work, according to which one should "give credit where credit is due". I then argue that epistemic justice can also fill the gap that currently exists vis-à-vis what norm pertains to knowledge sharing. Specifically, I argue that the moral-epistemic default norm, based on epistemic justice, is that knowledge must be shared. As I argue in this chapter, framing questions of whether or not knowledge should be shared in terms of inclusion and exclusion makes us see that epistemic justice, when conceived as a universal ideal that should have real world application, requires an ethos of inclusivity and solidarity, and that sharing knowledge is a moral-epistemic norm because of this ethos.

2. Moral-Epistemic Norms that are Grounded in Epistemic Justice

The starting point, that I established in Chapter III, is that epistemic justice is the proper alignment of truthfulness and trustworthiness. The point of there being this kind of

alignment is given by the social nature and generation of knowledge, i.e. by the fact that such an alignment is needed for epistemic cooperation to be sustainable and successful. That is, epistemic justice is needed for social knowledge generative practices to function as they should. The "alignment of truthfulness and trustworthiness" is that there is a reciprocating in terms of truthfulness and trustworthiness the one for the other and to a like extent: we ought to be truthful when and in so far as others are trusting, and that we ought to be trusting when and to the extent that others are truthful.

On this view, when a practice is such that epistemically just conduct is achieved to a high degree in it, it means that there are moral-epistemic norms which are being observed which sustain the reciprocation of trust and thereby the cooperative nature of the practice. When epistemic justice characterizes a social epistemic practice such as testimony, it means that there is a high level of reciprocation of truthfulness and trustworthiness. In turn, this means that both *placing trust in others and being truthful* are rightly viewed as the *default positions* to take. For, in the circumstances whereby there is a high level of reciprocation, trusting others and being truthful are – exceptions aside – what the moral-epistemic norms based on epistemic justice require. The circumstances described here, where there are high levels of truthfulness and trustworthiness and whereby these are reciprocated, may be characterized as the situation of "the circumstances of normal trust" (Williams, 2002, p. 110). That situation is sustained by epistemic justice, i.e. by people's duly reciprocating.

However, we should note that epistemic justice as an ideal or standard for proper epistemic conduct (and the proper allocation of epistemic power) does not depend on there being high levels of trust and reciprocation. Epistemic justice does not equate to, or depend upon, "the circumstances of trust" obtaining. While we may not characterize an epistemic practice as just unless there is some high degree of the observance of the norms of epistemic justice, i.e. it being the case most of the time that epistemic power is used well and allocated properly, epistemic justice remains an ideal and continues to have normative force when the circumstances are not characterized by high levels of trust. Indeed, one may think that attending to epistemic justice, and instituting it, is especially important in precisely those circumstances.

In this section and the sections that follow it, I will argue that some of the moral-epistemic norms that we already recognize are the norms of epistemic justice, and that, indeed, they are the default norms. They are the default norms, I argue, because they apply when the normal circumstances of trust obtain. These circumstances do not always obtain, however, and when they do not a deviation of these norms is allowed, justified, or even required. The need to deviate from the default norms will be discussed in the final sections of this chapter; in part to delineate the limits of these norms, and in part to validate them as being rightly deemed the default norms.

Intuitively, the following two categories of norms can be explained by epistemic justice. In terms of the proper use of epistemic power, the two-fold rule by Williams

is that a respect for truthfulness implies (even) for academics that "they take care, and they do not lie" (Williams, 2002, p. 11). To take care includes, obviously, "to not cheat" in any way, but taking care is broader, as it also means one should not rest content with less than the full truth. It captures, then, that one should expend efforts at uncovering the truth and also correct others where they go wrong. This is what truthfulness comes to, in its most basic senses of Accuracy and Sincerity, respectively (Williams, 2002, p. 11). In terms of the proper allocation of epistemic power, there is the norm of credibility – pointed to by Fricker's work but also by Onora O'Neill – that "we should give credit where credit is due", or alternatively phrased, "trust the trustworthy" (Fricker, 1998/2011, pp. 60–61; Fricker, 2007, p. 47, 123; O'Neill, 2017, p. 405).

Why is the moral-epistemic norm of truthfulness as spelled out by the twofold rule "to take care, and to not lie" backed by epistemic justice? The answer is straightforward if we assume that Williams's notion of truthfulness is the broader notion of epistemic justice, which I argued in Chapter III, section 5. For in that case, the norms of truthfulness as explained by Williams's genealogy are norms of epistemic justice. What supports this reading is that the intuitively appropriate way to understand "truthfulness" is narrower than Williams suggests when he says that truthfulness can mean that one does not speak the truth. That is, when the norms of truthfulness are correctly captured by the narrow two-fold rule that one should take care and should not lie.

Truthfulness in this narrow sense captures one of the default norms of epistemic justice, as it guides the proper use of epistemic power in the context of knowledge generation and distribution. It is a default norm, because it is what epistemic justice requires, when the normal circumstances of trust obtain. In other cases, i.e. in the cases where we cannot assume that the circumstances of trust obtain or can assume that they do not obtain, we fall back on the principle of epistemic justice; for only that broader principle gives us leeway to consider what to do when the requisite levels of moral or epistemic trust are lacking, or how we can go about sanctioning an abuse of moral-epistemic trust with the aim of re-instating justice in the social generation of knowledge.

Indeed, I believe that we should cast Williams's genealogy as explaining that we need an ethos of epistemic justice, in order socially to generate knowledge. Saying this is not to deny Williams's central claim, which is that a functional epistemic practice requires that people are intrinsically motivated to care for uncovering truths and for speaking the truth. An intrinsic motivation is needed, as the genealogy explains, because other options than being truthful may be more tempting from a self-interested perspective. It is why truth must be intrinsically valued if epistemic practices are to be successful in generating knowledge. My claim, however, is that the normative force of the moral-epistemic requirement of truthfulness (in the narrow sense in which I argue it is to be understood) lies in its being the spelling out of epistemic justice in circumstances of trust.

Notably, however, truthfulness was defined, not only by the two-fold rule that Williams gave, but also as "a deep caring for (the value of) truth". That notion with its ethical valence that truth is what matters, and hence that we should strive for it in all sorts of circumstances, comes closer to the ideal of epistemic justice. But it does not equate with it or with the ethos of epistemic justice, for the reason that the ideal of epistemic justice also includes caring for justice – and more precisely, for justice in the generation of knowledge. A just use of epistemic power is about giving the right amount of credence to sources as well as about resisting wishful thinking and speaking the truth when one acts as an epistemic agent and as a putative source for others within the epistemic practice. That one should care for truth is a function of the fact that everyone in the practice is interdependently reliant upon one another. Because epistemic practices are social, truth-conduciveness is not our only nor our best option as a standard for evaluating epistemic practices: epistemic justice is broader and is a more suitable candidate.

I take it that it does not actually require much additional argument that epistemic justice also backs the norm of credibility by Miranda Fricker, which comes to the norm "to give credence where and to the extent that credence is due", or most simply put as O'Neill did: "to trust the trustworthy". This norm can be explained by the rationale of epistemic justice, namely that giving due credit is a matter of justice; not doing so is a harm which is underserved, i.e. an epistemic injustice, as Fricker has claimed, while the aim of doing so is epistemic as well as moral. To see this last point, it is important to note that even while the injustice of a credibility deficit may be suffered by an individual, it is also an epistemic loss for all those who partake in the practice themselves: by not adequately allocating epistemic power, they hamper the process of knowledge generation and end up with worse epistemic results than otherwise.

Neither one of the above norms speaks to whether or not we should be disposed, as a default, to share knowledge with others. Nor do other recognized, legal norms, such as the norms of freedom of (self) expression or of transparency. As O'Neill argues, the latter concerns "a requirement on information handling, rather than a requirement to communicate, or an entitlement to knowledge" (O'Neill, 2009, p. 173).

The difference between requirements on information handling and communicative speech acts, i.e. knowledge sharing, is even clearer with the norm of freedom of expression, as there are no rules that bind freedom of expression, where epistemic and moral norms ought to be in place. For freedom of (self) expression is not centrally about communicating truths to others, but exists to allow everyone to express themselves, as long as it does not harm others (O'Neill, 2009, p. 171). These rights, O'Neill notes, are therefore 'quite reasonably not restricted by duties to express oneself effectively, sensibly, accurately, or in any other specific way' (O'Neil, 2009, p. 171).

When it concerns claims to truths, and powerful parties who make it, O'Neill argues correctly that the legal norm that there should be freedom of expression is inadequate (O'Neill, 2009, pp. 169–170). There have to be restrictions on freedom of expression in

terms of the norms of truthfulness. Otherwise, anyone can say anything, and this is not the case when the generation of knowledge is at stake. As Bernard Williams said: 'in institutions that are expressly dedicated to finding out the truth, such as universities, research institutes and courts of law, speech is not at all unregulated' (Williams, 2002, p. 217). This point is one O'Neill cites and endorses, and applies to the media as one of these institutions (O'Neill, 2009, p. 172). Likewise, transparency is not sufficient for a (proactive) sharing of knowledge. For, as O'Neill notes, making information available, rather than hiding it, does not ensure others will find it or grasp its relevance (O'Neill, 2009, p. 172). She notes, too, that complying with transparency rules can be a way of ensuring that no attention is being paid to it, when hiding the information would draw attention to it (O'Neill, 2009, p. 170).

According to O'Neill what is required is that there are norms, both epistemic and ethical, that concern knowledge sharing, in terms of the content of what is said and who says it. That argument is wholly in line with my point that we need to attend to rules vis-à-vis knowledge sharing. Her argument serves to show that legal rules on transparency and on freedom of speech, do not yet give us sufficient handles on how, when and why knowledge should be shared – nor, more specifically, by whom, and with whom; and, importantly, when and why, and to what extent, censorship is appropriate. That she calls for ethical and epistemic norms, which may not usefully be classified as either ethical or epistemic as they tend to be both, is in line with my arguments so far (O'Neill, 2009, p. 176).

What is lacking in O'Neill's account, however, is an explanation concerning what makes for the urgency for why norms of knowledge sharing should exist. From O'Neill's discussion, it seems that the urgency is mainly or only political, e.g. that "the public life of democracies" is sustained by accurate reporting (O'Neill, 2009, p. 172). Though it is true that much of our politics is informed by what passes as knowledge, I have argued that the *need* for truthfulness is inherently epistemic as well as ethical.

The generative account of testimony explains why moral-epistemic norms vis-à-vis knowledge sharing are especially important: because we co-create knowledge, we must protect our epistemic practices against undue influence of the socially powerful. Undue social influence can have a direct bearing on our ability to acquire knowledge. This is not to deny that what passes as knowledge *also* has an impact on what issues are tackled by a society and how they are tackled. Indeed, as I see it, that is the reason why epistemic justice is a means to justice overall, as well as a part of it. However, the existence and point of moral-epistemic norms, as well as the importance of upholding them, is best understood when we realize that knowledge is socially produced. That is what I have argued throughout this thesis. In the below, I will turn to arguing that there is, indeed, a moral-epistemic default norm that pertains to knowledge sharing.

3. Knowledge Sharing as a Norm of Epistemic Justice

Can a (default) norm that knowledge must be shared be established on the basis of epistemic justice? We have seen that Williams has said that in its most primitive form, the ethical-epistemic disposition of Sincerity is openness, spontaneity. This is linked to his arguments about assertions being interpreted as true by those who hear them, and to our propensity to come out with truths, if we do not think about what we say but speak "unprompted" (Williams, 2002, p. 75). But the idea that sharing, rather than keeping information to oneself, is the natural way to use one's epistemic power is also supported by the idea that we cooperate to create knowledge.

Clearly, some of the reasons to share knowledge that stem from cooperation are available on an individualistic account of knowledge. This is because an epistemic individualist or a veritistic social epistemologist, who maintains that knowledge is true belief like Goldman, can make sense of the idea that people engage in epistemic cooperation (Goldman, 1999, p. 4). As I will make clear in the below, however, that approach does not yield a norm of knowledge sharing. Nor can it explain the moral-epistemic valence that I argue is the force of the norms that are backed by epistemic justice. In the below, I will argue that the norm of knowledge sharing exists, and that it has that dual moral-normative valence, while the rationale that backs it is that of epistemic justice. This norm, like the norms we saw above, is to be understood as a default norm, like the other two. It is the norm when the circumstances of trust obtain.

A first reason for sharing knowledge that may come to mind is that in cooperation, we share the fruits of cooperation. Knowledge is a desirable good, so we ought to share it with others, if we want others to share their knowledge with us. However, that line of reasoning does not yet presume that knowledge is created in the process of information being shared. It also does not provide a rationale for knowledge sharing that is specifically epistemic. These last two points are linked, so I believe, since it is hard to explain why knowledge ought to be shared if one fails to see that knowledge is socially generated whereby the sharing of putative knowledge is part of the cooperation required for knowledge in the first place.

That presumption is what makes for a second reason, which is that knowledge is shared because we need others to come by it. In sharing putative knowledge with others, claims to knowledge are validated or indeed rejected as such. On this view, supported by the generative account of testimony which explains why knowledge is a "steep" (social) good, knowledge sharing is not only a moral question, but is also an epistemic one. However, the idea that we share putative knowledge merely to use others as instruments of validation does not quite yield a moral valence to the reason for why knowledge should be shared. It does not yet explain why such a norm is moral-epistemic, i.e. a norm of epistemic justice, as opposed to merely instrumental to knowledge claims being validated. The reason would appear to remain purely epistemic, even if we add that we require validation from multiple different

viewpoints in order to be sure that no evidence is overlooked, and that there are no biases.

A third reason for sharing knowledge as a way of cooperation is so that time can be allocated better. Being actively helpful saves another time, and if we make sure work is divided and the outcomes shared, then it is not done twice our collective cooperative efforts are going to be more productive. Incidentally, saving another the effort of having to validate a truth for himself or herself, is one reason why testimony is used purely for transmission in some cases – of the form "take it from me". If we do not rely on others, we cannot work together and build on one another's results to further the results we have. Like the second reason, however, this reason for why knowledge sharing should occur also remains practical and epistemic. It is true that we can start to discern some inkling of the idea that we need to be and need others to be (pro)actively helpful if we are to socially generate knowledge. Yet, that knowledge is socially generated need not be assumed. The idea that we help others and that they (may) help us in turn does not yield a moral-epistemic norm that we should share knowledge. In so far as there is a moral obligation to reciprocate, it is based on prior agreements, which are wholly voluntary: others cooperate with us if they wish to do so. What is lacking, here, is the idea of widespread, unavoidable, mutual reliance, that pertains to the idea of knowledge being generated through our social epistemic practices. Instead, the assumption is that of people freely and autonomously contracting others with whom they wish to enter into a knowledge exchange, and any moral obligations follow from that – which obligations, however, are not specifically moral-epistemic.

A fourth reason for sharing knowledge is that in doing so we invest in our community having a common ground; a common frame of reference. Shared knowledge is required for shared understandings and is a way of consolidating social bonds. This reason, however, appears merely politically motivated. While it is true that shared understandings will facilitate communication and social bonds, knowledge sharing when done for this reason is not done for a reason which is moral-epistemic in nature. That is, it is not done from the viewpoint of ensuring the generation of knowledge. For if it were, then epistemic criteria for knowledge are leading. Here, the generation of shared understandings – of knowledge which everyone accepts as such – seems more important here than the idea of validating truths; and the requirement that knowledge should be true may fall away.

The fifth reason for knowledge sharing, which I think *is* moral-epistemic in nature, builds on some of the former reasons. For, while knowledge sharing may certainly serve the aim of inclusion, it is inclusion of a special kind. For it is inclusion which is motivated both morally and epistemically at once. The idea is that others are to be included because their viewpoints are needed to arrive at knowledge and because their viewpoints matter to us as they are morally equal participants in a shared social epistemic practice. This is to say, as I will argue in the below, that knowledge sharing is fundamentally viewed as an act of inclusion, which is proper on moral and on epistemic grounds: it ensures that we together arrive at knowledge that is deserving of its title, whereby we relate to one another as moral equals within the epistemic practice.

That knowledge sharing is an act of epistemic justice follows naturally from the intuitive definition that a proper use of epistemic power comes to the proper inclusion and balancing of all available sources. For other human beings are putative sources, i.e. putative contributors. By sharing knowledge, we ensure that they can act as sources, both by validating our knowledge claims and by sharing them with others. There are moral and an epistemic reason for why others are to be included in the generation of knowledge, then, and when we cast knowledge sharing in the light of inclusion, we see that it is rightly required that we share our knowledge with others, and let others partake in our practice too – at least in circumstances of trust.

I should stress, however, that the idea of inclusion, as it applies here, is not backed by an idea of "using people" as instruments, which more readily appears to be the case when knowledge is purely individualistic. Then the thought is, rather readily, that any inclusion of another proceeds on what we need from them, being instrumental to our own purpose. On that picture, we share knowledge with those whom we deem worthy – which casts some light on a quite different and higher standard for our deciding who else "needs to know". In contrast, what the idea of inclusion here defended captures is its opposite. We come to see knowledge sharing is a requirement of epistemic justice, as opposed to a privilege, in part because sharing is required when we take seriously another's standing as a moral equal, in epistemic matters as much as in other areas of life.

More generally, to understand why knowledge sharing is a default norm that is (indeed) required by epistemic justice, we need to understand *both* that there is an important ethical side to knowledge generation (as Williams argued) *and* that there is an epistemic side to the social generation of knowledge which is a matter of ethical concern. The ethical side is that knowledge generation requires that we sustain trusting social relations. In order to do so, when already assuming the circumstances of trust, we should show at a minimum civility and respect. So, we should be willing to share knowledge, deeming others worthwhile and placing trust in them as a default. The epistemic side to knowledge sharing, which admittedly has a moral valence too, is that it is not possible to have an epistemically just community when truth, as a criterion, is discarded; e.g. when we merely aim to come to shared views no matter their epistemic value, or believe sources no matter what they say. Indeed, I think it is fair to regard knowledge sharing as fundamentally a form of trusting others, namely as a way of epistemically empowering them through an act of inclusivity whereby they are included into "a community of knowledge" (Kusch, 2002, p. 60; Welbourne, 1986, p. 1). [1]

That sharing knowledge is an act of inclusion into a circle of knowers is the central idea of Welbourne's thought that testimonial exchanges help to build a community of knowledge. This is a community within which people have come to share a body of

---

[1] In fact, Welbourne considered sharing knowledge is a way of creating a community of knowledge. However, in so far as the knowledge "pre-exists", such that it is being "shared as knowledge" it is assumed that there is already a community of knowledge. In that case, the way to view knowledge sharing is as a way of enlarging the community. This idea, that knowledge sharing quite often means enlarging the community of knowledge as opposed to creating one, has also been defended by Kusch (Kusch, 2002, p. 61).

propositions as common knowledge. I have also noted his Davidsonian point that, in the very process, they come to accept a shared world: a framework of agreement that constrains their disagreements and rules out inexplicable disagreements (Welbourne, 1986, p. 81; Davidson, 1973). What I want to add to Welbourne's arguments is that those who are included in the community of knowledge thereby have epistemic power which outsiders lack. Of course, in order to move from outsider to insider, one must be judged trustworthy and then be entrusted with knowledge. When both have occurred, one has become an insider, "a knower", where this status confers upon one epistemic power.[2]

With respect to insiders, this power consists in being able to exercise epistemic influence: one can epistemically engage with the propositions which one is told are true, as well as include relevant others into the circle – in a real sense, one has become an "equal" in epistemic power to others who know that P too (in respect of knowing that p). With respect to outsiders, one has acquired a new epistemic status or given that they may not be aware of the existence of the knowledge or the privileged community, at least factually a position which outsiders lack.

Knowledge sharing proactively and openly then seems to be what epistemic justice requires. As fellow humans we should give others – at least when we consider them as moral equals – the chance of benefiting from knowledge and the opportunity to bring to bear their perspective on its appraisal. Here, as I emphasized before, the idea is that the norm to share knowledge is the default: it prescribes that knowledge is to be shared in so far as the normal circumstances are circumstances of trust.

Notably, the way the requirement of knowledge sharing has been presented, as a default requirement, goes beyond it being a matter of benevolence. This is because, on a social conception of knowledge as I have defended it, in order for there to be knowledge, we require others to be cooperative members of our species just as we are reciprocally regarded as such. Of course, there may be considerations that speak against knowledge sharing too, and we should consider these, in order to validate the idea that the default is that others are to be included on most issues – or at the very least the issues that may affect them and to the extent that these others may reasonably expect such efforts to be not too exigent.

4. Limits to the Default Norm of Knowledge Sharing

Arguably, the best use of epistemic power is one that seeks to allocate it equally, such that epistemic power is not monopolized but diffused. One reason is that this leads to

---

2 This story fits with Craig's genealogy of the concept of knowledge, whereby being a knower is first and foremost a status of social significance (Craig, 1990).

lesser abuses of power, i.e. the power being used outside its proper sphere (see Walzer, 1983, p. 19). Another reason is that, when more people are empowered, there is more chance that all relevant viewpoints are brought to bear, which is important to filter for undue biases (Longino, 1995, p. 384). However, as we will discuss in the next chapter, there must be some limits, some epistemic criteria, in place too. For not everything that passes as knowledge is knowledge, and we have a responsibility, too, to be critical in accepting knowledge and in accepting epistemic authority. In this section, we will look at what exceptions exist to the default norm that knowledge should be shared: is it really so that an individual should share knowledge in circumstances of trust? And if so, to what extent?

A first thing to say is that, in the case of knowledge, it is likely that I would not abuse my epistemic power if I decide to share it. For in a case of knowledge, I am not misleading the recipient, or misinforming them. This could be the case, however, if I personally doubt that the piece that is in my possession is knowledge and yet present it as such – without expressing that personal doubt. Hinchman notes that it is possible to carefully avoid or "hedge" claims which are not to be seen as acts of telling (Hinchman, 2005, p. 566). He uses this in support of his view that acts of telling are those whereby a speaker may be assumed to give assurance to the hearer, even if it is done implicitly, through the illocutionary force of the speech act in question; for it is possible to say "Do not take my word for it" or "I am not sure but" when presenting another with putative knowledge.

Indeed, one can avoid telling another something, too, by merely asserting information that make them draw their own conclusion, in which case one is not responsible for having told them (Hinchman, 2005, pp. 570–571). One may, indeed, exploit Gricean principles of conversational implicature. Consider this snippet of dialogue: person A says to B "how many of these tablets is it safe for me to take?" B replies "I took three and they did me no harm". Has B answered A's question? Did B do so in a way that undertook assurance to A that "three" is the correct answer?

Even so, we may think that a case of knowledge sharing too could be an abuse of epistemic power, in a number of ways. First, it could be an abuse of epistemic power if I (purposely) provide the knowledge in order to distract my interlocutor from gaining access to *other* knowledge, or knowledge of some other kind. In that last respect, one can think of knowledge that is particularly discouraging for the interlocutor, or knowledge that aggrandizes an aspect of a problem (relative to expected benefits), or knowledge that (re)frames the inquiry at hand.

Second, it could be an abuse of my epistemic power if the knowledge I offer is offered for *other ends than epistemic ones* (Williams, 2002; Goldman, 1999). Here it may be thought that the abuse will impact, mainly, on the distribution of some other form of power, rather than epistemic power. For through my use of epistemic power, I may illegitimately acquire political power or economic power, for instance. Often, however,

in cases where epistemic power is used to further other ends than epistemic ones, there is also something epistemically questionable bearing on the case. This is so because the proper use of epistemic power is motivated by furthering knowledge of truths, so other kinds of motivations will likely have an epistemically adverse effect.

One may think, for instance, that a father sharing expert knowledge with his son, so that the son is best placed to be picked for a certain office, is not at all epistemically problematic. However, that kind of privileging may backfire. For the person who received inside knowledge will give a false impression at the job application about his actual abilities. Even as he may be hired, it may prove to be a wrong decision of the selection committee, such that he finds himself out of place, while a more suitable candidate was overlooked.

Third, there may be examples that straightforwardly combine the two elements: they are epistemically problematic and also impact on other power. For instance, if the intended aim of sharing knowledge is to change the subject that is debated or inquired into, or if someone deliberately ignore aspects of an issue in order that the way of life she enjoys goes unchallenged.

Many people will doubt that all knowledge should be shared, and rightly so. I will look at three reasons that may speak against sharing, in turn[3]:

1. privacy and other considerations against sharing, such as secrecy;
2. relevance;
3. exigency.

Privacy and other consideration against sharing knowledge, such as secrecy. For legitimate reasons, knowledge that pertains to persons' private sphere is a special case: often it is not to be shared. A private sphere is needed, to protect freedom of speech and of thought (Mill, 1898, p. 15). The idea here is that we should be free to test our ideas, even bad ones, without censorship or (social) repercussions. Related to this is the protection of the private sphere is the point that some knowledge about a person can impair their social status.

There are many things which others simply should not know, because such knowledge can be used against them, or even used to (epistemically) dominate them. It is especially clear that knowledge concerning another person should not be shared when it may limit a person's freedom or hamper their development, or deflate their self-image or self-confidence.

An exception is that we may need to share knowledge of a person's conduct that rightly undermines their undue reputation as a knower; it should be made clear when

---

3 Efficiency may also be thought to be a good reason. It will be discussed in the next section, where I consider the objection that knowledge sharing should proceed on a need-to-know basis.

people speak without caring to tell the truth. Yet, great caution is required, as it is an injustice, too, to incorrectly undermine the epistemic authority that a person has and should enjoy.

In general: sharing knowledge should aim to empower others, not undermine their personal confidence. It is quite difficult, in some cases, to ensure that one's motives are pure and to know what to do, given that the consequences may be multifarious. It will be possible to make mistakes. There is some room for discretion then, for there being no obvious one right thing to do. At the same time, the default helps one to understand that one's responsibility to others is the main consideration: one should share knowledge *except where* doing so undermines other people's self-confidence and their ability to be a source of knowledge. In cases of doubt, when a person seems a little arrogant say, such that others epistemically suffer from it, or such that the person is hampered in their self-development, it will be best to simply ask the putative recipient what knowledge, or what criticism, they are comfortable with, and to communicate in a way that makes clear what intentions are served.

Relevance. It may be claimed that not all that we can communicate ought to be communicated: that, in fact, relevance is a criterion for what is to be considered worth knowing. Yet, it is worth noting that while relevance is indeed important in communication, defining what makes for it proves elusive. While Grice considered the concept of crucial importance, he was unable to define it (Grice, 1989, p. 27). As Longino argues, relevance is context-dependent, even when it concerns the question what scientific theories or models are empirically adequate to furnish us with knowledge. For it is always contestable what data – or even whose data – are relevant for the acceptance of a theory or model (Longino, 1995, p. 394). Moreover, which additional criteria are called for, to determine what theory is to be accepted, depend on the socio-political context (Longino, 1995, pp. 396–397).

The elusiveness of relevance then, as noted by Grice, and its standing in need of further interpretation (even in the case of what qualifies as knowledge), as Longino argued, make it problematic for relevance to be a criterion for knowledge sharing. But if relevance is relevant, how do we proceed? What should or should not be considered worth sharing, worth bothering another with?

The answer is that by lack of a criterion for relevance, it might be best to check with the other, what they seek to know, to start a conversation; maybe one that appears to lead nowhere at first – as Grice suggests. That allows one to find out what is relevant and to ask another to help verify the relevance. It is to be open to questions or suggestions, even while one is also trying to be helpful.

Alternatively, one may be open by allowing that relevance is not a limit on sharing: that things can become relevant, if only the other person learns more. Typically, schooling proceeds by offering knowledge that might give one the handles for inquiring further

into issues when and to the extent that they have sparked interest or have become because one's circumstances or life goals. Some of what we learn is not relevant, except that it provides a mental training, and that we are capable of learning and remembering difficult things, and making the needed connections, gives us the confidence we need.

But, to argue that relevance is a criterion that can be used instead of others, such as that of being epistemically just, seems misconceived. At most one may point at the fact that some ways of speaking are clearer in certain contexts. Indeed, this is what Grice made clear: that we do expect others to say those things that the context makes it relevant for them to say, and that we may draw the wrong conclusion when information is being offered which is not particularly relevant (Grice, 1989, pp. 30–31).

On my view, relevance cannot usefully serve us as a criterion for delineating, and limiting, what knowledge is to be shared. Even so, relevance does pertain to questions of effective communication given the fact that any linguistic act requires interpretation, whereby relevance is presupposed. Indeed, one's way of communicating may well be subject to scrutiny from a viewpoint of epistemic justice. For it is true that if one seeks to mislead or deceive or if one is not making enough effort to be understood, then one is abusing one's epistemic power. A norm of epistemic justice, then, could well be that sharing knowledge requires that one communicates in a way that enables the transfer of knowledge (O'Neil, 2009, pp. 177–178). That also includes, I would think, that one is open to understanding what another person seeks to know, i.e. what another person finds relevant.

Exigency. The point that the moral norm is that knowledge must be shared, rather than not, should not be misunderstood as a norm demanding that one shares all of one's knowledge, all day long. Nor is it the idea that one must share knowledge without asking for a reimbursement of one's (reasonable) costs and efforts. Instead, the default norm requires that one gives others access where doing so is possible, and that one does not keep knowledge from others for no good reason. As we saw in the above, one can have some good reasons not to share knowledge. One is not required to ensure others know everything that one oneself knows or does, but, instead, that one acts as a good, reliable, source, and shares one's knowledge when others indicate they seek it. Quite notably, knowledge distribution is not the sole responsibility of one individual, as the recipient has a role, and there are others who may share the knowledge requested, too – to whom one may also defer.

5. Need-Based Knowledge Sharing?

A rival criterion to the default that knowledge be shared is the default that knowledge not be shared unless our interlocutor "needs to know it". This competing need-to-know

basis for knowledge sharing is arguably an approach whereby knowledge is more readily not shared than shared. There are two reasons for assuming this criterion. One reason is efficiency. If efficiency guides us, we put an emphasis on informing people only based on whether the information is *directly* relevant for them and/or whether, when we give the information, it will yield us some *direct* benefit. This is efficient from the viewpoint of someone who must make some effort (or risk their reputation) by giving knowledge (or even in deciding to give it).

It is also efficient from the perspective from the hearer, it may be thought, because being informed about facts that do not have much direct relevance is burdensome and makes one responsible epistemically without there being any good, immediate or apparent social cause. There are lots of things people do not want to know, simply because it overwhelms them to know too much, or because they will have to remember it while they are given some (implicit) responsibility to use and share that knowledge at some relevant time and with some relevant others (who need to know it). And besides, being informed of things takes up time.

A second reason why a need-to-know basis defies knowledge sharing as a default norm is that it assumes that knowledge is rightly held by those who have it, and hence that it is not a common good. Moreover, the decision as to what is relevant for whom is made by those who have it. That means that the powerful can evaluate whether they should extend their powers, without it being morally required. The need-to-know basis also allows for all forms of social bias to enter the assessment of whom else is deserving to know. In fact, the idea of deservingness does not so much as enter the stage of a decision to share or not share knowledge: the question of who needs to know may be determined solely by a self-interested calculation as to who is the most "useful" person to involve, or who cannot without risking one's own status be ignored or passed over – such that structures of other forms of social power replicate at the epistemic level, in terms of epistemic power had by members of a community.

Notably, the idea of a *need-based distribution* has been found to be problematic, quite independently of the subject matter, because "needs" tend to be defined by the powerful, and to work to the disadvantage of subordinated groups, who are dependent on the powerful in the process, i.e. "needy" (Fraser, 1989, p. 169). Talk of needs also avoids any questioning of the entitlement of those with power and of the idea that others should have (equal) rights (Fraser, 1989, p. 182). This is an argument that I find important to mention, even as I will not go into it. It does suggest that conceiving of knowledge sharing on a need-to-know basis may actually be a way for those who have knowledge, and hence have greater epistemic power than others, to justify their having greater epistemic power in a way that allows them to retain it – even to those who might want to have it. The reasoning would be that there is no need for others to have the knowledge in question too (whereby the ones with the power define that need). Moreover, it is self-justifying in the sense that it allows one to keep one's knowledge to oneself, or share it

only with others as a special privilege, while not recognizing that having power which others lack constitutes a form of dominance. It makes it seem that one is entitled to be another's (moral-)epistemic superior. Part of what feeds the idea of entitlement is the idea that inquirers are epistemically autonomous, as opposed to being interdependent on others.

Of course, efficiency itself – depending on what goal is sought – may not actually speak against knowledge sharing. For in many cases, efficiency will speak in favor of it. What efficiency requires depends on what goal is pursued. If the goal is the least effort for the person who has the knowledge, then it may be efficient to not share knowledge. But that can backfire in the long run, when people stop making an effort for you, to involve you or let you know what is going on. Even from a perspective of rational self-interest, there may be reasons to share knowledge (Goldman, 2001, p. 360).

For example, in a business environment, it appears that knowledge sharing is one of the key features that makes for the success of the business. This is so even if workers may be set to compete in such ways that they fear sharing what they know, as others could then, too easily, get by without them. One role of management, besides the exercise of discretionary authority, is to ensure a flow of knowledge to enable senior management to make decisions. Yet, precisely because of their position of power, a concern with the idea of informing others on a need-to-know basis is that (middle) managers, in fact, block the knowledge flow, despite their putative role in facilitating it, to sustain their own position.

A point to note is that the need-to-know basis, as I have described it, derives its main point from an idea of knowledge that is static and private. The ideas of efficiency and of relevance, then, which are embodied in the idea of a "need to know" in order to consider sharing knowledge, are coupled with a starting assumption that knowledge is private property, whereby the individual who possesses it may decide what it be used and useful for. He or she, who is deemed the proprietor of the knowledge has the capacity/right to decide what to do with it. He or she is free (i.e. does not have any obligations, duties or binding moral reasons) to share it or to keep it for himself/herself. Keeping it to himself/herself is deemed perfectly defensible if doing so furthers his or her individual goals or is simply the least amount of effort required as it pre-empts having to even think about the needs of others (if one pushes the line that the need to know must be apparent to one, in addition to its existing to start with).

To conclude this section: the idea of a need to know basis is an intuitive criterion for knowledge sharing: on the face of it, it makes sense. Time is limited, and we better not waste other's time, or our own. However, the default it sets is one whereby the possessor of knowledge can freely decide whom to share knowledge with and when, and where that judgement is based on their personal assessment of relevance, which cannot be faulted on moral grounds. Because, as a principle, the putative receiver has no righteous/valid claim, receiving knowledge is viewed as a privilege.

As I have argued, then, the criterion of informing someone on a need-to-know basis is a criterion that is heavily tainted by epistemic individualism, i.e. by the assumption that individuals are epistemically self-sufficient. It is such epistemic individualism that I have resisted throughout this thesis. This "need to know" basis might be voluntarily adopted, by agreement between all relevant parties, in specific circumstances. However, it ought not to be the default. The essential commonability of knowledge puts the onus of justification on those who want to limit access to knowledge.

It is by rejecting the idea that knowledge is first and foremost an individual's property, that we can see why we should care that others may want to know, and why we should care that others know and see to it that they do – for both moral and epistemic reasons. It is important to note that what the need-to-know basis makes us overlook are the reasons why we care that others know. The reasons for sharing, so I argue, and for why we should be fair when we decide on whether to share knowledge, are given by the moral-epistemic norms that exist *because knowledge is socially generated*. Without others' help, their teaching and supporting us, we would not know a thing in the first place. This is why Welbourne stressed that within a community there should be reciprocity (Welbourne, 1986, p. 67). Without a default of knowledge sharing, there is no room for sustaining epistemic justice and furthering knowledge. For doing so requires, both from a moral and from an epistemic perspective, solidarity and inclusivity, whereby openness, or "spontaneity", as a default is key to epistemic cooperation.

6. Instituting Trust: Inclusivity and Solidarity

Up to this point I have discussed norms that appear to be candidates for giving substance to the content of a theory of epistemic justice. I have argued for three norms and have explained of the first two how epistemic justice is their rationale; in the case of the third, so far, I have assumed it. The notion of epistemic justice is abstract. My hope is that it is less abstract, now that I have explained three moral-epistemic norms that give it content. Something I have not yet argued for is why or how epistemic justice grounds the norm of knowledge sharing. Why should inclusion in the epistemic community, which I have argued is the rationale of knowledge sharing, be a demand of epistemic justice? Part of that answer will be provided in the next chapter, where I discuss how we can arrive at proper criteria for knowledge and credibility. But part of the answer is that on my view, epistemic justice is not only a political notion, but also an ethical one. That is, there is an ethos of epistemic justice that exists to sustain and to promote epistemic justice.

One function of this ethos is to ensure that one does what one is expected to do as a cooperative member of the species. But another function is to ensure others are enabled and motivated to further epistemic justice, too. The norm of knowledge sharing

then, is to be cast as form of empowerment *and* as a form of instituting trust in a way that strengthens the epistemic community. In turn, if the trust is reciprocated by the involvement of the new member, there is a stronger epistemic basis for the knowledge the group holds in common.

Strikingly, the idea of empowerment is one which requires an understanding of one's act of knowledge sharing as an act of solidarity. I think this is because we view empowerment as not a strict requirement of justice, but rather as an act that builds, fosters or grounds and enlarges a community. I believe there are reasons of solidarity, then, for including a person in a circle of knowers, even when it may be thought "unnecessary" from a perspective of individual rationality as dictated by efficiency or from a (more socially inspired) "need-to-know" basis.

What is solidarity, and why is it relevant? As I see it, solidarity is prior to empowerment. While empowerment assumes an asymmetry, which is overcome through empowerment, solidarity is the idea that we are essentially similar, all of us. It allows one to treat others as if they are equals, even if our circumstances differ, and to act in concord to achieve a common goal.

Solidarity is the third of a three-part slogan used by the French revolutionaries. But, while liberty and equality have received much attention, solidarity not as much (Carrabregu, 2016, p. 507). A putative reason is that deontological accounts of ethics fit into the rationalist individualistic approach that has dominated most of the sciences and the humanities, an approach whereby rights are central to justice, and where these rights – as well as the correlative obligations shaped by them – are conceived as pertaining to individuals (Carrabregu, 2016, p. 507). This excludes moral-epistemic norms that appeal to one as a member of a group or community, whereby the actions taken require reliance on such a group or community.

Another reason that framework excludes solidarity is that solidarity presupposes actions not only *by*, but also *for the sake of* a group or a community (Carrabregu, 2016, p. 507). In the individualist rationalist tradition, with its universal principles, any actions that are not either self-directed or universal in spirit are supererogatory (Carrabregu, 2016, p. 507).

In feminist literature, however, there is an insistence on the importance of attending to the power and interests of groups. Political philosophers Young and Fraser have argued that theorists of justice should look at how power plays out between groups, given that often vested interests are secured for and by a group, to the disadvantage of another group. The focus on groups is a natural one for those who are concerned with domination, because only a group can affect social changes when the changes require new social structures. One of the insights yielded by seeing social groups as important units of political theory is that the focus on groups that is central to this way of understanding justice (as an ideal of non-domination) creates room for understanding the dynamic nature of power, and hence for the notions of empowerment and solidarity,

alongside that of domination (Allen, 1999, p. 56).

One of the first to write about solidarity from within a deontological conception of justice is Habermas (1993). As Gent Carrabregu remarks Habermas apparently aimed to answer criticisms of the existent rights-based conception of justice by feminists and neo-Hegelian critics (Carrabregu, 2016, p. 507). Feminists demanded a notion of justice that would give room to the ethics of care of others; neo-Hegelians sought a notion which would do justice to "the sociocultural rootedness of morality" (Carrabregu, 2016, p. 508). To answer that demand, Habermas created room for solidarity with his insight that solidarity naturally accompanies the rights-based conception of justice as its counterpart, or in his words, as its "reverse side":

> From the perspective of communication theory there emerges instead a close connection between concern for the welfare of one's fellow man and interests in the general welfare: the identity of the group is reproduced through intact relationships of mutual recognition. Thus, the perspective complementing that of equal treatment of individuals is not benevolence but solidarity. *This principle is rooted in the realization that each person must take responsibility for the other, because as consociates all must have an interest in the integrity of their shared life context in the same way*. Justice conceived deontologically requires solidarity as its reverse side. It is a question not so much of two moments that supplement each another as of two aspects of the same thing. (Habermas, 1993, p. 98 – my italics)

We, as members of a community, must stand up for the rights of others, if their rights, *which are also our rights*, are to be protected. This is the quintessence of solidarity, that someone's fate in a social setting is equal to, i.e. "stands for" that of anyone else. This idea of solidarity also secures its place as a natural corollary of rights. For it is necessary for there being any rights at all.

Without solidarity, no rights are upheld, as rights would then have a merely instrumental existence. The function of rights is to ensure that an individual's most vital of interests is secured through their being enforced on our behalf (Mill, 1910, pp. 50–51; James, 2005, pp. 149–150). Rights are a social call upon others to respect the autonomous space of those whose basic interests they secure; when the object of some right is threatened, invoking it is to call on third parties to intervene when necessary.

But such intervention presupposes solidarity on the part of others, a common cause, a social context of "a common life context" which yields the solidarity of others – e.g. other members of society, or other members of humanity. This also means that it is compatible with Mill's notion of rights in his *Utilitarianism*, which are the moral obligation of *justice,* in the sense that 'to have a right is to have something which society ought to defend me in the possession of' (Mill, 1910, p. 50). A right exists if and when an individual has a claim that society protect his possession of it, and hence a valid demand that those who violate it are sanctioned. Rights exist, according to Mill, to safeguard an individual's security, which is a first condition – alongside nutrition, he notes – to securing a person's livelihood (Mill, 1910, p. 50).

Mill rejects the claim that all moral obligations fall under the heading of justice (Mill, 1910, p. 47). On Mill's view, solidarity is not a moral obligation of justice, (in so far) as it is not directed to any one person and in so far as obligations would appear to attach to individuals as opposed to collectives, on his view in the day. At the same time, an ethical disposition towards solidarity is required to sustain justice. Thus, as I see it, Mill's view of rights is compatible with the idea that there is to be a shared ethos of solidarity that sustains justice that includes the moral obligation to ensure that justice is upheld, in the sense, at least, that rights are effected (Mill, 1910, pp. 46–47, pp. 49–50).

An ethos of epistemic justice, in turn, is one that focuses on ensuring that knowledge and the criteria for it and for credibility are sustained, through a deep, intrinsic valuing of truthfulness, and a concurrent investment in social relations that further this, including the sanctioning breaches of truthfulness and of trustworthiness where needed and defending the system against being undermined by unfounded attacks.

This idea of solidarity as a corollary to rights and as needed for upholding justice, whereby justice is understood as essentially "non-domination", fits with the understanding of solidarity defined by Allen as "the ability of a collectivity to act together for the agreed-upon end of challenging, subverting, and ultimately, overturning a system of domination" (Allen, 1999, p. 127). Solidarity is a positive imperfect duty to act against injustice, and is the natural companion of the right that one not be treated unjustly.

In her work, Allen has argued that solidarity is a modality of power, and specifically, of "power with" (Allen, 1999, p. 126). Her work has made clear that a conception of justice that centers on power has room for the notion of resistance and (self) empowerment, alongside domination. Yet, I do not fully endorse the idea that solidarity is an ability, or a power. Supporting this point there is the fact that we have no verb that matches the noun solidarity, in the way that we have the verbs "empower" and "resist" that match the nouns "empowerment" or "resistance". This might be a result of a society focusing on individuals as the unit of action, but as things stand, solidarity is associated less with (collective) action than with (individual) intentions and dispositions to act in accord with others. Thus, I would think that solidarity is best understood as expressed in or through individual actions, for instance by exercising one's power such that one *empowers* or *includes* others, or *resists* injustice on their behalf. Thus, I see solidarity as a precondition of acting in accord, rather than the power of doing so or of collective action. This is so even if it has to be recognized that collective action is a power, as it can bring societal changes.

In my view, solidarity gives content to the ethos required for including and trusting others who are not already part of some pre-given community. I regard solidarity, then, as the ethos of (social) justice, which ethos is required for people to uphold the rights of others, identifying with their humanity, and to ensure that these rights are socially effectuated. Solidarity contrasts with mutual respect or mutual recognition in that it is not necessarily exercised within a set relationship as it can *serve to create the grounds for*

it. That is, I think solidarity does not presuppose a community. It can be used inside a community to promote justice, but can also serve justice by including others who have previously been excluded for no good reason, i.e. to enlarge one's community. That is, it can fit with an ideal of justice that is universalist in nature, and is an ethos that we require to institute epistemic justice and inspire trust in others.

To conclude, solidarity is essentially the recognition of connectedness that makes for relations and communities, and for enlarging these. It also contrasts with fairness, in so far as fairness involves that one believes there is a "cooperative venture" or "scheme" that is going to repay one. In the case of solidarity, there is no such scheme, but only a common cause. So, it is believed, instead, that one morally has reasons for making an effort, and that these efforts will, if all goes well, be supported and paid forward. An act of solidarity is best seen then as a social investment, whereby there is no expectation of a return by specific others acting in a like manner, even if there is an implicit request that others will join one's cause because it is the common cause, e.g. in the name of justice, and hence act out of solidarity or in solidarity (simply to show support), too.

7. The Epistemic Responsibility of Knowledge Recipients

I have emphasized in chapter V that on a social understanding of knowledge whereby a generative account of testimony is central to it, the knowledge giver and knowledge receiver share responsibility for the epistemic state that they end up with vis-à-vis the matter in question. This section spells out what epistemic justice requires on the side of a putative recipient of knowledge. As may be expected, the expectations on a receiver are not for them to be passive.

Recipients should recognize and reciprocate, in so far as possible, the efforts extended to them that enable them to be included in the epistemic process. This means that in a case of testimony the recipient must be critical in his or her acceptance. Even so, it is also important to emphasize that a recipient has to value being entrusted to start with, by listening in a constructive and open manner, even before they engage with the content: they should be open-minded. They should value and recognize the speaker as a putative source and make clear that they interpret a communicative effort as a speech act with the particular intention invested in it.

Second, if the piece of knowledge is accepted, it also means that the recipient who accepts a given piece of knowledge as such acquires the responsibilities that the previous knower had, when it comes to sharing (even if sometimes there is an added responsibility of ensuring that the original source remains known as well). In the special cases, where the knowledge has been hard-gained, or shared with dire consequences for the sharer, the position of the recipient will be one where they should extend moral-epistemic

leniency in their judgement to those who have suffered or will suffer the consequences of sharing. To illustrate this, I will discuss two cases: whistle blowers and those who come forward with allegations of sexual assault.

One way to respect the personal sacrifice of whistle-blowers is to assume their good intentions and to disregard the inevitable attacks on their personality driven by those who have political power and/or economic stakes. Because of the personal risk carried by a whistle blower, it is especially important to be open-minded and to regard the whistle-blower as a putative source. This does not mean that one cannot be skeptical of the way whistle-blowers choose to use their epistemic power. It may indeed be the case that they try to gain influence in politics – disagreeing with and attempting to change decisions and policies that were not theirs to redesign. It is also true that whistle-blowing itself (as long as it has not been recognized as actual whistle blowing) is at least prima facie an abuse of epistemic power, as it will be classed that way by those who have the power to set the rules on the use and necessary confidentiality of the knowledge in question. These rules, however, could well be epistemically unjust themselves, and even if the rules are not, there may be good reasons to defy what the prevailing epistemic norms demand in light of weightier considerations, given that the issues brought out cannot be assessed for their injustice in any other way, since they would have remained unknown. In a case of whistle-blowing, then, it seems especially important that the public, or those who are addressed, retain an open mind. They should consider that a genuine whistle-blower seeks to address the abuse of power (which is non-epistemic in nature; political or military or economic, though an abuse of epistemic power may be an additional problem, too) and that their sharing of the knowledge in question stems from a moral conviction that what they have experienced or found out needs to be known.

The second case is that of victims of sexual assault. In recent years, we have witnessed that more moral-epistemic respect is being extended to victims of sexual assault who testify in court or outside court. We have come to understand that secondary victimization, in the form of not being believed even despite one's best epistemic efforts to relate one's story is a moral wrong, too – adding insult to injury. As such, it has become a commonplace that we should extend to victims' testimony credibility as a truthful source, as a moral-epistemic default.

At the same time, we must also recognize that considerations of justice overall are in play. In the cases of whistle blowers and of victims of sexual abuse, epistemic power tends to be used to expose the abuse of other powers. That one may hesitate to know whom to believe is then also a function of getting to grips with what interests ought to take priority, from the point of view of justice. In a case of whistle blowing, the nation's safety may be at stake. And in a case of an allegation of harassment, we might consider that we should not, as citizens, move to a verdict too soon, as we tend to value the (legal) presumption of innocence. However, it is also true that the alleged abuse of power can often persist for as long as the powerful parties in question can lay a claim to the

knowledge of the abuse, defining what qualifies as abuse or ensuring that knowledge of the abuse does not go out.

In other words, some cases are difficult and lead to divided loyalties, precisely because we sense that there are moral, legal or political interests that make our epistemic verdict particularly consequential – while we may also sense, quite rightly, that those factors hamper our reaching the correct conclusion.

8. Limits to the Default Norms

My aim throughout has not been to argue that we must share knowledge in all cases, but that the values of epistemic justice set a default that knowledge ought to be shared. Defaults can, of course, be overridden. In this section, the aim is to delineate justified exceptions, i.e. when, why and how a proper moral-epistemic response or disposition may deviate from the default norms.

As I have argued for the default norms, their default-status depend on high levels of trust characterizing an epistemic practice. The condition that the "normal circumstances of trust" obtain makes placing trust in others reasonable, such that it becomes the default. Likewise, they make it morally imperative to be truthful in the sense of being honest and open, while it is in order to invest in increased levels of trust and in order to sustain the circumstances of trust that (proactive) knowledge sharing is rightly viewed as the default. This picture implies that the defaults are overridden when the circumstances of trust do not obtain, or – as in the case of knowledge sharing – when we may assume that investments will be misplaced and erode trust rather than further it.

It may be argued that we are justified or even required not to speak the truth when the trust we thereby bestow on others will be abused. Here one can think of the murderer at the door, who wants to know if your friend is home. In this context, Williams has insisted that some people "have no right to the truth" (Williams, 2002, p. 105) and that, the murderer, given his objective, does not "deserve the truth" (Williams, 2002, p. 114), which I have argued should be understood as a requirement of justice. This is to say that "someone deserving the truth" is not a judgement that is about the person in question, or that should reflect the preferences of the person in a position to share: instead it is governed by the moral-epistemic norms that are part of epistemic justice.

Another example where we cannot both be truthful and live up to the expectations of trust, and where the latter takes priority, is given by contexts where speaking the truth would undermine the trusting relationship(s) in question for no good reason. Here one can think of a friend asking if you think she is pretty, whereas you don't think she is particularly pretty but that saying so – or even hesitating to give an answer – would cause offense. In such a case, it would appear Williams's stance is that some questions

should not be asked to start with, as in the adage: "Who asks no questions gets told no lies" (Williams, 2002, p. 212). This is to say that one should not expect others to tell the truth when asking too much. So, the justification in this case for not speaking the truth is that truth is not what is called for: the hearer has no right to the truth, as they ask too much, and in doing so, they prove that truth is not what concerns them. The point of the exchange is not knowledge, and the relations in question cannot serve it.

The same goes for the case adduced by Bernard Williams of the forgetful grandmother who inquires after the health of her son, whom she keeps forgetting has died. Informing her of this fact may amount to being cruel, and there is no "being cruel to be kind" case to be made. There is no real question of epistemic justice here, it may seem. There is no point, let alone an epistemic one, in informing someone of a hurtful fact, if it is clear that the answer given is one she will immediately forget. So again, truthfulness is side-lined by considerations of trust.

Even so, both the situation of the friend and that of the grandmother are cases that are less clear-cut than Williams's take on it suggests. First, lying remains something to avoid when there is no compelling need to lie; if "white lies" are condoned it seems we do not care for the truth enough. Second, it is rather tricky for one to decide for others whether they seek the truth or not, and we are here not providing what they ask for, at least prima facie. I think my hesitance to condone lying in these cases – even if they are "white lies" – can be explained by the fact that these instances occur within a practice which is generally characterized by the circumstances of trust, while the relations in question are also good ones. So, not only do we feel uneasy to deviate from the defaults that our practice sets. There is a sense that giving the dishonest answer means that we do not take our friend, or our grandma, seriously enough; i.e. a sense that they actually do deserve the truth, no less.

Next to this, there is a sense identified by Williams that we should not be dishonest because truthfulness, our caring for the truths, is compromised by it. If we think that the friend has asked too much of us, we may actually believe that the truth serves her right. Likewise, even though the grandmother will be hurt and will forget rather immediately what our answer is, we may believe that the true answer is best as we have then at least managed to relate to her as an autonomous being, in the moment of sanity that she had. In both cases, it seems we may assume too much if we assume that these people cannot handle the truth, or that our relationships with them cannot sustain truth.

In all of the cases, all we can try and do is find the solution that restores the balance, i.e. that furthers epistemic justice by being appropriate to what is demanded of us in terms of telling the truth, within and given the social relationship and particular circumstances. If telling the truth is too blunt or hurtful, one may still refrain from lying, by deflecting the question or by hedging one's answer or by saying, as one may opt to do in the case of the friend, that one is not a good judge of such matters.

Conversely, however, truthfulness can also take priority over trusting another for the truth. An example here is that when we trust others to be truthful because our social

relationship dictates that we should trust them, we can only go so far. Here one can think of Faulkner's example: that trust can be given affectively, to inspire truthfulness, disregarding that person's bad track record. Even so, however, as Faulkner has pointed out, the value of truth will still be operative (Faulkner, 2014, p. 339). That is, I can decide to trust someone so that they speak the truth, whereby my trusting them is a reason for them to be truthful, but such trust can be given only up to an extent. If there are signs they will not take the trust invested in them seriously, they will no longer be trusted for the truth. This is so *even if* that trust was merely affective, i.e. based on the social relation and the expectation that the other would live up to the expectations of truthfulness. So here we can see that epistemic justice governs the situation, too. It explains both why we should trust, such that others can restore their reputation for instance, but also why we cannot ignore how they respond to such trust.

Similarly, we can have faith in the social mechanisms which exist to discern good from bad sources; indeed, in the normal circumstances of trust, we are expected to sustain the system. But when it turns out that the criteria for knowledge and authority are skewed or biased towards one social group, we no longer have that responsibility. In such a case, they effectively undermine the best chances we have as a society to achieve truths (Fricker, 1998/2011, p. 62).

In that case, the idea that we should care for truth means that trust in the social system – and our willingness to sustain the relations as they have been – is to be given up. When the system is not epistemically just, new criteria are needed: criteria that will promote truth and ground trust in the system, i.e. criteria that are epistemically just. An example of a plea for new criteria is Longino's list of alternative values for the appraisal of scientific theories, which will be discussed in the next chapter.

As knowledge sharing is a way of investing trust, we may believe that even when the circumstances are not those of trust we should still share knowledge. Doing so is not required, it seems, if one puts oneself or others at risk. One has an epistemic justification for not sharing knowledge, too, when one may expect that the knowledge is going to be mishandled or undermined or used to undermine one's epistemic power and that of others. And yet, even if one incurs a personal risk that is substantive, knowledge sharing may be the right thing to do, morally. For, one important way to address injustices – of both an epistemic and a non-epistemic kind – is to "speak truth to power". In such a case, the moral value of one's sharing one's knowledge may outweigh the consequences of one's doing so.

We may, indeed, find ourselves in a situation where the reasons to share knowledge and not to share it are evenly balanced. But, as a default, it is important to realize that without trusting another party, it is not possible for them to prove their good intentions and epistemic abilities; i.e. it is not possible to prove their trustworthiness. In a case of doubt, then, it is most often self-protection that dictates caution and non-sharing, whereas it is solidarity, openness, fairness, equality of opportunity, spontaneity, that dictate the opposite.

However, here too, if one cannot trust that the knowledge which one shares will be well-received – i.e. received as (putative) knowledge by those whom one seeks to address – then one is epistemically justified to not share the knowledge, and especially so when doing so comes with a risk to oneself. That is, when sharing knowledge we require that there is sufficient ground in the form of trust extended to us, for us to be able to do so. Sufficient credence should be rightly expected to be given for the effort to be expended. But also, we should feel socially safe to share our knowledge. Epistemic justice, then, encompasses some important exceptions to the default norm.

Conclusion

This chapter focused on the ethics of knowledge generation and sharing, i.e. the ethical side of epistemic justice that deals with the proper use of epistemic power. I have discussed moral-epistemic norms that are relatively uncontroversial, such as the norms Williams provided for accuracy and sincerity, and the norm which is key to Fricker's work, that one should "give credit where credit is due". I also argued that it is important that we should also fill the gap that exists with respect to knowledge sharing. As I argued, knowledge sharing is one of the uses of epistemic power, and arguably one of its most important and justifiable uses.

To show how epistemic justice is involved, I have argued that knowledge sharing is often an act of inclusivity and of solidarity, and hence serves to institute the trust that is needed for social epistemic cooperation. It is importantly also an epistemic act and required for epistemic reasons, such as increasing the diversity of the group, which allows for the improved filtering of biases. This last point will be discussed more in the next chapter.

In this chapter I have argued that epistemic justice can explain why knowledge sharing is the default norm. I have argued that acts of knowledge sharing can be viewed as a matter of solidarity, and a matter of reciprocity or fairness. In the case of solidarity being the reason, sharing is a form of empowerment, as it means extending one's power to another who can pay it forward. Such an act strengthens the ethos of the (already existent) community while it also enlarges it. It can make good on universalist aspirations that are contained in justice as an ideal, and of knowledge as in principle shareable and contestable. In the case of reciprocity as reason, sharing knowledge is a form of mutuality, in the form of due recognition of another as a moral and epistemic equal, who deserves to be trusted. It can also be cast as an act of fairness. Others are to be treated as cooperative members and should benefit from the knowledge that a community yields because of its efforts. Knowledge sharing then can be explained by the ideal of epistemic justice: in the form of an act of solidarity it serves to promote and sustain epistemic justice; cast as reciprocity or fairness it directly follows from it.

Chapter VIII | Towards a Society Well Ordered by Epistemic Justice

Introduction

In the previous chapter I examined which moral-epistemic norms follow from epistemic justice – that give the notion content. These norms, if they can be realized, indicate one way of assessing what it is for our epistemic practices to be epistemically just entailing that epistemic power is used well. In this chapter, I will focus on how epistemic power is to be allocated. This is important for understanding how epistemic power is properly used when allocating credence and assessing knowledge claims. The chapter will also argue that epistemic justice is useful as an ideal by which to evaluate our social epistemic practices.

This chapter develops a Critical Theory Test of the credibility or authority of a source. What I do not spell out here is what amount of credence is proper or when accepting a content as knowledge is proper. It is not possible to provide substantive answers to these questions in the abstract. I will, however, provide a perspective on how we can evaluate our epistemic practice.

Given the social generation of knowledge, my approach means I consider the formal aspects of the constitution of an epistemic, knowledge-generative, community. What are the proper procedures for arriving at criteria for including others into the epistemic community that generates knowledge, and for others to receive epistemic authority? Who should be able to co-determine the criteria for knowledge, and what arguments for a revision of criteria are reasonable? To answer these questions, I focus on the special community of scientists, and at the criteria for knowledge and for participation that are used there. I argue that Longino's proposals for structural requirements are guided precisely by a need for epistemic justice, which is intrinsic to knowledge generation.
In this chapter, the following questions will be central:

1. What is needed for our social epistemic practices to be epistemically just with respect to the allocation of epistemic power?
2. What is the point of evaluating our social epistemic practices in terms of epistemic justice – or of striving for them to be epistemically just?

The first question concerns the way we should understand epistemic justice when it concerns structural issues, and mainly the question of how epistemic justice is to be

achieved in that respect. In response to this first question, I will argue that the following conditions must be met to ensure that epistemic power is allocated on a proper, epistemic basis:

1. Other forms of power than epistemic power should be kept in check. Ideally, the social political context is one where there is no abuse of any form of power, or, at least, social relations should not be characterized by oppression and domination;
2. The criteria that are used in practice for knowledge and for authority need to withstand social scrutiny, i.e. they are to pass a "critical theory" test;
3. The structures of the community should be such as to ensure that it is possible for anyone who wishes to do so to contest the criteria that are used for knowledge and for authority, whereby they are to be changed when they do not pass the Critical Theory Test;
4. Breaches of the norms of the proper use of epistemic power must be sanctioned, i.e. a public ethos of epistemic justice must be upheld.

The above points are discussed in consecutive order in sections 2–6 of this chapter. They have an important procedural bearing on achieving a proper epistemic basis for the allocation of epistemic power, without stating in substantive terms what that proper basis is.

Answering the second question, I seek to diminish any remaining skepticism that epistemic justice is not a subcategory of justice, by clarifying what role epistemic justice can play as an ideal that may be used both to evaluate our epistemic practices and to strive for a society that achieves justice overall. I will argue, more precisely:

1. that an evaluation of the epistemic justice of our practices is valuable, as it offers a form of appraisal that has advantages over an appraisal of the practice in terms of its "truth-conduciveness".
2. that striving for epistemic justice serves to further both knowledge and justice, through securing justice in the generation of knowledge, while knowledge and justice are notions that are to be understood as the epistemic success term and universal ideal that we intuitively liken them to be – but also as essentially contestable notions.

While any conception of epistemic justice presupposes a realist notion of truth as I have argued in Chapter III, epistemic justice has the advantage that it offers a way to appraise and question the ways in which social processes shape our knowledge. By moving beyond a "direct" appraisal of our social epistemic practices in terms of the likelihood that its methods uncover truths, it can question the status of the criteria that are used for any

such appraisal without undermining the idea that what we centrally seek is "objective" knowledge, knowledge deserving its title.

Epistemic justice is a concern for those who take seriously the idea that the notions of knowledge and justice may be misappropriated by the socially powerful. Such appropriation means that epistemic power comes to serve their particularized interests and ends, amongst which their continued domination. While epistemic justice is a form of justice in its own right, and hence a necessary element of justice overall, on my view, it is also a quite important means for achieving justice overall. This will be argued by explaining the way different forms of power conceptually relate to each other.

2. Conditions for Epistemic Justice

How can we determine the proper allocation of epistemic power? I will focus first on the proposal that this goal is secured if the other forms of social power are kept in check. This is compatible with Miranda Fricker's point that a root cause of the more serious types of epistemic injustice is the undue influence of other forms of social power (Fricker, 2007, p. 7). It is a special case of Walzer's general thesis that each kind of power is to be used only in its proper sphere (Walzer, 1983, p. 19). He points out that one way to prevent injustices is to diffuse power, such that one kind of power cannot be monopolized and used to illicitly obtain other forms of power, i.e. influence in another sphere (Walzer, 1983, p. 20).

The proper basis for epistemic power is that it is distributed solely on the grounds of epistemic reasons: that is its correct normative basis that is captured by the idea of epistemic *justice*. On a social conception of knowledge, reasons are epistemic when they derive from a perspective of ensuring that knowledge is generated by the community. Because knowledge is socially generated, the perspective will include attending to ethical and political questions of justice in the generation of knowledge, too.

The aim of epistemic justice is for epistemic power to be regulated in such a way that there is no abuse of it. That abuse includes the misappropriation of epistemic power. Epistemic power is to be had by those who have knowledge and/or credibility, based on expertise. In society, it seems to be natural then that it accrues to those who have the social reputation and/or titles that serve to signify that such authority is rightly ascribable to them; namely those who are educated, as is made visible by their titles, and depending on the importance accorded to that title, or those who may be reasonably assumed to enjoy special privileged epistemic access.

In the next section I will look at what criteria may be proposed as proper epistemic criteria for knowledge – and, in order for me not to beg the question, I will explain what distinguishes such "criteria" from the "conditions for knowledge" that feature on the

traditional analyses of knowledge that I have rejected.

While epistemic power is power of one form, other powers may be used to illicitly obtain it and, conversely, they may be the reason why epistemic power is abused. The social generation of knowledge, and the social attachment of the statuses that signify credibility, means that much can go wrong in our use of epistemic power. However, often, the misuse of epistemic power signals the misappropriation of another type of power. This mechanism is apparent in the example Marx gives whereby the capitalist class can bribe the members of the dangerous class (Marx and Engels, 1983). Here epistemic power is obtained illicitly (by the economically powerful), and abused (by those who are bribed and hence are enticed by receiving some economic power in return). Moreover, it seems right to assume that those who have extensive political or economic power can change the structure to their benefit so that they are more easily able to gain epistemic power (Fricker, 1998/2011, p. 62).

There is, then, a risk that economically or politically powerful groups can distort the proper allocation and use of epistemic power. They can do so by abusing their powers – e.g. commodifying knowledge and education such that only the rich can afford it, or changing the language of instruction to a language only the wealthy can learn to speak, or by controlling the media such that they have a greater influence on what comes to pass (even if wrongly) as knowledge. But the economically or politically powerful can also limit others' epistemic power by diminishing its appeal, targeting – for instance – the efficacy and outward appearance of epistemic power. Undermining the authority of universities as "leftist bulwarks", undermining science as "(mere) ideology", and attacks on mainstream media as providing "fake news", are cases in point. This is not to say that we should always accept the authority of those institutions and people who enjoy it, for our structures are not always epistemically just. Rather, it is to say that epistemic justice is neither a luxurious nor a redundant notion.

3. The Criteria for Epistemic Authority

The practical criteria that I will discuss are not "conditions for knowledge" of the sort that the analyst of knowledge seeks to codify. The criteria for epistemic authority discussed here serve a practical purpose within social practices. They are indicators that signal that the accordance of these statuses is appropriate. On my understanding of the terms, criteria are used to pick out something, i.e. they serve to recognize it. For criteria to be "proper" they must be appropriately sensitive to whatever conditions make a thing the thing it is – but only on condition that the conceptual understanding of the thing in question was correct to start with. In the case of knowledge, as we have seen, some understanding of the nature of the concept is available, but that understanding is not readily codifiable. But notably, even if there *were* a set of conditions, these would still need to translate to criteria that are

appropriately sensitive to them, unless the conditions given are epistemic to start with.

As an example, one may consider that we may understand the verificationist idea that truth is what lies at the end of inquiry as a statement which is merely epistemic, i.e. merely about the criteria for recognizing what truth is. In that case, as has been Misak's reading of Pierce, there is no objectionable metaphysical equation according to which the truth of some matter is whatever is being agreed upon (Misak, 1991/2004, p. 42).

Yet, what this example also shows is that one may wonder whether some of the epistemic criteria are not in fact limiting, or dictating, our metaphysical options. That may be regarded as problematic to the extent that one wants to idealize the idea of what the idealized agreement of a community comes to; if one is of the opinion that truth must be such as to outstrip even our best potential human reach, then a reduction of truth *is* implied.

That criteria can have implications for the conceptual understanding of a term, also means that in the case of criteria for knowledge one may be hesitant to accept some of the practical criteria for it as proper. What truth and knowledge have in common here is that such practical criteria are especially important, in regard to our understanding of these concepts, because of the conceptual difficulties in codifying them.

So, the key question is: what are the *proper* criteria for the allocation of epistemic power in our social-epistemic practices? There is a platitudinous answer: that epistemic power is properly ascribed when it is due or deserved. In such a case, epistemic power is given to those people and sources who actually have knowledge and the requisite authority. Yet, stating "truth-conduciveness" itself as a criterion for the accordance of knowledge and authority, conceived of as social statuses, is not going to be much help.

More practical indicators which may be used to properly accord the status of knowledge to a claim are, for instance, that the claim:

- passes as knowledge in a community (Welbourne, 1985);
- fits in with other, existent knowledge (Quine and Ullian, 1970);
- is used as a basis for action or further theory-building (Williamson, 2002);
- has survived sufficient (scientific) contestation from multiple viewpoints (Longino, 1995; Mill, 1989; Popper, 1979);
- is plausible independent from the socio-political context that provides ideological motivation (Longino, 1995).

Criteria for authority being appropriately accorded include:

- the social status that has been accorded to the source with respect to their competence and reliability – such as titles, prestige, their reputation, their track-record when it comes to answer questions on a particular matter (Williams, 2002);

- the ability of a source to render account for the knowledge they claim (Douglas, 2009. p. 153);
- a source's track record for honesty and their propensity to spontaneous openness and transparency (Williams, 2002, p. 12, p. 75);
- there being no apparent personal motives or interest that motivated the source's thinking or their coming to espouse their particular position; (Williams, 2002);
- the source's expected ability to resist bias and wishful thinking – e.g. given what is at stake for them in terms of their reputation (Williams, 2002; Fricker, 2007, pp. 118–119; Shapin, 1994, p. 75).

However, as we saw, one problem with any criterion that is being used to aid us in allocating epistemic power is that the criteria – or "indicator properties" – may be rigged by the socially powerful, serving to sustain their power. My claim is not that we should reject all criteria altogether. Instead, we should ascertain, in as far as we can, that the criteria that are being used, examples of which we saw above, are not biased, i.e. that they guide us to (relevant) truths, as a requirement of epistemic justice.

This point concerning the correctness and appropriateness of the criteria used is even more important when they are used in a theory of justice, because a good theory of justice should not sustain social domination. For, if the criteria are rigged by the powerful, then epistemic justice itself can be used by the powerful as a means to secure their interests. The difficulty we face here is that both justice and knowledge are concepts which are socially powerful, even as they are – as a result – essentially contested when it comes to their content and the criteria that make for their proper application.

Despite these difficulties, I think we can approach the matter indirectly, through developing a test that any good criteria are to meet. That test will then spell out what criteria may be used when epistemic judgements as to the allocation of epistemic power are proper. First, we will consider how even criteria that were considered innocuous may be regarded as social and political, even in the domain of science where objectivity is aimed for, and what possible remedies have been proposed, for which we turn to Longino's work (1990, 1995, 2002).

4. Longino's Work on Criteria in Science

Before explaining Longino's insightful criticism of the prevailing criteria (or as she calls them "cognitive values"), a first thing to explain is how and why there is a need for criteria in science in the first place. This will mean I need to explain first some of the key elements of Longino's theory of scientific knowledge, and how her work has an important bearing on understanding epistemic justice when it comes to the allocation of epistemic power.

In her work, Longino has argued that the acceptance of scientific theories and models is not a matter solely of observation and logic. She explains that criteria are needed in addition for scientists to decide between theories and models that are comparable but different. That there is a gap between the evidence and what hypothesis confirms it is the so-called "problem of underdetermination" (Longino, 2002, p. 128). In addition, and relatedly, there is the fact that what we regard as evidence for a hypothesis is mediated by our background knowledge and background assumptions (Longino, 1990, p. 75). All in all, to determine what evidence confirms what hypothesis and to choose between different comparable hypotheses, so Longino has argued, we need more than observation and logic can provide.

Longino has argued that a false individualist approach to knowledge has led to undesirable consequences. The main one is a dichotomy between truth and rationality, on the one hand, and sociality on the other. She points out that this dichotomy, which philosophers and sociologists of knowledge share, makes it hard to understand that the social generation of knowledge is possible, without this leading to a relativized conception of truth (Longino, 2002, p. 138).

In the false dichotomy between rationality and sociality, philosophers are concerned solely with the purely normative issues of truth and rationality. Sociality in the other hand, is the focus of sociologists of knowledge, who look at social interaction to find causal explanations of how knowledge acquires the status that it receives in a social group. Seeking to reconcile the validity of their different approaches to the subject, and salvage the concept of knowledge as a normative one, Longino has argued that the fact that sociologically we can see that observation and logic are not fully explanatory of what we come to accept as scientific outcomes (and as knowledge) does not mean that "good reasons" are irrelevant (Longino, 2002, p. 128). The normative project of justification that philosophers engage in remains important.

But Longino stresses that the "problem situation" is a different one than philosophers (epistemologists and philosophers of science) have hitherto assumed (Longino, 2002, p. 128). This problem situation is to be one that includes the social, and that recognizes that the normative questions of what qualifies as good reasons to accept or reject a theory should move beyond observation and logic: it should also look at uncovering the assumptions and commitments that are being tacitly assumed. As we will see this brings Longino to the view that we should also look at the ethical assumptions and political relations in play. The influence of the social should be a good one, as opposed to yielding (ideological) bias:

> The logical problem of underdetermination shows that empirical reasoning takes place against a background of assumptions that are neither self-evident nor logically true. The identification of good reasons is similarly context-dependent, whether this is accomplished by the scientist in action or by the philosopher in reflection. (…….) The point, however,

> should not be that observation and logic as classically understood are irrelevant but that they are insufficient. The sociologists' empirical investigations show that they are explanatorily insufficient. The underdetermination argument of philosophers shows that they are epistemically insufficient. Rather than spelling doom for the epistemological concerns of the philosopher, the logical problem of underdetermination together with the studies of laboratory and research practices changes the ground on which philosophical concerns operate. This new ground or problem situation is constituted by treating agents or subjects of knowledge as located in particular and complex interrelationships and acknowledging that purely logical constraints cannot compel them to accept a particular theory. That network of relationships – among other individuals, social systems, natural objects, and natural processes – is not an obstacle to knowledge, but can be understood as a rich pool of varied resources, constraints, and incentives to help close the gap left by logic. This philosophical concern with justification is not irrelevant, but must be somewhat reconfigured in order to be made relevant to scientific inquiry. (Longino, 2002, p. 128)

To avoid the false dichotomy, then, Longino defines knowledge differently to highlight the commitment to the need for scientific, methodological rigor to uncover the truth as well as the unavoidably social context within knowledge is shaped. Her account of knowledge does not eradicate the standard notions that are key to scientific practice: it assumes that both logic and observation are needed. Yet it "proposes to add to those norms traditionally studied by philosophers (and about whose nature there is ongoing disagreement) those that ought to govern the social interactions also potentially constitutive of scientific knowledge" (Longino, 2002, p. 139).

In her social account of knowledge, Longino stresses that the social generation of knowledge is more than the sum of its parts: through social interactions – not merely shared standards – objectivity is achieved, precisely because "subjective" individual perspectives are transformed into a shared outlook (Longino, 1990, p. 76).

This is also a view I have defended when arguing that knowledge is socially generated through testimony, since the idea is not one of an extended body of warrant whereby individual justifications are weighted and added up, but whereby we are *inter*dependent on one another and on the socially set criteria being applied, which are used so that knowledge claims and purportedly authoritative sources that do not meet the standards are falsified – even if, in incidental cases, we may decide that the criteria are themselves to be revised. Both Longino and I look back to the work of Mill, although I have emphasized throughout this thesis that I do not take Mill's dialogic account of rational warrant to form part of any analysis of knowledge.

In Longino's theory of knowledge called "contextual empiricism" the notions of "epistemic acceptability" and "conformation to reality" play central roles (Longino, 2002, p. 9). The first relates to the social processes, namely critical discursive interactions that are needed to accept knowledge claims (Longino, 2002, p. 158). The second idea of "conformation to reality" relates to the fact that knowledge is to be constrained by the external world, whereby there is room for scientific pluralism, in the sense that different theories and models can explain different aspects of some part of reality, depending on the social context of inquiry, without there being

any one theory that proves to be better (Longino, 2002, pp. 109–111). Instead of a condition of truth for knowledge she uses the term "conformation":

> Conformation introduces an element of constraint by and responsiveness to the world, but not so tight as to exclude non-reconcilable alternative accounts. (Longino, 2002, p. 138)

My argument in what follows is that Longino's proposal for how scientific communities should be structured, so as to ensure that they achieve objectivity, is a paradigm of how, from a political-epistemological perspective, we should structure the institutions that generate bear on generating knowledge and according epistemic authority. An important feature of her theory of (scientific) knowledge (her "contextual empiricism") is her emphasis on the inclusion of a diversity of perspectives, which I will discuss in the next sections.

Objectivity as Longino describes it, is a precise match to Bernard Williams's rendering of truthfulness in his sense of "Accuracy". For as we saw, while Accuracy is the truthfulness that is needed for uncovering truths, he described it as involving "resistance to wishful thinking, self-deception, and fantasy" (Williams, 2002, p. 127). Longino characterizes objectivity as follows:

> Objectivity is a characteristic ascribed variously to beliefs, individuals, theories observations, and methods of inquiry. It is generally thought to involve the willingness to let our beliefs be determined by "the facts" or by some impartial and non-arbitrary criteria rather than by our wishes how things ought to be.(Longino, 1990, p. 62)

Since objectivity equates to one of Williams's two forms of truthfulness on this conception, even while it may be attributed to non-humans, I think it is apt to say that in characterizing the objectivity of a social epistemic practice, Longino was in fact describing what it is for a community that seeks to generate knowledge or, which is equally important, to eradicate falsehoods, to be *epistemically just*.[1]

To make a case that even criteria for theory acceptance that were considered innocuous in the past have a social, political or moral valence, Longino has compared the traditional list of cognitive values with what she calls "feminist values" (Longino, 1995, p. 392). The aim of this comparison is to show that currently endorsed cognitive values reaffirm the socio-political status quo (Longino, 1995, p. 396).

Her point is that all values that are used in science to appraise theories and models are embedded in a socio-political context, which they can serve or disrupt, and therefore all supposed purely cognitive values have moral-political valence. Longino's list of

---

1 Longino maintains that both of these – generating knowledge, and falsifying knowledge claims, are goals of scientific inquiry, that may conflict (Longino, 1990, p. 32-34). She also notes that the goals of science determine what criteria are used to evaluate its success (Longino, 1990, p. 18).

comparison is as follows:

Feminist list	Traditional list
Empirical adequacy	Accuracy
Novelty	Internal/External Consistency
Ontological heterogeneity	Simplicity
Complexity of interaction	Breadth of scope
Applicability to human needs	
Diffusion of power	Fruitfulness

(Longino, 1995, p. 392)

The traditional list of values are those that philosophers of science would accept as values that enhance the likelihood of truth (Longino, 1995, p. 383). As Longino insists, for a long time she herself thought that these values were socio-politically neutral, i.e. purely "cognitive values" that have a "solely epistemic or cognitive basis" (Longino, 1995, p. 384). However, she states she is no longer convinced that this is so, and argues in the paper that the values import non-cognitive moral or political considerations, as they are shaped by the status-quo and serve to retain it (Longino, 1995, p. 396).

The "new" cognitive values which she offers as an alternative are feminist in the sense that they serve to counter male domination and gender bias, and more generally, to end oppression and marginalization in the many ways it is sustained in the social structures (Longino, 1995, p. 396). The label "feminist", so she explains relates to the aim shared by feminist of ending oppression (Longino, 1995, p. 391). The list she gives is a derived from the issues feminist scientists and theorists address and the ways they go about doing so. Longino claims that the values she has distilled are the values feminists endorse in their (scientific, scholarly, and philosophical) work. While her point is to draw a contrast, she is open to amendments on what exactly the feminist values are and does not claim that her list is complete (Longino, 1995, p. 391).

The value of heterogeneity, for instance, is an antidote to the value of simplicity. It prevents a standardization that serves both to obscure relevant ontological differences and subordinate those who do not conform with the standards. The value of novelty is more likely to promote theories and research that changes certain long held beliefs and disrupt the status quo rather than perpetuate the values of internal or external consistency. For when consistency is valued, it means that the theories that are being accepted are the theories that fit in best with the existent ones. Similar points about supporting the status-quo of science can be made about the values of breadth of scope and of fruitfulness. Breadth of scope means the theory has to be broadly applicable, but this can mean that the subject matter chosen is one that glosses over differences all too readily, while standardization is being looked for.

Fruitfulness, which admittedly has less of a secure status among the traditional values, means that the theory should be one that serves science itself to yield more findings. But this in effect means that the theory has to work well with existent theories, and hence that it will not question but serve to support the already existent theories. Longino pits theoretical fruitfulness against an alternative that she labels "diffusion of power". Indeed, that value seems to be the opposite, as it seeks to further theories that are able to work for society, and particularly to make it work for those who have suffered from society's oppressive structures, i.e. for groups who had little or no access to it, or who did not matter to its objectives.

I think this list of alternative criteria, contrasted with the traditional ones, demonstrates that there are alternative criteria to the ones in use. It shows that criteria which were always taken for granted can be criticized and replaced. Longino's list also offers a resource for achieving enhanced epistemic justice by providing criteria that aim to counter domination and to be more inclusive. Scientific explanation is still answerable to logic and empirical adequacy, but the broader range of cognitive values allows for ensuring that theories which are accepted are not simply those that fit best with the status-quo, which tend to favor the socially powerful.

5. Structural Requirements for an Epistemic Community

Longino arrived at the need for norms that structure a scientific community and other knowledge communities through a realization that scientific knowledge is socially embedded and its putatively "pure" cognitive values reflect broader societal norms. Her main point has been that, even while we are right to think that "knowledge" is a success term, such that the aim is a normative one of arriving at truths, the conditions of generating knowledge and criteria that are used for it are social. An important aspect of this claim is that the sociality of our criteria for identifying proper claims to knowledge means that these criteria are determined, in part, by those who are deemed to be qualified to decide on them. It explains why the constitution of the community is an important matter.

As I have stressed throughout this thesis what is deemed knowledge within a community depends on communal standards even while those standards may be mistaken (This represents a general commitment to fallibilism.) Accordingly, the proper constitution of a community is key for achieving what may be called "objective knowledge", i.e. knowledge not tainted by biases or by the promotion of particular group interests. Such "objective knowledge" is arrived at through ensuring that the generation of knowledge is epistemically just. This seems to be also what drives Longino's insistence on criteria for knowledge that further non-domination and her emphasis that objectivity requires that the community should be inclusive.

Even so, as Longino has also made clear, it is undeniable that we cannot achieve knowledge without there being social standards for knowledge: not just anything can qualify as knowledge because what properly qualifies as knowledge is to be constrained by the way the world is: our scientific models and theories have to "conform" to reality, and our everyday knowledge claims must be "true" properly to qualify as knowledge.[2]

Openness to criticism from a diversity of viewpoints, then, is required for the epistemic reasons of scrutinizing hypotheses and falsifying those that are to be discarded, and to resist bias (Longino, 1990, p. 76). It is to counter, specifically, bias which is a result of socio-political influences on the outcomes of science, given that scientific knowledge is a social product. As Longino puts it: "certain features of community structure are important to the knowledge productive capacity of a community" (Longino, 1995, p. 384)

To ensure that criticism is expressed within the community, Longino provided the following four recommendations for the structure of a scientific community:

1. Venues. There must be publicly recognized forums for the criticism of evidence, of methods, and of assumptions and reasoning. (…)
2. Uptake. There must be uptake of criticism. The community must not merely tolerate dissent, but its beliefs and theories must change over time in response to critical discourse taking place within it (…)
3. Public standards. There must be publicly recognized standards by reference to which theories, hypotheses, and observational practices are evaluated and by appeal to which criticism is made relevant to the goals of the inquiring community. (…)
4. Tempered Equality. Finally, communities must be characterized by equality of intellectual authority. (…). (Longino, 2002 p. 129–131)

The reason Longino included "Tempered Equality of Intellectual Authority" is that she thinks it will ensure that even majority values can be made visible, through increased diversity (Longino, 1995, p. 385). This idea of inclusion of a diversity of perspectives could be viewed as a requirement of political justice, because we want fair representation, inspired for instance by looking at justice from the perspective of that which Iris Marion Young has called the "politics of difference" (Young, 1990, 2000). But it could also be a self-standing epistemic reason, as Longino appears to argue: more perspectives must be brought to bear and if they do not count equally in the process of knowledge formation, then they are not being taken into account at all. In the below I argue, it is both: it is in fact a requirement of epistemic justice. The link between inclusivity and objectivity as underscored by Longino makes that clear.

---

2 Conformation is a broader evaluative notion than truth. Longino has argued that a broader evaluative notion than that of a binary choice between truth and falsehood is needed in science, because models and theories typically relate to aspects of reality, which they conform to up to some degree (Longino, 2002, p. 117). The idea of capturing relevant aspects of reality is presupposed when we evaluate a proposition as true or false, and so truth is the correct evaluative term when we consider propositional knowledge as captured by the concept of "everyday knowledge" (Longino, 2002, p .109; see also Longino, 2002, p. 115-117) – even if, of course, there can be questions of accuracy that make clear that an everyday claim is not capturing the relevant aspects of reality, whereby the notion of "conforming to reality' would seem appropriate.

In the context of understanding these four conditions, it is worth noting that Longino argued in earlier work that they determine the degree to which a scientific community is *objective*, as they are "are necessary to achieve the transformative dimension of critical discourse" (Longino, 1990, p. 76). This also explains why, for Longino, diversity is key to epistemic success, even if it may mean that the norms that are used within a community cannot be taken for granted as: they may also be questioned, too. As such, diversity requires inclusion; i.e. a striving for a community that is open to welcoming newcomers and particularly those who have a different background, method, or point of view. It also means that she emphasizes that those who challenge the status quo, even if they represent a minority, are to be given a real bearing on the outcome.

So, on Longino's view, objectivity requires sociality to start with (Longino, 1990, p. 74). This is contrary to the mainstream view that sides with individualism, but it is akin to Mill's position (Mill, 1989, p. 31). It is explained by the role of sociality as filtering out undue influence that can be wielded in the process of the social production of knowledge. The guiding thought is that objectivity requires criticism from alternative viewpoints, "ideally as many as are available" (Longino, 1995, p. 384).

Because, on a social conception of knowledge the criteria for knowledge are socially set and maintained, there is a close link between the criteria for epistemic authority and the group that decides on these criteria. This is not only because the group decides on these criteria – whereby more diversity of viewpoints represented means there is less chance that the criteria are biased. It is also because the criteria that are in use determine who are qualified to participate in the knowledge generating processes. When new criteria are used, it may mean that new people are allowed in who did not qualify according to the older criteria.

The close link between the criteria and being regarded as a proper member of the community who is able to shape these criteria also means that for those excluded by the criteria it will be hard to contest them: they tend to not be taken seriously to start with, and their complaints about the criteria in use will also not be taken seriously. Even so, there is a point of resisting criteria for knowledge that are unfair, and for appealing to epistemic justice.

The need to appeal to justice in this epistemic context is what I regard as the main message of the paper "Marginality and Epistemic Privilege", by Bat-Ami Bar On, which predates the literature on epistemic injustice. What Bar On recognizes is that what is needed to gain epistemic power is not (or at least, not ultimately) for those who are epistemically marginalized to claim that they *do* have knowledge, even by the standards of the socially powerful, namely knowledge in the form of "epistemic privilege". Instead, Bar On argues that they should resist the criteria for knowledge that unfairly work against them and that serve to deny that they have knowledge to start with.

The argument she gives for this is that, as long as the rules of authoritativeness

are set by those in power, any claim to special knowledge can be easily dismissed. Instead, she argues that the way epistemic power is being allocated through the social structures needs to change. Thus, Bar On argues that what is needed to prevent and end silencing is that the perspectives of the marginalized are included *as a matter of justice* (Bar On, 1993, p. 97). It is, in the language of Fricker, to give everyone *due* credit.

Yet, despite the risk of perpetuating domination, when spelling out what epistemic justice requires in an individual case, we do need criteria for knowledge and for authority. Such criteria are needed for any verdict about what the "proper" use and allocation of epistemic power come to. We will look at Longino's proposal of Tempered Equality next, to see if it offers a solution. Then I will argue that what we really need is a check on the allocation of epistemic power in the form of "a Critical Theory Test" as developed by Williams (Williams, 2002).

6. Tempered Equality of Epistemic Authority

An interesting challenge to the idea that authority should be allocated on the basis of merit is Longino's requirement of "Tempered Equality". This requirement is backed by the idea of diversity, as Longino explains:

> The extension of equality of intellectual authority to all qualified participants is intended to require representative diversity in the community. (Longino, 1995 p. 385)

However, the requirement of Tempered Equality is not as radical as it may seem. This is mainly because the rule that all participants to the knowledge-generative process are to be taken equally seriously, unless reasons speak against it, applies to participants who are already included in the epistemic community. That is, it already concerns "qualified" participants. (Longino, 2002, p. 131; as noted too by Douglas, 2009, p. 128). So, whoever is to be taken equally seriously is already "properly" a member of the scientific community, i.e. someone who does not need to earn her credentials. It is not any random person, making a random remark at an inappropriate time.

But when is a person qualified? That requires a social judgement, too – while that is not inherently problematic, as long as the criteria can be contested, it is important to attend to the point that there is always a political question of: *who* decides that these persons are the knowledgeable, authoritative ones; that they are experts or that they meet the minimum requirements for equal participation; that they need and deserve to be informed; that they need and deserve to be believed? This is so, not least, because the idea of epistemic justice, too, is premised on the idea that there are better and worse social judgements about the criteria and their application.

As Douglas correctly points out, Longino's view that participants are qualified

requires that participants meet some objective standard for partaking: there is an assumed, largely hidden, barrier to entry (Douglas, 2009, p. 128). Indeed, Longino herself has noted this. Moreover, Longino recognizes that the relevant questions with respect to inclusion concern not only who may be included, but also *who* the "we" are that get to decide on this (Longino, 2002, p. 134). This concerns "membership", and how it bears on the criteria that are going to be created for inclusion (Longino, 2002, p. 133).

Instructive, here, is that the question of membership is one that pertains to justice as much as it pertains to setting the right social criteria for knowledge (epistemic justice), which is what we should expect if epistemic justice is a form of justice. Recently, in her political philosophical writing, Nancy Fraser has described membership as the third large issue of theory of justice. While she has called it a "dimension" of justice, alongside "recognition", and "redistribution", she has also clarified that it concerns a question at a higher level. For she labelled it not only as the dimension of "representation" but also and more insightfully, "the question of the frame" (Fraser, 2005, p. 47). That the question is a pertinent one, recognized by political philosophers, is also made clear by Walzer who writes that, conceivably, the community itself is the most important social good to be divided (Walzer, 1983, p. 29).

With respect to the question of the frame or of proper representation, I believe that a universalist conception of the different forms of justice may be upheld. As may have been surmised from what I have said about epistemic justice, so far, my notion of epistemic justice is one that is to be universal. Even as I recognize that the criteria for knowledge may differ in different communities, I think that some communities are better at achieving knowledge than others: there are universal standards that we should all aspire to. Having said that, it may also be true that it is dangerous for any society to feel confident that they have all the right criteria in place. There should be room for progress, as well as for resisting the views of those who are socially powerful. Therefore, I think that it is important to be critical of the criteria most of all: all knowledge should be subject to revision, and all epistemic criteria and norms of epistemic justice are so, too. But, *also* in order to be critical, we do need guidance. This is a gap that I think can be filled partly by Williams's "Critical Theory Test", which I turn to below.

First, however, it makes sense to discuss how "reasonableness" is understood in the epistemic and political domain. Criteria of reasonableness are currently what guide us in assessing claims of knowledge and of authority: they are the needed limits to our open-mindedness. For we should be able to contest knowledge claims and criteria for knowledge, and able to properly epistemically exclude those views and criteria ad sources that undermine the outcomes of epistemic practice. Some notion of reasonableness and of fair contestation is needed alongside open-mindedness precisely because not "everything goes".

As Longino has herself remarked, in addition to Tempered Equality (of Intellectual Authority) as a rule for participants who are qualified, her model of a properly structured scientific practice also includes "Public standards" of contestability, that serve to evaluate theories and "by appeal to which criticism is made relevant" (Longino, 2002, p. 130). These are the standards, then, which one's viewpoint must live up to be included in a critical community (Longino, 2002, p. 133). It serves to protect against a "cacophony" that would arise from having to take all views seriously with no standards, and someone who "reiterates the same complaint no matter what response is offered" will eventually no longer be deemed qualified (Longino, 2002, p. 133). In what follows, I will discuss and compare two complementary views on reasonableness, presented respectively by Douglas (2009) and Young (2003).

Heather Douglas maintains that our standards for critiquing and dismissing scientific theories must live up to the same kind of accountability that we expect from scientists themselves:

> We should be asking scientists, particularly those with outlier positions, what evidence could conceivably convince them of a specific claim about which they are sceptical. If no possible evidence could convince an expert of an empirical claim, we can surmise that values, rather than evidence, are the primary support for their rejection of the claim, an improper direct role for values and an indicator of junk science.
> (Douglas, 2009, p. 153)

This point is one that helps us see that when dismissing scientific claims, scientists are expected to play by the rules of science. One of these is that they should render account such that we have a grip on their doing so. This expectation that scientists are able to explain their position and justify it, may apply to scientists to a higher degree than to the rest of us, because of their perceived status as scientists, hence as having an equal footing and basing themselves on the same rigorous standards as their peers.

Yet, as Douglas has made clear since, this rule of contestation also applies to a lay person who seeks to reject some scientific claim, if they claim that there is no good scientific proof for it. There she insists the same question should be asked, namely whether any scientific evidence could convince the person to start with (Douglas, 2016). If no sufficient evidence could conceivably be given, while the claim is an empirical one, then that indicates that one is not prepared to change one's beliefs or one's ideological convictions. In that case, the rejection of the scientific view is purely based on value, and on dogma, and hence is not scientific but ideological (Douglas, 2016). In defending this view, Douglas is in effect defending the authority of science. While scientific methods are those that are most rigorous, a denial of a scientific finding has to live up to standards of science, unless one is prepared to admit that one simply does not *want* to believe the scientific outcome. It is to insist that, even though knowledge is fallible, to persist in "bucking the full weight of evidence is to persist in being unreasonable", as Quine and Ullian put it (1970, p. 40).

Similarly, Iris Marion Young has a notion of reasonableness in mind that refers to people's dispositions, rather than the content of what they say or their ability to justify it:

> Reasonable people can have crazy ideas, but what characterizes them as reasonable is their openness, their having "an open mind".
> (Young, 2000, p. 24)

Young adds:

> To be reasonable is to be willing to change our opinions or preferences because others persuade us that our initial opinions or preferences, as they are relevant to the collective problems under discussion, are incorrect or inappropriate. Being open thus also refers to a disposition to listen to others, treat them with respect, make an effort to understand them by asking questions, and not judge them too quickly. (Young, 2000, p. 25)

While similar in spirit to Longino's insistence on inclusivity and Douglas' insistence on open-mindedness, Young's position is quite strikingly more radical than either of their positions. As we will see, she seeks pre-emptively to dismantle any claim of superiority of those who have enjoyed education over those who have not. Her view is one that is admirably clear in rejecting any claim to have more right to speak or to be heard than those who have no or lesser education when that education is not actually all that relevant; whereby, of course, the socially powerful tend to be the ones who have enjoyed more education, and are often also the ones who make the decisions in question. Young's position will be discussed in the last section of this chapter.

7. A Critical Theory Test for Knowledge and Authority

We have seen that even while Longino offered a principle for bestowing epistemic authority that would seem to be what epistemic justice requires, so far, it is somewhat hard to know how much diversity her idea of "Tempered Equality" really allows for because the participants have to be "qualified". To develop this point, I believe that it is helpful to look at the "Critical Theory Test" which Williams developed in Chapter 9 of *Truth and Truthfulness*. That test seems particularly helpful, as it provides a check on the allocation of authority, and is framed from the viewpoint of seeking to do justice to those who may potentially be unjustly excluded or marginalized even if they may not themselves realize it, because the social structures are accepted by most people for the most part.

Williams's test is an adaption of what he calls "the Critical Principle", which might be thought sufficient for countering any wrongful allocation of power (Williams, 2002, p. 221).

The central idea of that principle is that any acceptance of an unequal distribution of power as legitimate should be uncoerced, because "coercion cannot constitute legitimation, and... some methods of belief formation are simply coercive" (Williams, 2002, p. 221). As Williams formulates the Principle it reads:

> Suppose that of two parties in the society, one is advantaged over the other, in particular with respect to power; and suppose that there is a story which is taken to legitimate this distribution, a story which is at least professed by the advantaged party and is generally accepted by the disadvantaged; and suppose the basic cause of the fact that the disadvantaged accept the story, and hence the system, is the power of the advantaged party: then the fact that they accept the system does not actually legitimate it, and pro tanto the distribution is unjust. (Williams, 2002, p. 221)

Williams thinks that that principle needs some amendments, because 1) those within the system may be oblivious to advantages that accrue to those with greater power, or believe that these advantages are deserved or that they are offset by other advantages accruing to others, while from an external viewpoint this is not so; and 2) the claim that the causal history of the belief is what is problematic about it runs counter to thoughts that a belief can be sound even if it is caused by someone's exerting power (Williams, 2002, p. 226).

For, as Williams remarks, the force of the better argument is a force too – indeed, an exercise of epistemic power, as I see it – and so it will be necessary to distinguish good and bad exercises of power, which will be hard to do from within a practice whereby the authority of those with power is accepted (Williams, 2002, p. 227). Secondly, the originating source being coercive may not speak against the soundness of the beliefs. For, as he remarks, students may be coerced to go to school, but that is no reason to doubt the soundness of what they are taught (Williams, 2002, p. 226).

Accordingly, Williams proposes a refinement, what he calls a "Critical Theory Test", which should be suitable for parties inside a practice (Williams, 2002, p. 227). To tackle the first obstacle, it may be good not to think in terms of advantages, but in terms of pointing to a given power distribution being unequal, which even people inside the practice cannot deny, while such an unequal distribution of power is not in itself always illegitimate (Williams, 2002, p. 223). The Critical Theory Test which Williams proposes, then, starts from there being an unequal power distribution and there being a socially accepted justification for it. The test is a question aimed at the insiders who have lesser power. It asks them regarding their accepted belief that the power differences are justified: "If they were to understand properly how they came to hold this belief, would they give it up?" (Williams, 2002, p. 227).

This question is one that tackles the second obstacle, since the focus comes to lie not on what has caused the belief, but on its normative force (Williams, 2002, p. 224). In using this test, as Williams notes, it will be hard for insiders to question their belief concerning the justification of a power distribution, and so often an influence from outside is required in order to question how they came to have the belief, and what its normative force is (Williams, 2002, p. 227). Moreover, challenging authority of the socially powerful will tend

to be a gradual process.

The crucial point is that the normative force of any justification of an inequality of power is undermined when there is no independent ground that can justify the unequal distribution. The process of realization that such a ground is lacking is analyzed by Williams in three steps:

1. there is no other ground to accept it than the say-so of the powerful (Williams, 2002, p. 227);
2. the power or authority of the powerful derives from the system which they claim is just (Williams, 2002, p. 228);

These first two steps mean that there is no independent ground that furnishes the inequalities in power other than the fact that this belief is promulgated by those in power.

3. that there are good explanations both why those who are powerful think that the system is just, which is that they benefit from it and it may not even occur to them to think otherwise, and why they should teach that it is just (Williams, 2002, p. 229).

The belief that the inequalities are just can come to be undermined, when people realize that 1-2-3 are all true: the story told by the powerful serves their own purposes, and there is no independent other ground that justifies accepting the unequal power distribution. The core of the Critical Theory Test, then, is that any belief that a given unequal power distribution in a society is just should not be the result of the exercise of power, but should be based on good reasons. Good reasons here are the reasons which do not rest on the exercise of power by those who have more of it. They may be reasons that are caused in some way by the powerful exercising their power, but their normative force ought to be independent from the exercise of power: it should explain, or even justify, why these people have power in the first place.

This is especially important because having greater power tends to be an advantage, and most power is retained through the idea that the powerful are entitled to their power; i.e. that their authority is proper and that their advantages are deserved or that the advantages they have are offset against advantages that others receive as a corollary.

What is interesting about Williams's Critical Theory Test is that its central concern is the "rejection of power that falsely presents itself as cognitive authority" (Williams, 2002, p. 231). This means that the test is devised precisely so as to avoid and challenge any undue allocation of *epistemic* power, which may be claimed by those who enjoy other forms of social power and who by claiming epistemic authority are able to retain and increase their position of domination.

Further, the test is devised so that it offers us a way of rethinking claims of epistemic power that are made by those who enjoy more social power than others in our own epistemic practices – starting with their very claim that they deserve to be regarded as

authorities. If there are independent reasons for it, then all is good. But if the supposed independent reasons turn out to be self-serving, then this will undermine the authority that was previously accorded.

Accordingly, I think the Critical Theory Test can work as a tool for rethinking the criteria for epistemic authority that we tend to take for granted in everyday life (and in scientific practice). I also think that it helps us formulate a core principle that lies at the root of an epistemically just practice, when it comes to the allocation of epistemic power. That principle, which the test embodies, is that authority must always be backed by some independent reason for it, other than self-serving reasons which the powerful might give. The test, however, is a negative one. Rather than telling us which criteria are proper, it provides us a with a means of questioning and, if necessary, challenging justifications for any type of unequal distribution of power. – It is worth mentioning that Williams considers it a virtue of the test that no assumptions are made that "the way of reason" leads us to truths about justice (Williams, 2002, p. 229). The test can also be used to challenge criteria for knowledge or authority that serve to exclude people for no good reason, even while these criteria might be presented as epistemically just.

Previously, I suggested that the Critical Theory Test can be applied to reinforce Longino's point that criteria that have shaped the status quo in science are not thereby, necessarily, correct or just: the values in science used for theory acceptance tend to sustain and reinforce the status quo (Longino, 1995, pp. 392–396). This is problematic because the social and political context has been one that readily excludes the participation of women and minorities; and unfortunately, the criteria in use support this exclusion (Longino, 1995, p. 392). What the Critical Theory Test helps to show, I suggest, is how and why Longino's argument that we should consider alternative values to the traditional ones is actually a powerful one: her argument shows that the traditional criteria do not meet the Critical Theory Test.

First, it raises awareness that the criteria that are in use are not to be taken for granted, and that their legitimacy can be questioned. Secondly, it makes a persuasive point that while they are generally accepted as legitimate criteria, their authority needs an independent basis. Their current basis appears to be that they are the criteria that are used by the main figures who enjoy epistemic authority in the scientific field, and whose authority goes unchallenged as a result of these values being upheld as the only proper criteria. While it may be true that science as it has proceeded has a good claim to authority, Longino has convincingly argued that these criteria, at least in our current socio-political context, are not innocuous: they have a political valence that "naturalizes" the socio-political relations of domination (Longino, 1995, p. 393–394). This then brings us to see that the traditional values do not pass the Critical Theory Test, at least not if they are presented as the exclusive criteria for science, because for the claim that they are superior to other criteria they have no independent justification. Moreover, any superior status these traditional values continue to enjoy hampers the inclusiveness of science, in terms of not having room for those who seek to challenge the status quo, while also pre-emptively blocking alternative theories. It means

that science is in fact not working as well as it should. As Longino illustrates, the result of the excluding of other views and viewpoints is that scientific practice has for a long time failed to grapple with societal issues that have urgency and relevance, while it also misses out on new and different alternative theories and models for doing so.

Conversely, with respect to the Critical Theory Test, Longino's argument illustrates how that test serves its purpose. It can be a devise for making visible epistemic injustices that tend to remain invisible, in cases where epistemic as well as non-epistemic power have been safeguarded through a skewing of epistemic criteria, such that the differences in power seem justified. It can be used to question the epistemic authority of those who tell us that the power differences are justified, and to ascertain the merits of the justification provided.

Concerning the questions "who should be deemed qualified?", and "what are the proper criteria for knowledge and authority?", the outcome of the Critical Theory Test shows that we have reason to be both careful and humble. For if *even* relatively uncontested criteria that were used in science can be found to be unjustly exclusive, while science is supposedly open to revision and should strive to be fair in its appraisal of all available sources, then in everyday life, too, we should be careful not to dismiss others too readily. Whatever we think are good criteria for excluding others as putative sources – or for regarding ourselves as epistemically superior – may not actually be justifiable, in the sense that they would pass the Critical Theory Test.

It is important for us to understand, then, that upholding epistemic justice requires not only a care for consistently applying our existent criteria, but also, that we should always be willing to rethink the criteria in use, and see if they are not in fact unduly excluding others. Similarly, as Young has argued, which we will turn to next, there is a need to rethink any formal criteria that stand in the way of people with lesser education to be heard and represented properly in a deliberative democracy. While that may appear a political point, it is, so I argue, a point of epistemic justice.

So, what are good grounds for an unequal distribution of epistemic power? Here the most plausible ground is an argument that those who enjoy more epistemic power than others have been granted that power because they have more to contribute and that the very possibility of their being able to contribute more *requires* that they are allocated greater authority than others – even aside from requiring any rewards (including the enjoyment of the privilege of epistemic status) that may incentivize them to put in the effort.[3]

There are arguments for and against this idea, that greater authority is needed to be able to contribute more. The argument for it, which we are most likely familiar with, and which is consistent with Craig's genealogy too, is that even while everyone has a voice, not everyone will be "as well positioned" as another vis-à-vis the facts. To obtain

---

3 Note that there is a parallel here with Rawls' Difference Principle, which also assumes that the default ought to be equality as we are moral equals and taking part in a cooperative scheme (Rawls, 1999). Taking that as a background assumption, any difference in treatment (in receiving greater status or material goods) has to be justified; both Young (1990, p. 14) and Fraser (2000, p.108) also assume this starting point in their theories of justice.

truths, then, we should care that we get our information from the most reliable source. A similar thought is Williams's own idea that speech acts – and participation more generally – have to be regulated in any knowledge-productive community: we have to "filter against cranks" (Williams, 2002, p. 217). This is because if we were to listen to what they have to say, it would likely hamper or undermine our ability to obtain truths, even if Williams recognizes that in some cases, it may prove that the cranks were right in retrospect (Williams, 2002, p. 217).

What the argument for giving greater authoritative status to some sources than others is based on, then, is the idea that we cannot take every source equally seriously. This is not only because doing so is not efficient and because time is limited, but because it is assumed that only one account of the facts is ultimately to be believed: only one of the views will be right or the most accurate.

In her work, however, Longino takes issue with this last assumption. She believes that "scientific monism", which is the view that there is only one true or accurate account of the world, in terms of representing the facts of it, is unduly restrictive: to assume such monism is to rule out other accounts prematurely (Longino, 2002, p. 95). Such a prejudgment is not a scientifically apt attitude, as it does not further our ability to gain knowledge but instead blocks possible avenues. The scientific pluralism for which she argues makes it possible that there are competing accounts of the facts, which may co-exist, in so far as they may each be useful to understand different aspects of reality (Longino, 2002, p. 141). On her view, scientific theories are too complex to be evaluated as simply "true" or "false". That binary evaluation is what lurks behind scientific monism, and is also why people believe it is not possible to be a metaphysical realist whilst being a scientific pluralist (Longino, 2002, p. 115-117). But she thinks that such a stance is mistaken, and introduces, instead, the broader evaluative category of "conformation" that can be used to appraise how well theories fare in representing relevant aspects of reality.

While I do not have scope here to go into the full ramifications of Longino's claim, I think it is worth noting that a social account of knowledge which I have defended so far is compatible with Longino's liberating stance on the issue of the possibility of allowing for epistemological pluralism in science. This is because we do not assume an account of knowledge whereby knowledge is reduced to "justified true beliefs", which assumes that – the limited concept of – truth is the arbiter of knowledge. On a social account, whereby knowledge is a social status which is assigned properly only if there is the right kind of connection to the external world that is to constrain it, the broader, evaluative category of conformation can feature in that role. And, importantly, this is not a loss of realism or of a "caring for truth", because we will *still* be resisting wishful thinking and not allowing that anything goes.

An argument against according greater authority to one party from the outset, is that if it is rightly assumed that that party has more to contribute, then surely, that should

suffice for their making a greater contribution. Why, in such a case, would greater authority actually be needed? The only plausible answer here would be along the lines that it serves convenience. For if the greater value – the more likely rightness – of a person's contribution cannot be told from the nature of that contribution itself, then it is actually rather questionable whether that contribution is of greater value. As Douglas emphasizes, one difficulty is that "lay people" may not be able to ascertain the validity of scientific claims and that telling who the experts are, and which experts are to be trusted most, is impossible (Douglas, 2009, p. 45). In such cases, we rely on other scientists to vet the claims, even while we may also expect scientists to be open and willing when explaining their views as well as the choices they made along the way (Douglas, 2009, p. 153).

But in everyday life, it seems natural to expect that those who have more to offer are also able to explain their views better, or provide the more insightful answers more promptly, and therefore, there is no need – so it appears – to rest one's case squarely on one's authority. Nonetheless, it can be *convenient*, as Craig has pointed out, that we have social epistemic statuses. This is because it saves all of us time if and when those who are able to steer the discussions in the right direction are easily identified – and this without having to first listen to all the different views being given and defended. Thus, a case can be made for assigning epistemic statuses, but the case is far less clear than one may think. Indeed, the statuses enjoyed by people, while important, may be more readily overridden than we think. This is because they exist only for convenience, and whatever authority a person has always still needs to be deserved – e.g. by their being knowledgeable about the very question at issue, and by their keeping their knowledge up to date.

8. Towards an Epistemically Just Society

In response to the second question posed at the outset of this chapter, namely, the question what point there is in evaluating our epistemic practices from the viewpoint of epistemic justice, my response is that epistemic justice serves to secure that both knowledge and justice are achieved together. Epistemic justice is an ideal that guards against the threat that notions of knowledge and justice are misappropriated by the socially powerful, whereby they come to serve their particularized interests and ends amongst which their continued domination.

Epistemic justice and other forms of justice are interdependent. Epistemic justice requires for its existence that the social relations more generally are characterized by justice, in the sense that other forms of justice – e.g. political, legal, physical, economic – are also upheld. At the same time, epistemic justice is itself a necessary element of

justice overall, while appeals to it are an important means for achieving it.

A starting point is the default assumption that all should have equal powers, unless there are good justifications for the difference. A reason for ensuring diffusion of all forms of power is to avoid their abuse. Here it is assumed that, generally, when a power is allocated well, it is not being abused – even if it is no guarantee that it will not be abused. It is when a power becomes a means of dominating others, i.e. when it is monopolized by a group in some way, that abuse of it is likely. Any abuse tends to serve the point of securing an improved quality of life, i.e. higher standards of living, for the group who have the power, at the expense of others.

But the abuse is also more likely than not characterized by an attempt to consolidate the power which is had in some other domain. That is, a concentration of power may lead to domination and oppression of others in the relevant domain, but – pivotally – by illicitly acquiring power and abusing it in some other domain, whereby it appears epistemic power is especially fruitful, as it means one can lay one's hands to justifying that one's greater power is deserved (Williams, 2002; Fricker, 2007). As we noted, the appropriation may go unnoticed; it can happen quite subtly through a skewing of the social mechanisms that tend to be susceptible to these other powers.

I have argued that a just society is one that allocates all powers correctly, whereby its use and allocation are subjected to norms specific for the domain in question. It is also a society that, countering domination, will insist on equality as a default for its allocation. This is because, as I have argued, avoiding an abuse of power means that there is no accumulation of power(s) by the few – or by specific groups in society. Accordingly, a just society it is to be characterized by such notions as: inclusion and diversity, empowerment of those who have lesser power, which includes education (in a society where this is important) and effective arena's for (political) contestation, but also one which ensures that all can and will be heard and are able to partake in the production of knowledge.

With respect to that last point it is important that any criteria used for allocation of knowledge and credibility should be subjected to tests. Equality of all powers, including epistemic power, is to be the default; therefore, we need reasons for divergence and these are to be based on more than self-servingness (packaged as stability of the social order, fruitfulness of research or efficiency in communication). Moreover, when assessing the criteria that are apparently justified, for excluding people from knowledge generation in any way, we have to take into account the socio-political context that shapes our criteria and try and discount it. This will mean that much more inclusivity, diversity, and open-mindedness is called for. That is not only a call based on epistemic prudence but is also a call to further justice.

A point of attention here is intellectual ability. For it may be thought that greater intellectual ability is a ground to give someone more epistemic power than another. The problem with this is that intellectual ability tends to be judged by those who are powerful, and that the criteria tend to have been set by those. Having a good education

and being a good learner may depend on fitting into a system, and knowing what that system requires. It is hard to defend that intellectual ability is always inherited – and yet, many people end up receiving the same level of education (and income) as did their parents.

So, can we structure society such that is epistemically just such that the norms delineated in the last chapter are maintained? Can we transpose the Longino's findings with regard to an epistemic community to society at large? How does epistemic power relate to other forms of powers in our world? While these other powers need to be kept in check, there is a case to be made that epistemic power sometimes is to be used to interfere with these powers in order to keep them in check. This is when journalists seek to check on the way political power is used, for instance, or when whistleblowers share knowledge despite its being a breach of trust, in order that a verdict can be cast by those whom they believe should know and be able to resist what is done in secret.

With respect to the last point it appears we should add one good use of epistemic power to the ones I have already documented. In the last chapter, I argued that intuitively, a proper use of epistemic power is: (1) to not lie, and to take care; (2) to give credit where credit is due; (3) to share one's knowledge proactively, in circumstances of trust; and (4) to be openminded, sharing one's knowledge and bestowing credibility in cases of doubt. Yet, there is another slogan, which captures the great importance of epistemic power in the political domain, which is: "To speak truth against power" (Williams, 2002). The fact that epistemic power can be used to resist political domination, or even economic oppression is quite an important insight. It shows that one may use epistemic power to correct for the abuse of other powers. Indeed, one important use of epistemic power is precisely to call out abuses of power and raise awareness of injustices.

9. Epistemic Justice in Politics: Representation

It may be thought that lesser emphasis is to be placed on the moral-political in the field of science than in the field of politics; and, conversely, that an emphasis on inclusion (in democratic deliberate processes) is a political requirement of representation rather than one stemming from the nature of knowledge. In this section I argue that in fact in the political domain epistemic justice too plays a key role precisely to underscore both why representation is important and why the bar for being heard should not be unduly high.

Iris Marion Young has focused on the importance of inclusion as a "normative ideal of democratic communication" (Young, 2000, p. 6). Inclusion is important, she maintains, as a democratic norm, as a matter of political equality and as a matter of showing respect, so that the decisions are accepted by others even if they do not agree with them (Young, 2000, p. 52). Young argues that it is because there are differences in values and in

life experiences that we cannot assume shared understandings and that we need to ensure that "internal" exclusion does not occur (Young, 2000, p. 55). Such internal exclusion is the informal exclusion that can happen as a result of the fact that the dominant group may be inclined to resist the influence of those who are not like-minded, not taking others' views seriously:

> Having obtained a presence in the public, citizens sometimes find that those still more powerful in the process exercise, often unconsciously, a new form of exclusion: others ignore or dismiss or patronize their statements and expressions. Though formally included in a forum or process, people may find that their claims are not taken seriously and may believe that they are not treated with equal respect. The dominant mood may find their ideas or modes of expression silly or simple, and not worthy of consideration. They may find that their experiences as relevant to the issues under discussion are so different from others' in the public that their views are discounted. I call these familiar experiences internal exclusion, because they concern ways that people lack effective opportunity to influence the thinking of others even when they have access to fora and procedures of decision-making (Young, 2000, p. 55)

Note that the kind of equal respect that Young here speaks of is *moral-epistemic* respect – what she points out is that the dominant group fails to see others as authoritative sources, or even as sources at all. The sense in which epistemic power is involved here is that there is a lack of epistemic power on the part of those who are deemed not to make sense, from the viewpoint of those who are socially powerful. The trouble is that when internal exclusion occurs, it appears epistemically justifiable, in so far as the epistemic criteria that are socially accepted as the correct criteria allow people to dismiss the views of others on the merits of these views.

What Young correctly pinpoints as a concern of justice is that the epistemic criteria that determine epistemic merit must make sense. So, she says, very concretely, that the mode of communication should not feature as a ground for excluding the content of what is said. Her plea is that even badly expressed arguments are arguments, even as we should all remain critical when it comes to filtering arguments for truths. I take this to be a point of epistemic justice, and that the Critical Theory Test, which importantly takes the perspective of the marginalized as the basis of the legitimacy of criteria, is to be used as a means of dismantling claims of epistemic superiority, which are both false and epistemically unjust, as they lack the basis that they need and exclude others.

As Young argues, it is because we must take others seriously, as equal members of society, that there is a responsibly to listen to others even if they do not present their arguments clearly or in a way that is immediately understandable to the rest of the community. This means that we need to ensure that we understand what others say, even where there is a lack of clarity, i.e. regardless of form:

> Participants in communicative democracy should listen to all modes of expression that aim to co-operate and reach a solution to collective problems (Young, 2000, p. 80)

What moves Young to focus on inclusion is that she thinks it is the best way, given the

power struggles that are, to achieve justice:

> In actually existing democracies there tends to be a reinforcing circle between social and economic inequality and political inequality that enables the powerful to use formally democratic processes to perpetuate injustice or preserve privilege. One means of breaking this circle, I argue, is to widen democratic inclusion. (Young, 2000, p.18)

Young's point is one that I am using to conclude my account of epistemic justice in this thesis, as it aptly indicates that epistemic justice is a means for achieving justice overall, while epistemic justice is one important element of it, alongside achieving other forms of justice – e.g. economic, political, gender, etc.

Two further points may be made on Young's work on inclusion. The first is that, while she appeals to inclusion, and to the ideal of equality, Young does not appeal to solidarity. Yet arguably one important aspect of the idea of a focus on power as an ability as opposed to a means of control or a form of domination is that there is room for theorizing the notions of empowerment, resistance and solidarity (Allen, 1999, p. 122, p. 126). I have argued in this thesis that there should be an emphasis on *solidarity* alongside *inclusivity* in ensuring epistemic justice. Both epistemic solidarity – e.g. sharing knowledge proactively, helping others to make their arguments such that these arguments will be heard - and inclusivity – listening to what others say, assessing the content of what is said with an open mind – are important in realizing the ethically proper use and the politically just allocation of epistemic power.

The second thing to note is that Young's position on inclusion seems to reinforce a key value in Longino's contextual empiricism, namely the emphasis on the importance that others are to be heard even and especially when their views differ. Even so, Young appears to be more radical. For the "bad" grounds which she cautions against, which may illicitly be used to exclude others, may appear to be reasonable from an epistemological viewpoint. They go further than what we would expect. While we have seen that the bar for inclusion should not be set too high, there are still to be standards for reasonableness; the standards for reasonableness that were adduced by Douglas seem preferable, in light of a concern with fact finding and the truth.

Here it is noteworthy that Young focused on political deliberative processes, where others than the learned/educated people *ought* to have a say, too, indeed. What affects people is what they should be informed about and should be given room to give input for. So, we could think that epistemic justice requires something different in the political deliberative domain than it requires in science, where the emphasis is less on representation and more on methodological rigor. This is not too surprising, so one may think, as more may hang on an inclusive representation of interests in politics than in science.

Conversely, however, there are moral and political reasons why the bar for participation in science must not be too high. That, too, is the take home message of

the works of Longino (2002), Douglas (2009), and Lacey (1999). It has already been argued that laypeople should be involved in setting the scientific agenda, because of its societal relevance (Kitcher, 2003). But in the actual practice of science, and more generally in our everyday epistemic practices, too, there is room for rethinking what criteria of epistemic authority are epistemically proper, i.e. epistemically just, and for considering the importance of solidarity and open(minded)ness.

Conclusion

In this chapter I have addressed two questions: first, how epistemic power is to be allocated in an epistemically just way, and second, why epistemic justice is a valuable means for evaluating our practices. I have used Longino's work to point to two things. The first is that even apparently neutral criteria that have been used in science for the generation of knowledge may be viewed as problematic. The second is that the fact that knowledge is social means that we should pay attention to both how epistemic power is allocated, and how our social epistemic practices are structured, which focused on a discussion on who should be included in the epistemic community, i.e. who should be regarded as qualified or credible.

What followed from that discussion is, on the one hand, that we need to scrutinize our criteria much more closely, whereby I suggested that we should institute and apply a Critical Theory Test for any claim to authority; on the other hand, what we need is to see that the knowledge generative community needs to be sufficiently open to criticism and inclusive of those who may not be thought to "qualify" as equal members, at first, if it is to screen for epistemic bias and if is to yield all the relevant truths that we should want to have within our society.

The conclusions of this chapter provide support for the norms that we earlier identified as the moral-ethical norms of epistemic justice: we need an ethos of epistemic justice that will give people the benefit of the doubt as argued in the previous chapter; and the bar for participation cannot be set too high. Even so, there need to be some criteria. Here my suggestion is that we should assess the ones that we have by subjecting them to critical tests that have a common core; namely that the knowledge and authority acclaimed – and conferred upon claims and people by the criteria in use – are to be backed independently of any particular interests being served. That is, they should be epistemically pure and just, by furthering no other ends than those of knowledge and justice, and hence, pivotally, by being motivated by truth and social equality. Epistemic justice is secured then through an ethos of epistemic justice and through the existence of social institutions that justly allocate epistemic power and that sanction violations of its use.

I have also discussed how epistemic justice relates to other forms of justice. Here

it is especially interesting to see what role may be played by epistemic empowerment, not only in terms of furthering epistemic justice but also in terms of eradicating other types of injustice(s); e.g. speaking truth to power. Conversely, epistemic domination is to be expected when other forms of justice are not had. The curtailment of others' use of epistemic power is one of the means of securing longstanding domination, and makes resistance to authoritarianism in politics much harder, as well as making any resistance harder for those who are socially marginalized versus those who are socially powerful. An important connection exists between securing epistemic justice and securing "justice overall", i.e. between securing knowledge and securing non-domination in all domains.

Conclusion

In this thesis I set out to show that there is a moral-epistemic norm that knowledge should be shared. In the course of my argument I have discussed the social nature of knowledge, the role and nature of testimony, and the nature and point of epistemic justice. In the final chapters, I articulated what epistemic justice means for our epistemic conduct, and have argued that knowledge sharing – rather than keeping knowledge to oneself – is the moral-epistemic default norm, while explaining the conditions needed for the norms of epistemic justice to be upheld.

I have explained and defended the claim that knowledge is strongly social. The main challenge is to explain that knowledge is social *without* denying that knowledge claims pertain to truth *or* debunking the status of truth as relative to a set of communal standards. I further argued that only a genealogical account of knowledge can provide such an understanding and that it has advantages over the justified true belief (or "JTB") conception of knowledge. This is because a genealogy does not reduce a social concept to it social function. It shows how it is not a contradiction to say that the concept of knowledge is not instrumental to our goals, even as the concept arises to meet the shared need to acquire truths and more truths than one person alone can muster using their onboard resources.

Like the "knowledge-first" approach, a genealogical account of knowledge can elucidate the point of the concept of knowledge via its role. But in addition, it can explain the concept's social nature whilst resisting reductions of the concept in terms of its function. Accordingly, it makes good on the claim by knowledge-first theorists that the concept of knowledge is indefinable, but also moves beyond it in explaining that there are socially determined criteria for the application of the concept, which may or may not be the correct ones for arriving at only or mostly truths.

The explanation for why the criteria for knowledge are socially to be set in some given way, namely such as to be "objective", crucially involves the pursuit of truths. Thus, the concept of knowledge is not reduced to social opinion or even to hypothetical or idealized agreement, even while we can explain that what passes as knowledge in a community, rightly or wrongly so, is relative to a community. Moreover, we can explain that it matters *both* whether a claim passes as knowledge within a community *and* whether or not it does so rightly or wrongly. The reason why it matters whether a claim passes as knowledge, in the first place, is that having this social status is required for it to be knowledge, whereby what rightly passes as knowledge is true.

That is, on my view there is no room for knowledge which is not possessed by a human or even the – forgotten or as yet to be (re)discovered – product of a human being, which merely exists as knowledge in some separate metaphysical realm. Instead of regarding knowledge as "objective" when it is "only of truths" and floating around in some metaphysical realm, I regard knowledge as that which human beings possess or have created for others to find, e.g. written in books (that are still physically accessible and readable, or at least known about).

Yet it also matters that what passes as knowledge does so rightly. For it matters that we have the right criteria and that they are applied correctly, because only then is what passes as knowledge within a community really knowledge: the concept of knowledge is *normative*, in that there is a presupposed link from the status of knowledge being accorded to a claim to its being the case that that claim is true.

Because of this, knowledge may be called "objective" – if one wants to use that term – when the criteria that are socially set for it are the ones that deliver truths when applied properly. However, we do not ordinarily claim that our knowledge is objective, for the reason that from within a practice, it is often assumed that the criteria used are the right ones. Instead, assuming the objectivity of the criteria for knowledge, we say that a claim is knowledge when it meets the criteria for passing as such.

One reason why passing as knowledge is a feasible criterion for knowledge – on condition that the criteria that are used for its assessment are in order – is that, by being shared and assessed for its truth from a variety of viewpoints, the claim has passed sufficient tests for truth, and hence qualifies properly as knowledge. Having assumed that the criteria are proper, however, such a test merely reflects the correctness of the criteria being applied to it, while it also assumes that the background knowledge used was a sufficiently safe and stable basis from which to judge the content of the claim for its truth. That these issues also require scrutiny for claims to properly pass as knowledge is too much to ask in the context of everyday life.

Yet, to ascertain (with greater likelihood) that what properly passes as knowledge is knowledge (i.e. is true), we may have to inquire into the criteria and the background assumptions, including the knowledge that is used in the process of assessing claims for their truth; whereby "claims" here includes hypotheses and theories. This larger project, however, is to be undertaken in science. It is why "scientific knowledge" is held to higher standards, and also why fewer claims pass as knowledge in science than in everyday life. While I am concerned with the everyday notion of knowledge, I nonetheless note this point, because and to the extent that what passes as knowledge in the everyday sense, as I present it, is modelled on scientific knowledge being attained through falsification. Another reason for noting it is there is also a "trickle down" effect, in the sense that science can yield new insights that can also debunk everyday knowledge claims, or can do so by assessing them more thoroughly. In addition, at least some of our everyday knowledge owes its existence to science.

I also argued that it follows from the social nature of knowledge that knowledge sharing is important to its generation and not merely its transmission. Testimony, conceived broadly, thus needs to be reconceptualized in two ways. It has to be seen, so I argue, as fundamental to knowledge; and moreover, as generative of knowledge, in the sense that, for any item of propositional knowledge, testimony has an important role to play in its social validation and accreditation as knowledge.

I have argued that *if* knowledge is social – in the sense that shared standards are to be upheld for it while social epistemic practices serve to generate and disseminate it, then epistemic justice is required for knowledge to be "objective", i.e. for social epistemic practices to be successful in helping us to acquire truths. The intuitive idea which I defended is that epistemic justice aligns truthfulness and trust, in aligning through morality, the social and the epistemic. Within our social epistemic practices, we can rely on one another, and on the criteria that are socially set for such reliance, even while their aim is the normative and ambitious one of securing truths.

I have argued that Williams's genealogy of truthfulness is to be taken as a genealogy for a broader notion, namely that of epistemic justice. I argued this on the basis of the fact that his conception of truthfulness is both too broad and too narrow. It is too broad because it involves considerations of justice. It is too narrow, because it would appear to focus exclusively on the ethical: on individual conduct and on people sustaining interpersonal relations. This is too narrow since there are political aspects to ensuring that truth is secured via social processes. These political aspects are brought out by Fricker's work.

I have argued that understanding the norms and responsibilities that prevail in testimonial practice provides an argument for knowledge's sociality and indeed for epistemic injustice. Even the most refined theories of testimony cannot explain these. I have argued that the reason is that they are not yet free of a methodological individualism that problematizes testimony's pivotal role in the generation of knowledge. Contrary to current conceptions, which focus only on how knowledge which already exists may be transmitted, and whether or not testimony is a self-standing "source" of knowledge, I argue that testimony is a social epistemic practice, and that its function is that it can add to or detract from the epistemic status of claims that are presented as (putative) knowledge. As I explained, only this kind of conception of knowledge can explain the "critical uptake" of testimony (see Faulkner, p. 19).

My claim throughout this thesis has been that a positive conception of epistemic justice is needed, which centers on epistemic power, conceived as an ability to exert epistemic influence, including the ability to enable or disable others from doing so. That ability is given by knowledge, on the one hand, and credibility – i.e. perceived rational authority – on the other. My conception of epistemic justice and of epistemic power has been derived from Fricker's work, even while Fricker was focused first and foremost on the undue influence of social power, and regards epistemic justice as a corrective virtue, rather than an ideal or principle from which norms can be derived.

While there are important connections between epistemic power and other forms of power, I have argued that in understanding epistemic (in)justice, we must focus on epistemic power and its use and allocation, rather than focus on the workings of social power that unduly influence epistemic outcomes. The point of doing so is, first of all, that it can help us to identify all and only those cases that are properly termed epistemic injustices or epistemic wrongs. Secondly, a focus on epistemic power is needed to show that there is indeed, as Fricker claimed, a specific dual moral-epistemic normativity that pertains to claims of epistemic (in)justice. Using the notion of epistemic power, I have argued that the notion of epistemic justice cannot be reduced or debunked, to purely issues of either justice or epistemology, as some critics have claimed.

Focusing on Fricker's work, I have clarified that an epistemic wrong involves a direct abuse of epistemic power, while in contrast, an epistemic injustice centers on the allocation, or distribution, of epistemic power. An epistemic injustice, on my view, is constituted by a person or group of people suffering an epistemic disadvantage compared to others, while there is no good epistemic justification for that disadvantage. While an epistemic injustice, as I define it, means one has a relative lack of epistemic power vis-à-vis others, it is also simultaneously a moral affront. It is an injustice, as opposed to merely "bad luck", since any differences in epistemic power ought to have an epistemic justification, while the default is that, as moral equals, no differences in any kind of power should exist without justification. Thus, I explain that even while an epistemic injustice does not describe an epistemic action, and may not follow from anyone's action in particular, an epistemic injustice is nonetheless a moral "bad" that is specifically epistemic.

So, the two categories of epistemic wrongs and epistemic injustices can overlap, in the sense that one's use of epistemic power, which constitutes a wronging, may be the cause of an epistemic injustice, when one wrongly allocates or fails to allocate epistemic power to others. At the same time, epistemic injustices can have other causes too, which are non-epistemic. The workings of social power, as Fricker has emphasized, are quite often the cause for epistemic injustices. Yet, what is important, is that the failing that constitutes an epistemic injustice is epistemic: there is a failing of the moral-epistemic norm that epistemic power ought to be allocated properly, whereby differences in epistemic power are to be justified on the basis of an epistemic rationale. This is to say, as I argue – against Bohman's position – that an injustice does not reduce to a case of social injustice, because of the fact that other than epistemic powers are involved. For, the form of domination in question, which the label of epistemic injustice pinpoints is epistemic domination.

Both this specific type of domination, and the fact that the sociality of knowledge requires an alignment between truthfulness and trust have led me to conclude that epistemic justice pertains to a specific kind of moral-epistemic normativity that cannot be reduced.

Epistemic justice ought to be incorporated into a framework of justice overall that

is ideal-theoretical and yet pluralist in nature. It is pluralist, as there are various forms of justice that exist alongside, each regulating a specific form of social power – physical, political, economic, legal, epistemic, sexual and the power of social identity. It is ideal-theoretical, in that there is the aspiration to build a theory that spells out what justice requires, for each form of power, whereby the ultimate aim is to show how an optimal situation of "justice overall" can be achieved, whereby each form of power is used and allocated as it should be.

The idea that such a conception of justice overall is possible is supported by the thought that if each form of power is used for the purposes for which it should be used, then these powers are also contained within their realm. That is, if and to the extent that the different forms of justice are each achieved, they can well co-exist alongside, and will keep one another in balance. It is when injustice of one form exists that other types of injustice will quickly arise too, because power of one form, when unrestrained is not contained within its proper confines: it will translate to power of other forms, defying the rules of (the proper basis for) power allocation in that "sphere of justice", and undermining the rule that use of a power is subject to the conditions of its allocation, such as the need to use it well and to render account when asked to do so.

Epistemic power is described as one form of social power alongside other forms of social power, such as physical power, economic power, political power, legal power and sexual power, and the power of social identity. Each of these forms of power, I maintain, is regulated by its own rationale; much in the way Walzer has argued that each social good is to be distributed according to its own social meaning. Yet I maintain that the regulation of any type of power concerns both its use and its allocation, and hence I do not condone the idea that justice is centrally about distribution. Justice is not only a political "thin" notion, as rights-based, deontological theories of justice have theorized it, but is also an ethical notion. This is so in part because the ethical use of power impacts on the distribution, such that the ethical and political cannot be clearly separated, and in part because acting justly is ethically required; so, justice is needed for ethical guidance, too.

In the closing chapters of the thesis (VII and VIII) I have articulated what epistemic justice involves, by providing a set of substantive default norms for the use of epistemic power and delineating the scope of these norms. Their having a "default" status means that they presuppose that the circumstances of trust obtain. While knowledge sharing may be a way of instituting and increasing such trust, it is also a default norm, because the investment that is being made, too, requires that it can be reasonably assumed that others are likely to respond well to that trust. By providing these norms and delineating their scope which is also a function of epistemic justice, I have shown that epistemic justice can serve as the rationale for the moral-epistemic norms that (are to) govern our social epistemic practices, such as testimony.

It is a further question whether the norms that are spelled out by epistemic justice are indeed our norms; that is, whether they are lived up to in practice to a sufficient extent.

This is not to deny their relevance, however, but to state that epistemic justice provides us with a rationale that can be used to assess to what extent our current epistemic practices may be deemed epistemically just.

I described that which epistemic justice requires when it comes to the allocation of epistemic power. I have argued that the question of what epistemic criteria are proper is important for delineating the proper use of epistemic power. I have also argued that an understanding of what a proper allocation of epistemic power involves, and what conditions are needed to achieve it, helps us evaluate to what extent our current epistemic practices are epistemically just. In addition, I have argued that epistemic justice as an ideal by which to evaluate our epistemic practices is needed given the social conception of knowledge which I defended. For we need a way of securing that sociality helps to further – as opposed to its being a threat to – the objectivity that is required for knowledge.

Regarding the criteria for knowledge and authority, I have argued that while they are needed in order to enable contestation, they also need to be open to contestation. This is to ensure that these criteria are not rigged, or the very notion of epistemic justice itself appropriated by the socially powerful. As the case made by Longino illustrated, even apparently innocuous criteria, and especially criteria that have been accepted for a long time as the only correct ones, may too readily exclude putative sources and viewpoints. As a remedy, I have pointed to Bernard Williams's Critical Theory Test. It creates room for questioning and challenging power distributions that are accepted from within a practice as simply the status quo. What that test embodies is a commitment to the idea that all power inequalities should be based on reasons that are independent of that power distribution being the status quo – the fact that all accept the differences in power is not enough. Instead, it is necessary that the normative force of the reasons that legitimate differences in power, if they are at all legitimate, should be independent of those who have greater power than others exercising it, e.g. by appealing to their authority.

Applied to epistemic criteria it means that the criteria ought to further epistemic justice – and hence truth – without unduly excluding others, and that claims of those who enjoy authority on the basis of these very criteria that other views, methods, positions or persons are rightly excluded from the practice should not appeal to their own authority, but instead be based on arguing that the exclusion is epistemically just or proper. The key point has been that the questions of criteria are important from the viewpoint of epistemic justice, because they bear on who is to be included and who is to be excluded as a (putative) epistemic source, and on what come to pass as knowledge in a society, rightly or wrongly.

Second, epistemic justice as an ideal allows us to evaluate whether our social epistemic practices are epistemically just. Furthering epistemic justice is important, as I have argued, both for achieving knowledge and for achieving justice overall. Feminist philosophy of science has developed insights into how our social epistemic practices should be shaped. These arguments ought to be cast in terms of furthering epistemic justice; and along with

it, both knowledge and justice. I have also applied this to the domain of politics, where epistemic power, too should be used well and allocated well, and where we may also need to rethink the criteria that are used for proper participation such that, as Young has insisted, they do not lead to internal epistemic exclusion.

List of Works Cited

Allen, Amy. (1999). *The Power of Feminist Theory: Domination, Resistance, Solidarity*. Boulder, Colo.: Westview Press.

Anderson, Elizabeth. (2012). Epistemic Justice as a Virtue of Social Institutions. *Social Epistemology, 26*(2), 163–173.

Audi, Robert. (1997). The Place of Testimony in the Fabric of Knowledge and Justification. *American Philosophical Quarterly, 34*(4), 405–422.

Audi, Robert. (2013). Testimony as a Social Foundation of Knowledge. *Philosophy and Phenomenological Research, 87*(3), 507–531.

Austin, John L. (1946). Other Minds. Reprinted in Austin, J. (1979). Urmson, J. & Warnock, G. (Eds.), *Philosophical Papers* (pp. 76-117). Oxford University Press.

Berenstain, Nora. (2016). Epistemic Exploitation. *Ergo: An Open Access Journal of Philosophy, 3*(22), 569–590.

Beeby, Laura. (2011). A Critique of Hermeneutical Injustice. *Proceedings of the Aristotelian Society, 111*(3), 479–486.

Bohman, James. (2012). Domination, Epistemic Injustice and Republican Epistemology. *Social Epistemology, 26*(2), 175–187.

Burge, Tyler (1993). Content Preservation. *The Philosophical Review*, 102(4), 457–488.

Carrabregu, Gent. (2016). Habermas on Solidarity: An Immanent Critique. *Constellations, 23*(4), 507–522.

Child, William. (1994). *Causality, Interpretation and the Mind*. Oxford: Clarendon Press.

Coady, David. (1994). *Testimony: A philosophical study*. Oxford: Oxford University Press.

Coady, David. (2010). Two concepts of epistemic injustice. *Episteme: A Journal of Social Epistemology*, 7(2), 101–113.

Craig, Edward. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Oxford University Press.

Craig, Edward. (2007). Genealogies and the State of Nature. In Thomas, Alan (Ed.), *Bernard Williams* (pp. 181-200). Cambridge: Cambridge University Press.

Dancy, Jonathan. (1995). Arguments from Illusion. *Philosophical Quarterly*, *45*(181), 421–438.

Darwall, Stephen. (2004). Respect and the Second-Person Standpoint. *Proceedings and Addresses of the American Philosophical Association, 78*(2), 43-59.

Davidson, Donald. (1973). On the Very Idea of a Conceptual Scheme. *Proceedings and Addresses of the American Philosophical Association, 47*, 5–20.

DeRose, Keith. (1992). Contextualism and Knowledge Attributions. *Philosophy and phenomenological research 52*(4), 913–929.

DeRose, Keith. (1995). Solving the Skeptical Problem. *The Philosophical Review, 104*(1), 1–52.

Descartes, René. (2000). *Meditations On First Philosophy*. South Bend: Infomotions.

Douglas, Heather E. (2009). *Science, Policy, and the Value–Free Ideal*. Pittsburgh: University of Pittsburgh Press.

Douglas, Heather E. (2016). *5th Descartes Lectures: Science, Values and Democracy*, Tilburg University (as yet unpublished)

Dretske, Fred I. (1971). Conclusive Reasons. *Australasian journal of Philosophy, 49*(1), 1-22.

Dretske, Fred I. (1983). *Knowledge & the flow of information*. Cambridge, Mass.: MIT Press.

Faulkner, Paul. (2000). The Social Character of Testimonial Knowledge. *The Journal of Philosophy*, 97(11), 581–601.

Faulkner, Paul. (2011). *Knowledge on Trust*. Oxford: Oxford University Press.

Faulkner, Paul. (2014). The Moral Obligations of Trust. *Philosophical Explorations, 17*(3), 332–345.

Foucault, Michel. (1972). *The Archeology of Knowledge and the Discourse on Language*. New York: Pantheon Books.

Foucault, Michel. (1979). *Discipline and Punish: the Birth of the Prison*. New York: Vintage Books, a division of Random House.

Forbes, Graeme. (1984). Nozick on Scepticism. *Philosophical Quarterly, 34*(134), 43–52.

Fraser, Nancy. (1989). *Unruly Practices: Power, Discourse, and Gender in Contemporary Social Theory*. Minneapolis: University of Minnesota Press.

Fraser, Nancy. (2000). Recognition without Ethics? In Garber, Marjorie & Hanssen, Beatrice & Walkowitz, Rebecca L. (Eds.). *The Turn to Ethics* (pp. 95-127). New York: Routledge.

Fraser, Nancy. (2005). *Reframing Justice*. Assen: Koninklijke Van Gorcum.

Fraser, Nancy. (2007). Identity, Exclusion, and Critique. *European Journal of Political Theory, 6*(3), 305–338.

Fricker, Elizabeth. (1994). Against Gullibility. In Chakrabarti A. & Matilal B.K. (Eds.), *Knowing from Words* (pp. 125–161). Dordrecht: Springer.

Fricker, Elizabeth. (2006). Second-Hand Knowledge. *Philosophy and Phenomenological Research, 73*(3), 592–618.

Fricker, Elisabeth. (2015). Know First, Tell Later: The Truth about Craig on Knowledge. In Greco J. & Henderson D. (Eds.), *Epistemic Evaluation: Purposeful Epistemology* (pp. 46-86). Oxford: Oxford University Press

Fricker, Miranda. (2003). Epistemic Injustice and a Role for Virtue in the Politics of Knowing. *Metaphilosophy, 34*(1–2), 154–173.

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

Fricker, Miranda. (2011). Rational Authority and Social Power: Towards a Truly Social Epistemology. In Goldman, Alvin I. & Whitcomb, Dennis (Eds.), *Social Epistemology: Essential Readings* (pp. 54-71). Oxford: Oxford University Press. (Reprinted from *Proceedings of the Aristotelian Society, 1998: 98*(2), 159-177)

Fricker, Miranda. (2013). Epistemic Justice as a Condition of Political Freedom? *Synthese: An International Journal for Epistemology, Methodology and Philosophy of Science, 190*(7), 1317–1332.

Gallie, W. B. (1955). Essentially Contested Concepts. *Proceedings of the Aristotelian Society, 56*(1)*,* 167–198.

Gettier, Edmund L. (1963). Is Justified True Belief Knowledge? *Analysis, 23*(6), 121–123.

Goldberg, Sanford C. (2011). If That Were True I Would Have Heard It By Now. In Goldman, Alvin I. & Whitcomb, Dennis (Eds.), *Social Epistemology: Essential Readings* (pp. 92–109). Oxford: Oxford University Press.

Goldman, Alvin I. (1967). A Causal Theory of Knowing. *Journal of Philosophy, 64*(12), 355-372.

Goldman, Alvin I. (1999). *Knowledge in a Social World*. Oxford: Clarendon Press.

Goldman, Alvin I. (2001). Social Routes to Belief and Knowledge. *The Monist*, *84*(3), 346–367.

Goldman, Alvin I. (2002). Knowledge and Social Norms. *Science*, *296*(5576), 2148–2149.

Greco, John, & Sosa, Ernest. (1999). *The Blackwell Guide to Epistemology*. Malden, Mass.: Blackwell.

Greco, John. (2007). The Nature of Ability and the Purpose of Knowledge. *Philosophical Issues, 17*(1)*,* 57–69.

Greco, John. (2009). Knowledge and Success from Ability. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition, 142*(1), 17–26.

Grice, Paul. (1957). Meaning. *The Philosophical Review, 66*(3), 377–388.

Grice, Paul. (1989). Logic and Conversation. In Grice, P. (Ed.), *Studies in the Way of Words* (pp. 22–41). Cambridge, Mass.: Harvard University Press.

Habermas, Jurgen. (1993). Justice and Solidarity. In Fisk, Milton (Ed.), *Justice* (pp. 89-102). Atlantic Highlands: Humanities Press International.

Haraway, D. (1988). Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Feminist Studies, 14*(3), 575–599.

Haslanger, Sally. (1999). What Knowledge Is and What It Ought to Be: Feminist Values and Normative Epistemology. *Noûs, 33*(s13), 459-480.

Henderson, David. (2011). Gate-Keeping Contextualism. *Episteme, 8*(1), 83–98.

Henderson, David & Greco, John (Eds.). (2015). *Epistemic Evaluation: Purposeful Epistemology.* Oxford: Oxford University Press. Kindle Edition.

Hinchman, Edward S. (2005). Telling as Inviting to Trust. *Philosophy and Phenomenological Research, 70*(3), 562–587.

Hinton, J. M. (1967). Visual Experiences. *Mind*, *76*(302), 217–227.

Hinton, J. M. (1973). *Experiences*. Oxford: Oxford University Press.

Hong, Lu & Page, Scott. (2012) Some Microfoundations of Collective Wisdom. In Landemore, Helene & Elster, John (Eds.), *Collective Wisdom: Principles and Mechanisms* (pp. 56-71). Cambridge: Cambridge University Press.

Hookway, Christopher. (2010). Some Varieties of Epistemic Injustice: Reflections on Fricker. *Episteme, 7*(2), 151–163.

Hornsby, Jennifer. (1997). The Presidential Address - Truth: The Identity Theory. *Proceedings of the Aristotelian Society, 97*(1), 1-7.

Horwich, Paul. (1990). *Truth*. Oxford: Oxford University Press.

Hume, David. (2000). *An Enquiry Concerning Human Understanding*. South Bend: Infomotions.

James, Susan. (2005). Rights, Moral and Enforceable: A Reply to Saladin Meckled-Garcia. *Proceedings of the Aristotelian Society, 105*(1), 149–153.

Kaplan, Mark. (2003). Who Cares What You Know? *The Philosophical Quarterly, 53*(210), 105–116.

Kelp, C. (2011). What's the Point of "Knowledge" Anyway? *Episteme, 8*(1), 53–66.

Kitcher, Philip. (2003). *Science, Truth, and Democracy*. Oxford: Oxford University Press.

Kornblith, Hilary. (1999). In Defense of a Naturalized Epistemology. In Greco, John & Sosa, Ernest (Eds.), *The Blackwell Guide to Epistemology* (pp. 158-169). Malden, Mass.: Blackwell.

Kusch, Martin. (2002). *Knowledge by Agreement: the Programme of Communitarian Epistemology*. Oxford: Oxford University Press.

Kusch, Martin & McKenna, Robin. (2018). The Genealogical Method in Epistemology. *Synthese*.

Lacey, H. (1999). *Is Science Value Free? Values and Scientific Understanding*. New York: Routledge.

Lackey, Jennifer. (2006). Book Review: Knowledge by Agreement: The Programme of Communitarian Epistemology. *Philosophy and Phenomenological Research, 72*(1), 235–238.

Lackey, Jennifer. (2008). *Learning from Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.

Landemore, Hélène. (2013). *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton, N.J.: Princeton University Press.

Lewis, David. (1996). Elusive Knowledge. *Australasian Journal of Philosophy, 74*(4), 549–567.

Locke, John. (1788). *An Essay Concerning Human Understanding: In four books.* (The eighteenth ed., Eighteenth century collections online). London: Printed for J.F. and C. Rivington, T. and T. Payne, L. Davis, B. and B. White, S. Crowder and 20 others in London.

Longino, Helen E. (1990). *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton: Princeton University Press.

Longino, Helen. (1995). Gender, Politics, and the Theoretical Virtues. *Synthese, 104*(3), 383-397.

Longino, Helen E. (2002). *The Fate of Knowledge*. Princeton, N.J.: Princeton University Press.

MacIntyre, Alasdair. (1985). *After Virtue: A Study in Moral Theory* (2nd corrected ed. with postscript). London: Duckworth.

MacIntyre, Alasdair. (1995). Truthfulness, Lies, and Moral Philosophers: What Can We Learn from Mill and Kant? *The Tanner Lectures on Human Values, 16*, 307–361.

Martin, Michael G. F. (1997). The Reality of Appearances. In Sainsbury, R.M. (Ed.), *Thought and Ontology*. Milan: FrancoAngeli. Reprinted in: Byrne, Alex & Logue, Heather (2009). *Disjunctivism: Contemporary Readings* (pp. 91-117). Cambridge, Mass.: MIT Press.

Marx, Karl & Engels, Friedrich. (1983). *Manifesto of the Communist Party*. New York: International Publishers.

McDowell, John. (1982). Criteria, Defeasibility, and Knowledge. *Proceedings of the British Academy, 68*: 455–479.

McDowell, John (1994). Knowledge by Hearsay. In *Meaning, knowledge, and reality* (pp. 414–445). Cambridge, Mass.: Harvard University Press.

McMyler, Benjamin. (2011). *Testimony, Trust, and Authority*. New York: Oxford University Press.

Mill, John Stuart. (1910). *Utilitarianism, Liberty, and Representative Government*. London: J.M. Dent & Sons.

Mill, John Stuart & Collini, S. (Eds.). (1989). *On Liberty and Other Writings*. Cambridge: Cambridge University Press.

Miracchi, Lisa. (2015). Competence to Know. *Philosophical Studies, 172*(1), 29–56.

Misak C. J. (1991/2004). *Truth and the End of Inquiry: A Peircean Account of Truth*. Oxford University Press.

Moran, Richard. (2005). Problems of Sincerity. *Proceedings of the Aristotelian Society, 105(1),* 325–345.

Moran, Richard. (2006). Getting Told and Being Believed. In Lackey, Jennifer & Sosa, Ernest (Eds.), *The Epistemology of Testimony* (pp. 272–306). Oxford: Clarendon Press.

Moran, Richard. (2013). Testimony, Illocution and the Second Person. *Aristotelian Society Supplementary Volume, 87*(1), 115–135.

Nietzsche, Friedrich W. (2001). *The Gay Science: With a Prelude in German Rhymes And An Appendix Of Songs*. Cambridge: Cambridge University Press.

Nozick, Robert. (1981). *Philosophical Explanations*. Cambridge, Mass.: Belknap Press of Harvard University Press.

Nozick, Robert. (2013). *Anarchy, State and Utopia.* (2nd rev. ed.). New York, NY: Basic Books.

O'Neill, Onora. (2009). Ethics for Communication? *European Journal of Philosophy, 17*(2), 167–180.

On, Bat-Ami Bar. (1993). Marginality and Epistemic Privilege. In Alcoff, Linda & Potter, Elizabeth (Eds.). *Feminist Epistemologies* (pp. 83–101). New York: Routledge.

Peacocke, Christopher. (1992). *A Study of Concepts*. Cambridge, Mass.: MIT Press.

Popper, Karl R. (1979). *Objective Knowledge: An Evolutionary Approach* (Rev. ed. repr. with corr. and a new appendix 2nd ed.). Oxford: Clarendon Press.

Popper, Karl. R. (1978). Three Worlds. The Tanner Lecture on Human Values. University of Michigan, Ann Arbor.

Prichard, H. A. 1950. *Knowledge and Perception*. Oxford: Clarendon Press.

Pritchard, Duncan (2003). Virtue Epistemology and Epistemic Luck. *Metaphilosophy*, 34 (1–2), 106–130.

Quine, Willard van Orman. (1951). Main Trends in Recent Philosophy: Two Dogmas of Empiricism. *The Philosophical Review, 60*(1), 20-43.

Quine, Willard van Orman. (1970/1986). *Philosophy of Logic* (2nd ed.). Cambridge, Mass.: Harvard University Press.

Quine, Willard van Orman & Ullian, J. (1978). *The Web of Belief* (2nd ed.). New York: Random House.

Radford, Colin. (1966). Knowledge: By Examples. *Analysis, 27*(1)*, 1–11.*

Rawls, John. (1951). Outline of a Decision Procedure for Ethics. *The Philosophical Review, 60*(2), 177–197.

Rawls, John. (1971/1999). *A Theory of Justice* (Rev. ed.). Oxford: Oxford University Press.

Rawls, John. (2001). *Justice as Fairness: A Restatement*. Cambridge, MA: Belknap Press of Harvard University Press.

Rorty, Richard (1984). Solidarity or Objectivity? *Post-analytic Philosophy*, *3*, 5–6.

Rorty, R. (1998). Is Truth a Goal of Inquiry? Donald Davidson versus Crispin Wright. In *Truth and Progress*: *Philosophical Papers, Volume 3* (pp. 19-43). Cambridge: Cambridge University Press.

Schantz, Richard & Seidel, Markus. (2011). *The Problem of Relativism in the Sociology of (Scientific) Knowledge*. Berlin: De Gruyter.

Schmitt, Frederick F. (1994). *Socializing Epistemology: The Social Dimensions of Knowledge*. Lanham, Md.: Rowman & Littlefield.

Sen, Amartya K. (1967). Isolation, Assurance and the Social Rate of Discount. *The Quarterly Journal of Economics*, *81*(1), 112–124.

Sen, Amartya. (2009). *The Idea of Justice*. London: Allen Lane/Penguin Books.

Shapin, Steven. (1994). *A Social History of Truth: Civility and Science in Seventeenth-Century England*. Chicago: University of Chicago Press.

Simmons, John A. (2010). Ideal and Nonideal Theory. *Philosophy & Public Affairs, 38*(1), 5-36.

Snowdon, Paul. (1980). Perception, Vision, and Causation. *Proceedings of the Aristotelian Society, 81,* 175–192.

Snowdon, Paul. (1990). The Objects of Perceptual Experience. *Proceedings of the Aristotelian Society Supp. Vol. 64*: 121– 150.

Sosa, Ernest. (2007). *A Virtue Epistemology: Apt belief and Reflective Knowledge, Volume 1*. Oxford: Oxford University Press.

Stevenson, Leslie. (1993). Why Believe What People Say? *Synthese, 94*(3), 429–451.

Strawson, Peter F. (1977). II — Scruton and Wright on Anti-Realism Etc. *Proceedings of the Aristotelian Society*, *77*(1), 15–22.

Strawson, Peter F. (1992). *Analysis and Metaphysics: An Introduction to Philosophy*. Oxford: Oxford University Press.

Tarski, Alfred. 1944. The Semantic Conception of Truth and the Foundations of Semantics. *Philosophy and Phenomenological Research*, *4*(3), 341-376.

Taylor, Charles. (1995). Overcoming Epistemology. *Philosophical Arguments*. Cambridge, Mass.: Harvard University Press.

Thomas, Alan. (2018). Should Epistemic Injustices be Redressed by the "Corrective Virtues"? In Mark LeBar (Ed.), *Justice: the Virtues* (pp. 209–236). Oxford: Oxford University Press.

Thomas, Alan. (n.d.). Stags, Hares and Knowledge: A Genealogy of the Knowledge System as a 'Mutual Assurance Game'.

Walzer, Michael. (1983). *Spheres of Justice: A Defence of Pluralism and Equality* (Repr. ed.). Oxford: Basil Blackwell.

Wartenberg, Thomas E. (1990). *The Forms of Power: From Domination to Transformation.* Philadelphia: Temple University Press.

Welbourne, Michael. (1986). *The Community of Knowledge* (Scots Philosophical Monographs, no. 9). Aberdeen: Aberdeen University Press.

Welbourne, Michael. (2014). *Knowledge*. Hoboken: Taylor & Francis. (Original work published in 2001)

Williams, Bernard. (1970/1973). Deciding to believe. In Bernard Williams (Ed.), *Problems of the Self*. Cambridge University Press, pp. 136–51.

Williams, Bernard. (1978). *Descartes: The Project Of Pure Enquiry*. Hassocks: The Harvester Press.

Williams, Bernard. (2002). *Truth and Truthfulness: an Essay in Genealogy*. Princeton N.J.: Princeton University Press.

Williams, Michael. (2001). Contextualism, Externalism and Epistemic Standards. *Philosophical Studies, 103*(1), 1–23.

Williams, Michael. (2015). Knowledge in Practice. In Greco J. & Henderson D. (Eds.), *Epistemic Evaluation: Purposeful Epistemology* (pp. 161-186). Oxford: Oxford University Press.

Williamson, Timothy. (2000). Skepticism, Semantic Externalism, and Keith's Mom. *The Southern Journal of Philosophy, 38*(S1), 149–158.

Williamson, Timothy, (2002). *Knowledge and Its Limits*. Oxford: Oxford University Press.

Wittgenstein, Ludwig & Anscombe, Gertrude E. M. (Eds.). (2007). *Philosophical Investigations: the German Text with a Revised English Translation* (3rd repr. ed.). Malden, MA: Blackwell.

Wright, C. (1992). *Truth and Objectivity*. Cambridge, Mass.: Harvard University Press.

Young, Iris Marion. (1990). *Justice and the Politics of Difference*. Princeton, N.J.: Princeton University Press.

Young, Iris Marion. (2000). *Inclusion and Democracy.* Oxford: Oxford University Press.

Zagzebski, Linda. (1999). What is Knowledge? In Greco and Sosa (Eds.), *The Blackwell Guide to Epistemology* (pp. 92–116). Malden, Mass.: Blackwell.

Zagzebski, Linda. (2012). *Epistemic Authority: A Theory of Trust, Authority, and Autonomy in Belief*. Oxford: Oxford University Press.

This thesis addresses a number of interrelated questions: what is knowledge, how can it be gained through testimony, and what, if any, are the norms that pertain to knowledge sharing? It argues that knowledge can be understood as related to a social phenomenon without relativizing truth to community agreement. What is needed is to understand that the concept of knowledge is indefinable and irreducible, as "knowledge first" theorists have contended, while we can elucidate the concept by using the genealogical method.

It is argued, further, that the genealogical elucidation of the concept brings out its social role and nature. In turn, the social understanding of knowledge yields a new dynamic understanding of testimony as a generative source of knowledge. Contrary to the two mainstream views of evidentialism and assurantism and hybrids, it is argued that testimony is a practice through which knowledge is co-created.

While we together shape and determine what claims pass as knowledge, we ought to ensure that what comes to pass as knowledge really is so. Thus, the social nature of knowledge and the generative role of testimony serve to explain why there are norms that guide our epistemic conduct, and why our epistemic practices ought to be such as to ensure the correct criteria for knowledge. It is argued that these norms follow from epistemic justice, positively conceived.

Epistemic justice is conceptualized as the proper use and allocation of epistemic power, and as fitting into a pluralist theory of justice. The norms for epistemic conduct and the criteria for knowledge and authority it yields are investigated in the final chapters of this thesis.

Machteld C. M. Geuskens studied philosophy (BA & MPhil) at University College London and Birkbeck College London, and law (LLB & LLM) at Erasmus University. In 2012, she obtained a personal grant from NWO to pursue a PhD in philosophy at the Tilburg School of Humanities and Digital Sciences. Between 2016-2018 she worked as a lecturer at Erasmus School of Law.