University of Edinburgh

Department of Philosophy

Engineering anti-individualism: a case study in social epistemology

Eric T. Kerr

This dissertation is submitted in partial fulfilment of the requirements for
the Doctorate of Philosophy in Philosophy

May 2013

DECLARATION OF AUTHENTICITY

I, Eric Thomson Kerr, the undersigned, declare that all material presented to the University of Edinburgh is my own work, or fully and specifically acknowledged wherever adapted from other sources. This work has not been previously formed as the basis for the award of any degree, diploma or similar title at this or any other university.

…………………………………… Date……………………

Abstract

This dissertation is a contribution to two fields of study: applied social epistemology and the philosophy of technology. That is, it is a philosophical study, based on empirical fieldwork research, of social and technical knowledge. Social knowledge here is defined as knowledge acquired through the interactions between epistemic agents and social institutions. Technical knowledge is here defined as knowledge about technical artefacts (including how to design, produce, and operate them). I argue that the two must be considered collectively both in the sense that they are best considered in the light of collectivist approaches to knowledge and in the sense that they must be considered together as part of the same analysis. An analysis solely of the interactions between human epistemic agents operating within social institutions does not give adequate credit to the technological artefacts that help to produce knowledge; an analysis of technical knowledge which does not include an analysis of how that technical knowledge is generated within a rich and complex social network would be similarly incomplete. I argue that it is often inappropriate to separate analyses of technical knowledge from social knowledge and that although not all social knowledge is technical knowledge, all technical knowledge is, by definition, social. Further, the influence of technology on epistemic cultures is so pervasive that it also forms or 'envelops' what we consider to be an epistemic agent.

Nomenclature

JTB	The Justified True Belief Account of knowledge
CAT	Collective Account of Technical artefacts
CCDSE	Constructivist Collectivist Descriptive Social Epistemology
DHB	The Deepwater Horizon Blowout, or the Deepwater Horizon oil spill of 2010
DNA	Dual Nature Account of technical artefacts
PET	Philosophy of Engineering and Technology
PLA	Production Logging Analysis
PLT	Production Logging Tool
RINSE	Rationalist Individualist Normative Social Epistemology
SE	Social Epistemology
SP	The Strong Programme in Sociology of Scientific Knowledge
SSK	Sociology of Scientific Knowledge
STS	Science, Technology, and Society Studies

TABLE OF CONTENTS

Chapter 1

Deepwater Horizon and technical knowledge

1.1 Outline of dissertation

It may be said that no other industry has had as significant an effect on the social and technological history of the 20th century as the oil industry. It may also be said that that which significantly affects social and technological history tends to affect the human condition and human concepts. Knowledge is a human concept and as such is not immune from such change. Despite this plausible suggestion, epistemologists have not paid close attention to social and technological history and have instead searched for general principles that govern when we should say that someone knows something (or, alternatively, when it is the case that someone knows something).

The focus of most epistemological research is on generating rationally coherent principles regardless of their practical application. In this thesis I want to argue that there is space within epistemology for practical accounts of knowledge that can decide in the here and now whether or not someone knows something. I will defend a communitarian/collectivist approach as standing the best chance of providing such practical application. I will also demonstrate that such an approach holds its own conceptually.

I defend my account by focussing on the risky world of petroleum engineering where being in possession of knowledge can be a matter of life and death. The search for knowledge that takes place in this epistemic community impressed upon me the importance of an epistemology that is grounded and useful. This is not to say that epistemology that lacks this immediate practical application should be abandoned but that it should not limit the boundaries of our subject.

With this in mind, I spent approximately 10 weeks training as a logging analyst in South-East Asia with an American multinational petroleum engineering company which provides services to the largest as well as some of the smaller oil, gas, and energy corporations in the region. The purpose of this venture was chiefly to learn myself the practice of acquiring knowledge in a particular setting and be less of a fly-on-the-wall and more of an inquirer on-the-ground or at-the-coalface. I also spent this time observing my colleagues as well as talking with engineers about their attitudes to their profession, training, the kinds of inquiry they do on a day-to-day basis, and the knowledge they acquire. I observed work done on oil rigs around the region (visiting various sites around Thailand, Malaysia and Brunei), how problems were approached, how discussions were carried out, and how consensus was reached about how to solve a problem.

I came to recognize two methods of knowledge-gathering: a trial and error approach carried out on the ground, involving back-of-the-envelope calculations, rough-and-ready tests, and finding the best available solution within a given time-frame. And a more deliberate consensus-based approach carried out back in the office surrounded by computer software, mathematical textbooks and reference tables, and other analysts and field engineers.

Of course, there is a near-endless variety of methods in-between these two caricatures but these observations highlight an important tension between the ways we go about acquiring knowledge. Some questions are pressing, in the sense that William James described.[1] (James 1890, ch. 24) They require quick, intelligent right action. Oftentimes, they involve testing a solution, tweaking and manipulating it, testing again, adjusting it again, and so on. Other questions are not immediately pressing although they will serve a purpose at some point in the future. They require careful consideration and, almost certainly, interaction with ones epistemic community. In engineering, they may involve following a specific procedure that has worked well in the past or is believed to work well. I believe that epistemologists can learn a lot

---

[1] James' example has us imagine we are being chased by gunmen to the edge of a ravine and must decide whether to face the men and their guns or jump. Fortunately, most engineering decisions do not involve guns or jumping off ravines but we can still consider them to be pressing: in need of an answer within a specified time-frame, using limited specified resources, and within a certain context.

about knowledge by looking at how we do, in fact, acquire knowledge and in Chapter 5 of this dissertation I expand on some of the lessons I brought home with me.

Ultimately, I aim to demonstrate that concepts such as knowledge are not found but made in the same way that we do not find technologies but make them. In other words, the laws of knowledge are less like the laws of physics and more like the laws of culture, linguistics, economics, and the like. I argue that what counts as knowledge, and who (or what) counts as a knower, can and has changed as a result of social and technological change, sometimes called sociotechnical evolution. (Bonen 1979; Trist 1981) Like technology, as we generate knowledge we are also defining what knowledge is. Since petroleum engineering has played such a crucial role in shaping our environment and defining our social world, it is perhaps unsurprising that its history and practices can also teach us something about our conceptual world.

1.1.1 The research problem: What does BP know and how does it know it?

On April 20th 2010 at approximately 10 P.M. CDT a large explosion and fire was observed on Deepwater Horizon, the Transocean drilling rig licensed to BP in the US Gulf, 50 miles from the Louisiana coastline.[2] 11 crewmen died. Over the course of almost three months, a joint government and industry task force had failed to stop crude oil flowing from a well causing widespread damage to the sea and coastline. The accident was also damaging for the energy companies involved – principally BP, Transocean, and Halliburton – who received increasing legal and political pressure to answer questions about why the accident occurred, and what they knew about the rig and the likelihood of an accident. The Deepwater Horizon oil spill raised political, financial, and ethical questions about what the organizations involved knew, what they could reasonably have been expected to

---

[2] New York Times, April 22, 2010, p. 13. Robertson, C. 'Search continues after oil rig blast.'

know, and what they had a duty to know. Some headlines from the time of the accident help illustrate what was being asked,

> What did BP know?[3]
> Did BP know more than they led on?[4]
> BP, Halliburton Knew Oil Spill Cement Unstable[5]
> BP and Halliburton knew of Gulf oil well cement flaws[6]

Specifically, politicians, journalists and members of the public were keen to discover if BP and other corporate bodies knew that a blowout was likely and preventable. On the 4th May 2010, Edward J. Markey, then chairman of the United States House Select Committee on Energy Independence and Global Warming remarked,

---

[3] Reported in *The Columbus Dispatch*, 9th June 2010. Lugarten, A. & Knutson, R. What did BP know? Retrieved from http://www.dispatch.com/content/stories/national_world/2010/06/09/what-did-bp-know.html.

[4] Reported in *NBC News*, 30th October 2010. Wilson, J. Did BP know more than they led on? Retrieved from http://www.ketknbc.com/news/did-bp-know-more-than-they-led-on.

[5] Reported in *Discovery News*, 29th October 2010. BP, Halliburton knew oil spill cement unstable. Retrieved from http://news.discovery.com/earth/bp-halliburton-oil-spill.html.

[6] Reported in *The Guardian*, 29th October 2010. Goldenberg, S. & Kollewe, J. BP and Halliburton knew of Gulf oil well cement flaws.

> BP has been in control of all the information and ultimately when the investigation is complete we will have to ask what did BP know and when did they know it.[7]

Since these questions often centre on what is, was or could have been known, epistemologists ought to have answers to who or what it is that holds the knowledge, what it means to say in the here-and-now that they have knowledge, how it is (or should be) transmitted, who or what kind of entity can be responsible for it, and so on. These 'real-life' events give rise to philosophical problems and, vice versa, solutions to philosophical problems can cast light on the 'real-life' event. Epistemologists ought to be able to say, regardless of their technical expertise on petroleum engineering or oil rig maintenance, whether a corporation can know something, or, if not, to explain what Markey means by, 'did BP know…?' As we shall see, the problem is not easily dismissed as metaphorical shorthand or a casual way of talking. It is, at least on the face of it, incongruous with the classical account of knowledge as justified true belief (JTB) since it is not at all obvious that

---

[7] Reported in *The Guardian,* 6th May 2010, p. 25. Goldenberg, S. 'Deepwater Horizon oil spill: Obama attempts to limit political fallout.'

corporations can possess mental states such as belief.[8] We can get at the answer to this question, I propose, by examining when and why we attribute knowledge to individuals and the role of technology and social interaction in generating knowledge.

Assuming there is some way in which Markey's question makes sense – whether because the group, BP, can know or whether some individuals belonging to that group can know – epistemologists ought also to be able to classify the kind of knowledge involved and specify conditions under which it would be possible for agents to know. It will be an empirical matter whether they do in fact meet those conditions but there is an entirely separate question of what must the agents be in possession of, or in what standing must they be, in order for them to know. Must they truly believe that, let us say, the well was likely to blowout? Must they be justified in this belief and in what way? Must

---

[8] I will explain the classical account of knowledge in the following discussion. For now, it will suffice to say that JTB was a theory of knowledge that held considerable weight up until 1963 when Edmund Gettier published one of the best-known papers in epistemology. (Gettier 1963) This paper has been taken to show that there is more to knowledge than justified true belief and much of the subsequent near-half century of analytic epistemology has been devoted to responding to this task.

they have some degree of confidence in their belief or even be certain in that belief? Ought they to have believed certain other things given the information available to them? What standing must they be in in order to have an epistemic responsibility to predict future events? There are countless questions one might ask from an epistemological point of view that have factual answers once we settle precisely what form an answer should take. Lastly, epistemologists ought to be able to classify the kind of knowledge that relevant agents may have had (again, the practical question of whether they did have it being a question of fact for engineers or lawyers). In short, epistemologists ought to be able to say what relevant agents could have known, under what conditions they would have known it, and when they could reasonably have been held responsible for knowing it.

The central research problem of this thesis will not, however, primarily be of interest to lawyers or BP engineers but to epistemologists interested in the view that technical knowledge, of the kind possessed by BP engineers, is a distinct genus of knowledge worthy of their attention. That it is not a matter of applied scientific knowledge or simply a branch of 'know-how'. Crucially, I want to argue that technology is not just something that we know things about – 'I know that a hammer is for hammering' – but that it changes what we know and how we know it. Technology – and technical artefacts – changes what knowledge is.

Now, it is obvious that some sociotechnical change can affect what we know and how we know it in a trivial way. No one could have known, for example, how to operate a microscope or the physical effects of a gyroscope before their invention and production (though in cases such as these we might sometimes be able to say that they would have the disposition to know it). This kind of epistemic change is uninteresting in the same way that I now have knowledge about a species of sea snail recently discovered that I could not have had before. This is uninteresting for epistemology because the new knowledge is a new *token* of the same *type*; a new page in the same book. There is nothing essentially different about my knowledge of sea snails than my knowledge of any other animal. Sociotechnical change, however, fundamentally changes what we know and how we know it in a way that, say, would merit a special chapter in the epistemologist's textbook called, 'Technical knowledge' alongside chapters on 'Perceptual knowledge' and 'Testimonial knowledge'.

In Chapter 2, I will look at what answers are available in the existing literature by means of a survey of social epistemology and the philosophy of engineering.[9] I pick out communitarianism, contextualism, and the Strong Programme in the sociology of knowledge as being particularly helpful in analysing technical knowledge. This is not to say that any epistemology of engineering and technology which employs more traditional analytic epistemology is wrong-headed or irrelevant. I pick these approaches out as they are useful in empirical studies in a way in which the alternatives are not and, although I do not argue that all epistemology ought to have some empirical utility, the heuristic value provided by these approaches seems germane to a study in

---

[9] Technology, I take it, is uncontroversially social. No societies, no hammers. But, for those who doubt this entailment, I aim to establish it in Ch. 4 of this thesis, which deals with the ontology of technical artefacts. Establishing that since technology is social, technical knowledge is social will be a harder slog but the reader should be more convinced than she is now by the ends of Ch. 5 and 6.

applied epistemology, particularly so with philosophy of engineering. As I said in my opening remarks, a traditional JTB-style answer to the question of knowledge is not able to give much practical guidance when attributing knowledge. In other words, it does not much help us to guide our behaviour as we rarely have access to the validity of many of the relevant criteria. Is the agent's belief true? We do not know; we are not infallible. Is the agent's belief justified? Well, it depends on what counts as a justification and any general principle of justification is unlikely to help us. In brief, such accounts are heavily dependent on stipulation. Given that the agent is justified in believing a true proposition, it is true to say that they know it. Given that the belief formed by the agent corresponds to the external reality in the right way, it is true to say that they know it. Variations on this theme hope to solve the philosophical questions about knowledge but do not have an apparent solution to the practical question: when can I attribute knowledge to this person?

This dissertation delivers a number of central novel theses which aim to provide a practical solution to knowledge attributions: I argue against accounts which point to the convergence of scientific and technological disciplines and in so doing ignore the heuristic value to conceiving of technical artefacts and technical knowledge as separate from scientific knowledge (notwithstanding their relation just as memory as a source of knowledge is closely related to perception but neither is reducible to the other or could be said to progress towards convergence). I argue that semantic epistemologies such as semantic contextualism and communitarianism ought to follow the implications of their thesis that 'knows' (and its cognate terms) should be analysed as semantic expressions. I argue that technical artefacts are artificial kinds and that knowledge thereof is consequently knowledge about artificial kinds.

As argued by Kusch and the Dual Nature Thesis (see §4.1), artificial kinds are kinds which combine aspects of the type of reference-talk relevant to both natural and social kinds. I analyse reference-talk about another artificial kind – tools – in order to develop an ontology of tools. In the later chapters I move on to consider what kinds of entity can be attributed with possessing technical knowledge. I consider, as a starting point, what tool use is, concluding that, in certain cases, animals as well as some technical artefacts (e.g. computers) can use tools. They can also design and produce tools. Consequently, these are not barriers to including them as bearers of technical knowledge. However, I note that the aforementioned do not, under some accounts, possess beliefs and mental states comparable to those of human knowers. This may prohibit them from truly possessing technical knowledge. Through an argument that focuses on perceptual knowledge, I argue that the bearer of knowledge can extend beyond the boundaries of human bodies.

1.1.2 The research questions

I will define the research questions for this thesis as follows:

1. What is social epistemology and how can it help contextualize the particular epistemic responsibility that a professional person holds. In particular, how can a communitarian epistemology – which treats knowledge as a kind of social status—define different kinds of knowledge?

2. What kind of knowledge do petroleum engineers have? What is technical knowledge?

3. What are technical artefacts and how does this affect technical knowledge?

4. Some technical knowledge appears to be perceptual. How do perceptual sources generate technical knowledge?

5. What epistemic agents can have this knowledge (individuals, groups, animals, artefacts, computers, etc.)?

1.1.3 The research answers

Answers to these questions cannot be generated by abstract general principles alone but we must refer to the real-life cases. That is, just as others have argued that philosophy and science are continuous, we should say that philosophy and engineering are continuous or, in this case, that epistemology and practices of inquiry and knowledge attribution in engineering are continuous. (see, e.g. Churchland 1986; Quine 1976, p. 233) Chs. 5 and 6 contain the empirical fieldwork research that aims to follow a continuous approach to philosophy and engineering. In answer to the questions outlined in §1.1.2, I will argue for the following,

1. Question 1 is addressed in Ch. 2 and 3. Social epistemology is seen to be a broad church with many non-complementary factions. I pick out communitarianism, contextualism, and the Strong Programme, as being of particular value to my project for the reasons given above. Communitarianism holds that

knowledge is a kind of social status and I argue in favour of extending this principle to distinct kinds of knowledge. One can have the status of a technical knower, meaning that one is attributed with the status of possessing technical knowledge. Typically, this status is held by engineers although there is of course a bleed-through from scientists and others who hold the status of technical knower. Secondly, communitarianism holds that epistemologists should explain individual knowledge with reference to the relevant community. Consequently, differences in social factors such as training, nationality, and profession, take precedence over individual factors such as reliability and epistemic virtue.

2. Although petroleum engineers, as human beings, will be in possession of many different kinds of knowledge – propositional, perceptual, scientific (roughly, knowledge of natural kinds, events, and processes), testimonial, and so on – this thesis will focus on their technical knowledge. As previously de-

fined, this is knowledge of how to design, produce or operate a technical artefact.[10]

3. Since I hold that technical artefacts can affect our human concepts (and thus what counts as knowledge) it is important to establish what technology is. I argue that technical artefacts are artificial kinds. Kinds are distinguishable by the reference-talk associated with them and so I base my ontology of technical artefacts on a theoretical and empirical study of technical artefact reference-talk. Ch. 4 concerns the theoretical basis of this and Ch. 5 contains some empirical studies.

4. Perceptual knowledge is particularly interesting as it appears to offer us a way to include animal and non-human technical knowers. In Ch. 6, I argue on the basis of extended cognition theory and neuroscientific research, that the bearers of perceptual knowledge can extend beyond the human body into technical artefacts.

---

[10] The ‘or’ in this sentence is typically an exclusive or (XOR). Although in some cases it is possible that an agent knows how to design, produce, and operate a technical artifact I do not mean to suggest that it is necessary.

5. Although I want a theory of technical knowledge to include certain non-human agents such as computers and other informational technical artefacts, this only allows 'extended epistemic agents' into the group of technical knowers. There is still, somewhere in the system, a human agent that many will argue is the true bearer of the knowledge. In Chs. 7 and 8 I explore an informational approach to knowledge which allows various other entities to enter the fray including artefacts and groups. Crucial to the informational approach is that the agent need not possess or be capable of possessing beliefs about what is known.

1.2 Chapter outlines

Ch. 2 is a literature review, albeit with an additional aim of setting up various concepts and problems for the remainder of the thesis. The first task will be to assess what answers are available to the research questions in current state-of-the-art social epistemology (SE) and the philosophy of engineering and technology (PET). This chapter offers a broad and inclusive survey, minimizing any preconceptions about what either discipline is or where to look for answers. I look at surveys and definitions of SE proposed by Lorraine Code, Alvin Goldman, Philip Kitcher, Martin Kusch and Frederick Schmitt and attempt to provide a more objective overview of the discipline. Existing surveys tend to be quite quick and simple and, however neat and elegant these simple surveys are, I intend my contribution to be more extensive. I am extremely inclusive with my review, attempting to cover anything that has gone under or been attributed with the name 'Social epistemology.' The advantage of this approach is to put various branches in stark contrast with one another so that the reader can see where their own commitments lie and who opposes them.

In the third section I describe some of the open problems in the philosophy of engineering and technology (PET), a discipline that stretches back as far as the early 1960s but has not yet made the necessary strides to establish itself as a recognizable discipline in philosophy. These problems include the nature of engineering knowledge (the professional corollary to what I am calling technical knowledge), the relationship between science and engineering (and scientific knowledge and engineering knowledge), the social aspect to engineering (the idea that engineering is slave to social factors in a way that science is not), the engineering method, among others. I conclude by setting trends in the philosophy of engineering within the framework of my social epistemology survey.

Ch. 3 investigates and evaluates the epistemological views of Martin Kusch and Keith DeRose. I see these as the best-equipped epistemological theories for my project as they allow the flexibility necessary to characterize technical knowledge as a species of professional knowledge, i.e. a various, context-rich social status.[11] This

---

[11] I use the phrase 'context-rich' here advisedly. This is a pedagogical term used to describe problems given to students set in a practical environment, either in its description or actually. An example would be to give an engineering student a triangulation problem and set that problem not in the abstract world of geometry but in the practical world of directional drilling. That is, to describe how triangulation can be used to determine the shortest path towards a target. I think this neatly captures a difference I wish to make between the (albeit worthwhile) activity of philosophical problem-solving set in the world of formal or informal logical inference and argument and the activity of philosophical problem-solving set in practical situations.

is necessary to make the link between the reference-talk of a community and the kind of knowledge it generates. Although Kusch's view has developed since its publication, I will focus on his 2002 book *Knowledge by agreement: the programme of communitarian epistemology*. (Kusch 2002) This book contains the groundwork for a communitarian approach to epistemology and its central tenets are necessary components of my own analysis of technical knowledge viz. that 'knowledge' is a word which we use to refer to agents who meet certain social criteria, and that these criteria are only discoverable through analysing the social conditions at the time.

A central virtue of Kusch's account is that it foregoes the tendency to ignore historical context and changes to what counts as knowledge in favour of a sociologically-informed theory of knowledge. It acknowledges further that our concept of knowledge changes over time.[12] Whilst this latter point is not essential to my own thesis it is, I

---

[12] See, by comparison, the recent work in the 'genealogy of knowledge.' For example, Craig 1990, 2007; Fricker 2008; Kusch 2011; Williams 2002. Craig suggests how our concept of

think, an inevitability once we decide that the concept is socially-dependent. Thus, any non-social account of knowledge fails to explain any change that is observed. My thesis demonstrates that one element affecting this change is technological change or evolution.

In the second half of this chapter I explore what contextualism has to offer a study of technical knowledge. I look at its history and concerns, focussing on DeRose's semantic contextualism which dovetails neatly with certain aspects of Kusch's book. In order to substantiate the claim that technology affects our concept of knowledge in a genealogical manner, we need to know what technology (and technical artefacts) is. If they were natural kinds, we would need to know what makes them different from other natural kinds. We would need to know why the new knowledge that technical artefacts (as natural kinds) enables is any

---

knowledge might have evolved from a concept of proto-knowledge (which may be thought of as a concept which can be used to flag reliable informants, for example). Note that Craig's suggestion, if it is taken as a claim as to how our concept of knowledge did originate, does not in itself entail that our concept of knowledge continues to change. To make this claim we would need to establish that there is continuous feedback between our behavior as epistemic agents and historical change.

different from the kind of knowledge that is already enabled by natural kinds: e.g. knowledge of physical processes, entities, and events. It is true to say that technical artefacts do resemble natural kinds in some respects: they exist in spatiotemporal physical space; they have a certain physical structure and can cause physical effects.

However, in other respects, they more closely resemble social kinds like money or religious ceremony. There is a clear sense in which a gas pipeline is not a gas pipeline merely because of its existence in spatiotemporal physical space but because it serves a human purpose, that of allowing safe transport of said hydrocarbon between two locations of interest. Without this purpose it would seem very odd to call it a gas pipeline at all.[13] In Ch. 4 I describe two attempts to take the ontology of technical artefacts seriously. One is the Dual Nature of

---

[13] One can imagine an object with the same physical structure of the Forth rail bridge in Edinburgh but floating purposelessly in Outer space. Here we would be more tempted, I think, to call it a coincidence than a bridge. By contrast, there is a sense in which a natural object, such as a rock, floating purposelessly in Outer space is still very much a rock, it just so happens to be extra-terrestrial.

Technical Artefacts account which describes technical artefacts in terms of combinations of natural and social descriptions and the other is influenced by the work of Kusch and Pablo Schyfter which I call the Collectivist Account of Technical Artefacts. This proposal describes technical artefacts in terms of their reference-talk.

Ch. 5 introduces the fieldwork case studies. I give a brief history of the engineering knowledge involved and provide some observations on problem-solving and the engineering method which endorse the hypotheses stated in §2.3. In the second section I provide an empirical analysis of reference-talk in these disciplines and develop a putative definition of what a tool is. I analyse the references made to a certain class of technical artefact – tools – in various kinds of documents pertaining to the engineering disciplines studied. I classify each of these references according to distinctions made in discussion with a group of engineers. In the third section I discuss some implications of this definition, problems and paradoxes. I look at each part of the definition of tool given previously and discuss some of their finer aspects in the context of the engineering disciplines. Technical knowledge is found to be more than simply knowledge of how to design, produce, and operate something. It is knowledge of how to design, produce, and operate something that provides information to the user of variable reliability and which has a proper (social) function.

One of the most fascinating discoveries of my fieldwork research was that the more knowledge someone possessed in a particular discipline, the more likely they were to describe the phenomenology of knowing in perceptual terms. With novices in production logging analysis, for example, I asked, 'How do you know *x*,' and the typical response was procedural in nature: 'this is how we calculate *x*,' or '*x* seems to correlate with other values I have,' or 'this is how I usually do it, if such and such is the case.'[14] However, in the case of more experienced analysts, their expertise and familiarity with the data led them to almost perceptually 'feel' their way to the answer. When one asks these

---

[14] As part of my fieldwork research I undertook training as a production logging analyst. I acquired a basic knowledge of the principles of analysis and the use of two computer software programmes: PLWin and Emeraude.

analysts, 'How do you know *x*,' they were more likely to answer, 'It just looks like *x*,' or 'I can see that *x*,' or, by instruction, 'You can see here that *x*.' Although this may be simply a metaphorical way of speaking that one acquires after so many years in the job, it seemed to me more significant. In Ch. 6 I argue that work in extended cognition may provide an explanation for this phenomenon. Through a discussion of research in cognitive science and neuroscience I come to the conclusion that it is quite possible not only that analysts are acquiring knowledge through perceptual activity but that the software used in the field attempts to emulate this activity. We thus have a viable analogy between human and non-human epistemic agents.

The aim of Chs. 1-4 is principally to emphasize the need to analyse knowledge as a social phenomenon. Chs. 5 and 6 give an example of what a social account of knowledge would look like in practice. However, I do not believe that a purely sociological approach would be enough. It is too capricious, too unstable. Surely, it enjoins us to ask, knowledge is not purely subject to what a group of engineers happen to agree upon at any one time. This would be a strong concession towards relativism and would force us to say that a systematically misled society could happily acquire knowledge on its own that could bear little relation to what was going on, as it were, in reality. Whilst I think this objection can be overplayed, Ch. 7 uses work in the philosophy of information to give a more nuanced account of how concepts become meaningful through a constructive process within a society and how the kind of work done by analysts in the fieldwork study can be compared to the information-processing of computer software that constantly threatens to make them redundant. In brief, I argue that neither a purely social account nor a purely informational account could give a comprehensive account of an epistemic culture. Instead, we need a combination of the two.

Finally, Ch. 8 returns to the original research questions and answers set out in §1.1. I discuss what it would mean for a group, such as BP, to know a proposition, to the extent that it is meaningful, and develop the idea of an epistemic community. I return to each of the problems described at the start of the thesis and draw together the selection of answers I have given in each chapter.

Chapter 2

Social epistemology and the philosophy of engineering and technology

2.1 A survey of surveys

2.1.1 The existing surveys

This chapter contains two literature reviews: a survey of work in social epistemology and a review of the philosophy of engineering and technology. In the first section, I discuss the aims and methods of social epistemology, broadly construed. Whilst many existing surveys are limited to particular approaches to studying social knowledge my aim here is to be inclusive and general. I suggest that the simplicity of existing surveys sacrifices detail for elegance. Instead, I suggest we can identify three 'sticking points' about which various accounts diverge and plot these on a Cartesian coordinate system. This provides a useful way of distinguishing between different approaches to social knowledge and the various ‘sticking points’ at which different approaches diverge. The third section reviews the history and current state of the philosophy of engineering and technology (as well as some related disciplines). I discuss the relationship between science and engineering and how attitudes to this relationship have changed over the last half decade. Finally, I set out an argument for studying engineering and technical knowledge in its own right rather than either a) as an epistemically uninteresting form of applied science, or b) as a hybrid technoscientific discipline mixing knowledge and methods from both science and engineering. (See Gille1986)

The term ‘social epistemology’ has been reserved by and for a broad range of actors over the last 40 years from library scientists to sociologists to analytic epistemologists to ethnomethodologists to Karl Marx and many more besides. (Goldman 2006) Here are some recent descriptions from the literature:

> Social epistemology is the study of the social dimensions of knowledge or information. (Goldman 2006)

> Social epistemology distinguishes itself from sociology of knowledge in its goal of providing a normative analysis of knowledge (Grasswick 2008, paraphrasing Fuller 1987)

> Social epistemology begins at the point of rejecting the individualistic reduction. (Kitcher 1994, p. 112)

> Social epistemology is the philosophical study of the relevance of communities to knowledge. (Kusch 1998)

> Social epistemology is the conceptual and normative study of the relevance of social relations, roles, interests, and institutions to knowledge. (Schmitt 1994b, p. 1)

> Testimony's relegation to last place on this list [of sources of knowledge] signals its lesser ranking in relation to the other sources - perception, memory, and reason (or induction) - because reliance on testimony has long been thought to compromise epistemic self-reliance. Social epistemology, I suggest, reverses this ranking.[15] (Code 2010, p. 29)

It is usually a little myopic to survey a fledgling discipline based on a brief history of development. In their own time, for example, none would have thought to group Descartes, Leibniz and Spinoza together

---

[15] By way of explanation, Code is here referring to what she considers much early 21st-century Anglo-American epistemology to be 'residually loyal' to. Namely, a tradition in which perception, memory, reason, and testimony are the principle, most reliable sources of knowledge, roughly in order of priority and reliability.

as one side of a debate with Berkeley, Hume and Locke. It was only posthumously that these names were organized as belonging to either a Rationalist or Empiricist movement. The quotes above demonstrate how a new field is often deliberately positioned against an existing field whether that is individualist epistemology, as in Goldman's case, analytic epistemology in general, as in Kusch's case, or the Strong Programme in the sociology of knowledge, as in Fuller's case. Nevertheless, there have, in recent years, been several general surveys of social epistemology from prominent scholars.

The first section reviews surveys from Goldman, Kusch, Schmitt, and Kitcher. I will show that each of these efforts fail to comprehensively and objectively characterize an emergent discipline either by being unjustifiably exclusive or overly simplistic in their metric. Each of these three bases their survey on what may be called a spectral metaphor. That is, all have in mind a spectrum of views with each view representing a fixed point along a single axis. Typically, each 'end' of the spectrum represents an extreme view (sometimes held by none).

The extremes are often dismissed in favour of a middle path between the two.

2.1.2 Goldman's spectrum of social epistemology

Alvin Goldman has been perhaps the most prominent advocate of SE in analytic philosophy. He also has a specific conception of what SE is or should be. In Goldman 2010a, he responds to a criticism from William Alston that 'much of the material in [his] book would be rejected by many contemporary epistemologists as "not real epistemology."' (Alston 2005, p. 5; and Goldman 1999) Of course, Goldman denies this and argues that his work is indeed 'real'.[16] What is real epistemology for Goldman? As it turns out, he does not much

---

[16] Goldman's book (Goldman 1999) is controversial in some circles for wishing to include non-traditional subjects of epistemological concern such as social organizations (e.g. research laboratories, law courts, democratic institutions, schools and universities, and websites such as Wikipedia). I am greatly indebted to this work in forming my thinking about 'knowledge in a social world' but suspect the debate over whether or not it counts as real epistemology is of more significance to shaping the future of a discipline than of collecting under one term attempts to study similar phenomena. This book is probably still the primary source in social epistemology despite being over a decade old.

disagree with Alston; he just thinks that his SE meets Alston's conditions. Real epistemology for Alston and Goldman is simple epistemology as it has been typically practiced in twentieth-century Anglo-American philosophy departments up until now. Goldman identifies six central tenets of this practice, some of which are overtly responses to work elsewhere that challenges traditional epistemology.

> TE1. Individualism: the epistemic agents are exclusively individuals. All stories, situations, thought experiments, intuition pumps, and so on focus on lone individuals without any need to describe their social context, role, status, history, etc.
>
> TE2. Normativity: it focuses on a narrow remit of evaluating distinctively epistemic states such as justifiedness, rationality and knowledge.
>
> TE3. Objective standards: these evaluative or normative standards are assumed to be objective standards. Justifiedness and rationality are 'not merely conventional or relativistic, but have some sort of objective validity.' (Goldman 2010a, p. 2)
>
> TE4. Truth: if a proposition is known then it is true.
>
> TE5. Realism: truth is an 'objective, largely mind-independent, affair.' (Goldman 2010a, p. 2)
>
> TE6. The 'central business' of traditional epistemology is the 'critical examination of doxastic "decision-making".' (Goldman 2010a, p. 2)

Goldman's proposed spectrum relates to these conditions. He argues that there are currently three groups of views in competition with each other over the same theoretical ground: preservationism, revisionism and expansionism. Preservationism broadly sticks to the tenets above; revisionism rejects them all; and expansionism (Goldman's preferred position) wants to go beyond the traditional tenets to examine social aspects to knowledge whilst preserving its general methodology. These positions can be located on a spectrum such as that represented in Fig. 1.

Figure 1. Goldman's spectrum.

Expansionism and preservationism belong in the category of 'real epistemology' and revisionism is a kind of non-epistemology or pseudo-epistemology.

2.1.3 Real vs. un-real epistemology

Goldman's spectrum is exclusive and (fairly) conservative. He seeks to exclude some areas of research which are, erroneously, bracketed under SE. For Goldman, these are 'deformed or bastard versions of epistemology' including postmodernism, deconstructionism, social constructionism, and various studies of science including the Strong Programme in sociology of science (SP).[17] Goldman states that these disciplines are incompatible with many, if not all, of the tenets above.[18]

---

[17] This phraseology is taken from an early incarnation of the 2010a paper retrieved from http://www.philosophy.stir.ac.uk/postgraduate/documents/GoldmanSEPaper.pdf.

[18] I disagree with Goldman's attempts to bracket these theories together as it is only their opposition to a form of analytic epistemology that seems to mark them out. I would argue instead that disciplines should be collected by the phenomena they study. It is true that we do not collect homeopathy or phlogiston theories in serious medical and chemical textbooks but to say that these revisionist theories are not sufficiently serious seems unproven. At the very least, Goldman ought to say more about why they do not deserve to be considered as possibly useful methods. All are explicit attempts to study the social character of knowledge. Whilst social constructionism and SP may be said to be incompatible with the tenets mentioned above, it is not obvious that others, e.g. Richard Rorty, are. Rorty would, at the very least, concur with TE4 and TE5 and his opinion on the others is the subject of my forthcoming paper with J. Adam Carter (Carter & Kerr, in preparation). (Rorty 1989, 1995, 1998) The sticking point is not said to be with a different conception of social knowledge but with divergence from current 'traditional' methods of epistemology which, after all, may or may not be correct and a case still needs to be made to show that they are the best methods for studying social knowledge. What more accurately, perhaps, characterizes the revisionist for Goldman is a rejection of what he calls the distinction between truth and institutionalized belief. (Goldman 1999, p. 7) For Goldman, veritism (giving appropriate place to truth in the sense of some kind of correspondence with reality) and epistemology are inseparable. Goldman has since tempered this position, noting that often aiming in non-truth-conducive ways can be epistemically better. For example, information about previous guilty pleas by the alleged in a court trial is non-truth-conducive but can be

Most of the practitioners of these subjects do not describe themselves as social epistemologists (SEs) but are sometimes called such or discussed as presenting challenges to SE. The revisionist camp extends to those who engage directly with traditional epistemology on its own terms but who reject some of the tenets. Martin Kusch for example, does not consider himself a social epistemologist precisely because of disagreements with SE. (see §3.1.1) Lynn Hankinson Nelson is a social epistemologist who denies that individuals are epistemic agents and so, presumably, does not practice real epistemology for Goldman. (Hankinson Nelson 1993)

Real epistemology, on the other hand, is characterized by the preservationists and the expansionists. Preservationists keep the traditional tenets but would like to add some questions for epistemolo-

---

epistemically better since it may prevent undue bias given to this fact. Similarly, the identity of Deep Throat was kept secret for some time in order to pursue the higher epistemic end of uncovering the Watergate scandal. (Goldman, in a presentation to the Episteme 2012 conference on privacy and secrecy, Delft, Netherlands) One problem, I submit, that remains with this approach is that it is not obvious why someone who does not hold the same intuitions about knowledge and truth should accept this characterization of proper research in social epistemology. As I show in later chapters the engineers as a group do not seem to factor truth into their decisions (largely because it is usually not possible to establish what is true in non-trivial cases) and there is a further case to be made that South-East Asians represent the kind of community that does not recognize a strong correlation between truth and knowledge.

gy that have a social dimension. Most commonly these include: What is the nature of testimony-based justification? What does rationality require agents to do, in terms of belief revision, if they find themselves in a situation of peer disagreement? And so on. Goldman's expansionist project is more ambitious. It seeks to identify activities which require some kind of communal doxastic-decision-making; social activities which are conducive to epistemic ends (e.g. true belief, justified belief, knowledge, etc.) Examples might include legal trials (especially regarding issues of evidence), political inquests and democratic structures, forensic medicine, journalism, electoral voting or Wikipedia.

In other words, they include any organization with an explicit aim to acquire knowledge collectively. It is unashamedly evaluative, and seeks to assess and, if possible, suggest improvements in the policies and procedures of epistemic or quasi-epistemic social systems. Epistemological inquiry into these organizations still qualifies as real epistemology so long as it coheres with most of the traditional tenets. My work in this thesis continues Goldman's efforts to apply epistemological principles to real-life problems and social institutions although I do draw on the work of revisionists such as Kusch and the Strong Programme in the sociology of scientific knowledge (SP) to support this problem-solving exercise.[19] I am also skeptical of Goldman's commitment to veritism and his own skepticism towards 'veriphobic' epistemology.[20] Whilst no one would surely doubt that we ought often to promote knowledge practices that are likely to lead to truth I will here offer two reasons why we may not want to put this goal at the centre of our epistemology.

The first is that there are cases where we are genuinely not interested primarily in the truth and that we would rather have a practice that was not likely to lead to the truth but did suit some other end. For example, I can inquire into whether a colleague is wrong on a certain matter of

---

[19] SP is a research area in the sociology of scientific knowledge developed by a group of sociologists at the University of Edinburgh in the 1960s including Barry Barnes, David Bloor, John Henry, and Donald MacKenzie. I discuss this work in detail in §3.1 in particular although these authors appear throughout the thesis.

[20] Veritism is the view that we should conduct our epistemic activities in a way that, generally, produces, generates or promotes true belief. Thus any theory that does not support this unified goal could be called 'veriphobic'.

fact not because I am interested in what the truth of the matter is but because I want to show my colleague to be wrong.[21] Whilst Goldman would probably concede this point, he can respond by saying that whilst there are such cases these are 'outliers' and, in general, we are typically interested in finding out the truth. Certainly, truth is not always the aim of inquiry. But it may also be demonstrable that it is not the constitutive aim of inquiry. Hongladarom (2002a) demonstrates that in some cultures truth is not the primary goal of inquiry and that some cultures have different intuitions on whether knowledge requires that the proposition be true. It may simply be a consequence of our own cultural determinations that we intuitively accept that knowledge is truth-entailing. Certainly it is not currently proven that there are cultures which are content for their inquiries to systematically lead to falsehood but we ought not to be complacent in thinking that the opposite is therefore proven.

---

[21] I would not like the reader to think that I ever engage in such vindictive activities but the example sprung to mind and seems apposite.

Further, my own fieldwork research indicates that engineers are often disinterested in whether or not something is true. This is not to say that they think it is probably false or that they would not care if they were told that it was true (although they may ask on whose authority you say that it is). Indeed, were I to provide evidence that what they believe is not true they would surely consider it relevant. Rather, it is to say that the concept of truth does no work in deciding how to act, in reflecting on how one did act, or in monitoring or adjusting one's epistemic practices. I simply cannot aim at the truth out of a desire to do so. I can at best pursue reliable, successful practices and hope that these promote truth but we have an entire history of scientific inquiry that shows us that we can spend a lot of time aiming at what is true only to find that very little of what we discovered was. This is the attitude of the engineers I spoke to – and to some extent the interviews with non-Western people – that I wish to emphasize here.

First, if we want to claim that truth and knowledge are linked we have to show that this claim is based on more than our personal intuitions. We must show that it is based on more than an appeal to the seemingly counter-intuitive fact that it sounds odd to say, 'She knows it but it isn't true.' As our scientific history shows us, our intuitions and common sense frequently turn out to be dramatically mistaken. I do not aim to bolster this claim in this thesis. In fact, I wish to put the issue to one side for the purpose of pursuing another goal viz. to develop a theory of knowledge that is useful to analysing and evaluating epistemic practices. I think only such a theory would be able to contribute to Goldman's efforts to extend the domain of epistemology into evaluating, improving and contributing to our epistemic practices.

2.1.4 Kusch's spectrum of social (and communitarian) epistemology

In *Knowledge by Agreement*, Kusch embarks upon a similar programme to Goldman (Goldman 1999, 2008, 2009, 2010a, 2010b) albeit towards different ends. (Kusch 2002, 2004) Kusch is more inclusive in what counts as real SE but wishes to distinguish his own position from many of the traditional projects in SE. For Kusch, SE is a spectrum of positions which give varying degrees of privilege to social and political explanations in the study of knowledge. He identifies two distinct movements in SE: one allied with the sociologist Steve Fuller which he calls the science policy programme (see Fuller 1987); and one allied with Goldman which he calls the complementary programme.[22] The science policy programme is an explicitly political attempt to influence policy-making decisions in scientific institutions in order to make them more democratic and accountable to the public. It considers knowledge to be a social phenomenon and approaches this phenomenon sociologically. The complementary programme is also evaluative of policy-making decisions but it distinguishes between individual and social aspects to knowledge. Traditional epistemology, it says, is going in the right direction but has largely neglected these social aspects.[23] Kusch's project – communitarian epistemology – distinguishes itself by being descriptive, rather than prescriptive, towards the epistemic make-up of scientific and other knowledge-

---

[22] Somewhat surprisingly, SE has an older history of systematic use in sociology than in philosophy – Steve Shapin, one of the key figures in the Edinburgh School's SP was using the term in 1979.

[23] Key proponents of this programme are Goldman, Philip Kitcher, Hilary Kornblith and Keith Lehrer.

based institutions. It uses traditional methods of epistemology but rejects the dualism of the complementary programme.[24] (see Fig. 2)

Figure 2. Kusch's spectrum.

The science policy programme and complementary programme are species of SE as opposed to communitarian epistemology.

2.1.5 Other spectra of social epistemology

Frederick Schmitt defines SE as follows:

> Social epistemology is the conceptual and normative study of the relevance of social relations, roles, interests, and institutions to

---

[24] The complementary programme is dualistic in the sense that it conceives of two aspects of knowledge – a social and an individual – which can be analysed separately without reference to each other. That is, in some cases I can describe the epistemic status of an agent without any reference to their social status.

> knowledge. Thus it differs from the sociology of knowledge, which is an empirical study of the contingent social conditions or causes of knowledge or of what passes for knowledge in a society. (Schmitt 1994b, p. 1)

For Schmitt, social epistemology evaluates the nature of knowledge *in itself and necessarily*, whereas the sociology of knowledge describes the sociology of what *passes* for knowledge in a given society. This is similar to the distinction made by Goldman between veritistic and veriphobic approaches. The worry in both cases is that a theory which concerns itself with what people happen to think is liable to run into various relativistic problems. Schmitt also cites the distinction between understanding knowledge individualistically or socially as a useful marker. (see Fig. 3)

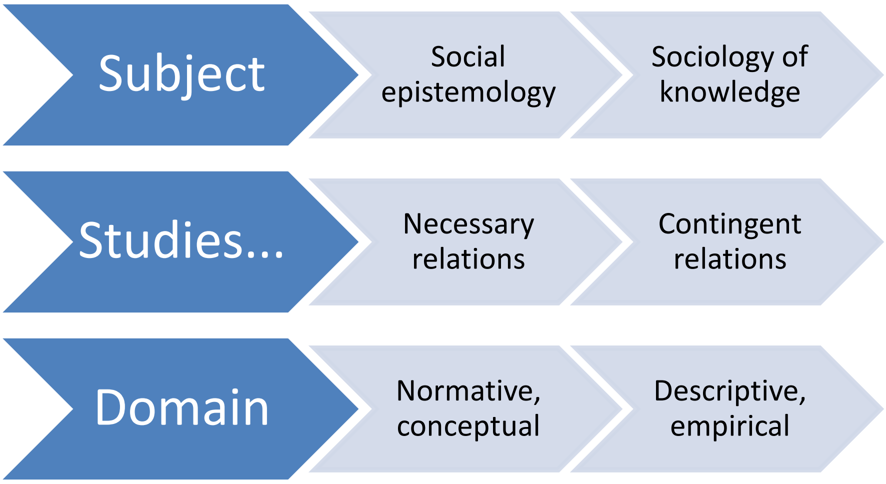

Figure 3. Schmitt's spectrum.

Philip Kitcher, like Goldman and Fuller, has ambitions that SE can be used as the 'theoretical wing' of science policy. (Kitcher 2002, p. 197. See also Fuller 1987, 1992, 1996, 2002) Like Goldman, his project is veritistic (truth-oriented). He argues that instead of formulating a context-free, general-purpose standard which evaluate the truth-promoting potential of any social practice, we ought to work by degree to develop 'standards of collective veritistic value in specific contexts and to offer formal analyses of special cases that might reveal hitherto unnoticed possibilities'. (Kitcher 2002, p. 197) He defends 'minimal social epistemology' against what he perceives to be more radical or extreme views. (Kitcher 1994) His view also rejects the traditional individualistic reduction which he believes derives from Descartes. Since philosophers are now aware of the theory-ladenness of belief we require a more sophisticated account of knowledge acquisition within a cognitive team. (See Fig. 4)

Figure 4. Kitcher's spectrum.

For the reductionist programme, belief and theory are conceived independently. For more 'social' SE and the relativist programme, belief is theory-laden.[25]

---

[25] Essentially, for Kitcher, individualistic reduction is the claim that, whilst much knowledge is achieved through a division of cognitive labour, this work is ultimately reducible to an aggregate or combination of the work of individual members of the cognitive team. Kitcher is marking the difference between social epistemology which considers only individual persons to be epistemic agents and a more radical view which contends that the sum of this individual work is less than the whole generated by the group or the even more radical view that the group itself may possess knowledge. (Kitcher 1994, 2002)

Whilst being of excellent value for their insight into the field from the perspective of weighty contributors to that field, I intend in the following to create a less entrenched survey of SE than is present in these surveys by identifying what their authors and others have in common and why they see incompatibilities between their own research and that of others.[26] I will use these sticking points to plot a more representative chart of positions in SE. This will be highly impressionistic but will serve as a 'bird's eye view' of the subject as it currently lies and may point to future developments and controversies. We can identify three sticking points and plot them as points on three axes. We can think of these axes as directions in which the author or group wish to take SE, and hence I have called this a compass rather than a spectral metaphor.

---

[26] 'Entrenched' is meant in the sense that the survey may itself be limited by the author's substantial involvement and contribution to the field. At the risk of overstating things and having made a less substantial contribution to the field, this is perhaps the one virtue I can offer here.

2.2 The compass rose of social epistemology

2.2.1 Advantages of the compass rose

Fig. 5 below is a representation of an (admittedly impressionistic) survey of most major figures in SE, broadly construed. Approximated values for each scholar were exported to MatLab (The MathWorks, Inc., Natick, MA, USA), calibrated along two axes, and designated a z-axis direction. In the first section of this chapter we discussed various surveys of SE. Some of these, such as Goldman's, are more exclusive than others. They seek to define what counts as social epistemology and what does not in a relatively narrow sense. This is a useful approach for those already involved in that narrow sliver of the field but it seems epistemically vicious to prohibit potentially fruitful approaches in the early stages of establishing a methodology and knowledge base for SE. SE does not currently have the rigorous, widely-accepted methodology, ontology, and knowledge base of the physical sciences or even of relatively new sciences such as psychology and sociology. It would be disingenuous of me to provide a survey of

SE which purports to be 'the study of the social dimensions of knowledge or information,' but does not include rivals to my own theoretical affections. Even epistemology as we know it, after all, does not have a monopoly on studying the social dimensions of knowledge or information. The survey consequently includes epistemologists, sociologists, ethnomethodologists, feminists, and cultural theorists. Those curious about the social dimensions of knowledge or information ought to consider these varied contributions as relevant.

Hence, whilst many existing surveys are limited to particular approaches to studying social knowledge my aim here is to be inclusive and general. I suggest we can identify three 'sticking points' about which various accounts diverge and plot these on a Cartesian coordinate system.

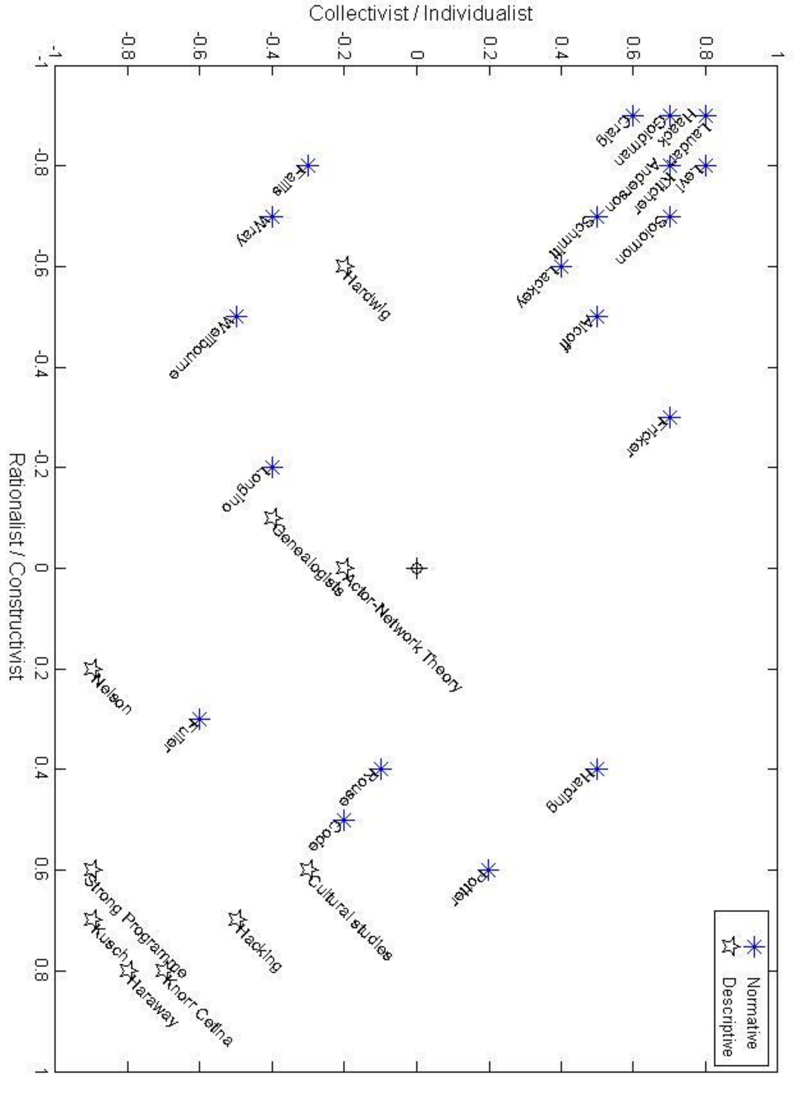

Figure 5. The compass rose of social epistemology.

The end points of these axes are incompatible. They are as follows,

i. X-axis (East to West) – rationalism and constructivism

ii. Y-axis (North to South) – individualism and collectivism

iii. Z-axis (Inwards to Outwards) – descriptivism and normativism

Of course, such values are far from scientific but the results do indicate two general groups which scholars tend to belong to. The majority of points occupy either the top left quadrant or the bottom right quadrant. That is, most rationalists are normative individualists and most constructivists are descriptive collectivists. Let us split the final section of this literature review into three sections: Rationalist Individualist Normative Social Epistemology (RINSE), Constructivist Collectivist Descriptive Social Epistemology (CCDSE), and Outliers.

2.2.2 RINSE

Each axis reflects recurrent themes in the literature. The X-axis derives from Kitcher but is also mentioned by Helen Longino who compares rational or cognitive approaches with sociological ones. (Kitcher 1994; and Longino 2002) Constructivism (or constructionism elsewhere) is often taken to be opposed to realism. Here it is opposed to rationalism because it is the name given to an idea proposed by Barry Barnes and David Bloor, encapsulated in the sound-byte: ‘there are no context-free or super-cultural norms of rationality.’ (Barnes & Bloor 1982, p. 27) By that token, rationality is the claim that there is one and only one norm of rationality summed up in Goldman’s third tenet above, viz.

> TE3: Traditional epistemology assumes that the normative standards of rationality and justifiedness are not merely conventional or relativistic, but have some sort of objective validity.

The rationalism component states, in some form, that the norms of rationality do not depend on the peculiarities of our social and political organizations. They are most readily comparable to the universalizability of mathematical axioms. There are non-contextual, invariant standards for what counts as knowledge, justified, proper or warranted assertion or attribution, and so on. Individualism regards the units of analysis as individual agents. Thought experiments, intuition pumps and descriptions of decisions, attributions, and so on, are populated by individual agents and their possession of particular properties and relations. Normative theories make prescriptive judgments of how agents ought to behave and what is or is not epistemically valuable. The combination of normativity with rationalism means that appeals to the epistemic good are typically given authority either by intuitive plausibility or derivation from unquestioned or unquestionable fundaments. Significant persons in this field include Alvin Goldman, Philip Kitcher, Larry Laudan, and Frederick Schmitt.

Descriptions of the Y-axis are present in the work of Bloor, Kusch, Schmitt, *inter alia*. (Bloor 1997; p. ix; Kusch 2002, p. 121; Schmitt 1994b & 1994c) Indeed, individualism vs. collectivism, in some form, is often portrayed as the motivation for SE. (See, e.g., Bird 2010, p. 1) However, even within SE scholars disagree profoundly on what collectivism and individualism are. In this survey, I have taken the following extremes to be emblematic of collectivism and individualism. Individualism focuses on the interactions between individuals such as through testimonial knowledge or in mass collaborations. (See, e.g., Brad Wray 2009; and Magnus 2009) A more moderate individualism might also analyse the epistemic product of aggregations of individuals, although it should always consider the primary unit to be the individual. (See, e.g., Lackey forthcoming a, and forthcoming b; List 2005, and forthcoming; and List & Pettit 2002, 2004, and 2006) Extreme collectivism reverses this order: the primary unit is now the group. (See, e.g., Kusch 2002; and Hankinson Nelson 1993) The real bearers and producers of epistemic ends (true belief, justified belief, knowledge, etc.) are communities, groups, collectives, institutions, organizations, and so on. A more moderate collectivism might also analyse individuals but would always relate this analysis to the individual's place, role and status in the group.

Another example is provided in the work J. Angelo Corlett, Margaret Gilbert, and Raimo Tuomela who endorse what Goldman would call a preservationist strategy and what Bird would call an 'analogical approach to social knowing'. (Bird 2010; Corlett 1996; Gilbert 1989, p. 313-14; and Tuomela 1992, and 2004) The general approach is to seek analogues for the individualist components of knowledge or, in other words, preserve the individualist structure. So, for instance, to take JTB, one may consider what in social knowledge is analogous to justification in individual knowledge. Most preservationists provide reliabilist accounts along precisely the same lines as reliabilist individualist epistemology namely endorsing some statement of a form similar to,

> S knows that *p* iff S's true justified belief that *p* has been produced by a reliable cognitive process (given other conditions obtain).

The "analogue" approach suffers from several problems as Bird observes. (Bird 2010, p. 39-40) First, committing oneself to an extant account of knowledge (and consequently a problematic account already subjected to counterexample and flaw) leaves one's social account open to the same objections. Second, it does not appear to reflect how we use the term "knows". There is little reason to suppose that just because we apply "knows" to both individual and social cases, that this propensity is based on our tacit understanding of a "structural isomorphism" between the two kinds of case. (Bird 2010, p. 40)

Last of all, we should consider the Z-axis which separates normative and descriptive SE. For the purposes of this survey, normative projects are those which attempt to answer questions such as: What is epistemically good or virtuous? How should we make and organize our doxastic decisions? How best should an organization be structured so as to make it epistemically optimal? As stated above, the approaches of Goldman, Kitcher, Fuller, and others have this ambition. Descriptive projects make no (or much fewer) evaluations of epistemic decisions, behaviours or organizations. It may provide social or political explanations and descriptions of what an individual or community considers being epistemically good but makes no (or few) judgments on how things *ought* to be.

2.2.3 CCDSE

Goldman's six tenets outlined in §2.1.2 provide an excellent measure by which to compare an alternative approach such as the Strong Programme whose subscribers would probably find objection in each one. SP is an example of a constructivist, collectivist, descriptivist SE and its practitioners are all situated in the South-East quadrant. (Fig. 5) Let us define what this means. Constructivism states that there are no supra-cultural norms of rationality. (Barnes & Bloor 1982) Rather, any identified norms are always relative to the specific political and social organizations within a culture or community. In some forms, this is taken to apply to all possible truths, even those of mathematics and logic. (Bloor 1986, and 1997, p. 35-42; Ernest 1998; Restivo 1992; and Triplett 1986, p. 439-50)

Collectivist accounts may view the primary unit of analysis as communities or may consider communities as a formative influence on the properties and relations of individuals. For example, Margaret

Gilbert, Lynn Hankinson Nelson, Frederick Schmitt, and Raimo Tuomela defend strongly anti-individualist epistemologies which posit collective belief and collective knowledge. (Gilbert 1987, 1994, and 2004; Hankinson Nelson 1993; Schmitt 1994a, 1994b, and 1994c; and Tuomela 1992 2004) Descriptivism aims to make no normative judgments about what agents ought to do but provide only a description of what individuals or communities do in fact take as being epistemically valuable. Most CCDSE is agnostic as regards truth and this presents a problem for traditional epistemology which takes truth to be a central component of knowledge. (Gettier 1963) The 'descriptive' element of CCDSE is essentially a non-normative, 'symmetrical' approach towards truth and knowledge thereby denying TE2 and TE3.[27] CCDSE is non-veritistic in the sense that the truth of a belief (meaning something other than 'mere' institutionalized belief, which SP roughly reduces truth to) has little or no explanatory value.

---

[27] See §3.1.3 for more on the role of symmetry in SP.

Others have called this attitude towards truth (and knowledge) social constructivism to capture the idea that truths and knowledge are *constructed* by communities iteratively, performatively, and contingently rather than representing an objective relationship between a proposition or mental state and the world. The history of constructivism spreads well beyond the reach of philosophy and sociology. If one were to argue in favour of it, one might point to its apparent uncontroversial acceptance in some disciplines, such as psychology and biology of the eye, for example. Ch. 6 considers in some detail a constructivist approach to theory of perception.[28] Perceptions have been shown in these disciplines to be constructs in the sense that they are fabricated by the perceiver's perceptual apparatus in various ways and not invariable representations of visual stimuli.

Constructivism about 'higher order' cognitive processes is much more controversial although I will defend it in this thesis. Constructivism about knowledge and social constructivism do not have the same

---

[28] Detailed discussions of constructivism in psychology are to be found in Raskin 2002, p. 1-25; and Balbi 2008, p. 15-27.

empirical scientific backing that constructivist perception and psychology do. Ch. 3 presents the constructivist approach to reference provided by the Strong Programme in the sociology of knowledge. In Ch. 4 I will argue on behalf of social constructivism about ontology and in Ch. 5-6 I argue on behalf of social constructivism about perceptually-sourced knowledge.

2.2.4 Outliers

Whilst almost all of our subjects fall into either the North-West or the South-East quadrant of the compass rose (Fig. 5), there are a few conspicuous outliers. The most notable of these I have represented as 'actor network theory,' an approach favoured by, among others, Bruno Latour.[29]

---

[29] An overview of this branch of STS which goes beyond the brief treatment here can be found in Biagioli 1999. For more on Latour's early social constructivism the reader is directed towards his 1979 (with S. Woolgar), 1987 and 1988. Other work of interest to social epistemologists is to be found in Bijker & Law 1992, Callon 1986, Knorr Cetina 1982a, 1982b, 1999; Latour 1998, 2005; Pickering 1984 and 1995; Pinch & Bijker 1984; and Rabinow 1999.

The theory was developed by sociologists of scientific knowledge partly as a critical response to those in the 'CCDSE' quadrant, particularly the Strong Programmers David Bloor and Barry Barnes. (Latour 2005) Michel Callon, Bruno Latour, John Law (and, for the purposes of characterizing a general social epistemology we might add 'semi-like-minded' contemporaries Wiebe Bijker, Trevor Pinch, and Steve Woolgar) who approached SSK and STS from an ethnomethodological, strongly constructivist perspective.[30] Of particular interest to this thesis is their work on technology, technoscience, and technical artefacts and so well shall consider some of this work in the review of PET literature, §2.3.5. (See Callon 1986; Latour 2005; and Pinch & Bijker 1984) Callon and Latour have argued that we should not conceive of nature and society (and, consequently, natural and artificial things) as distinct but of all items in our ontology as actors in a network. (Latour 1998, 2005; and Callon 1986) Indeed, Latour has gone as far as to say that non-human artefacts (such as other organisms and technical artefacts) be granted equivalent status with the human

---

[30] Latour has since distanced himself from social constructivism. (Latour 2003) For classic earlier work which inspired these relativist ethnomethodologies see Garfinkel 1967; Schütz 1962; and Schütz & Luckmann 1974. For a retrospective survey of this work see Lynch & Peyrot 1992.

observer. This became a severe bone of contention for Strong Programme researchers and produced a heated and extremely readable debate between Bloor and Latour within the pages of the journal *Studies in history of philosophy and science* (Bloor 1999a, 1999b; and Latour 1999). Bloor's criticism is taken by Latour to be a virtue:

> Latour makes no systematic distinction between nature and beliefs about, or accounts of, nature… It is as if he has difficulty telling these two things apart. (Bloor 1999a, p. 87)

> Yes, I have great difficulties in convincing myself that it is useful to create an artefact to get at the facts. (Latour 1999, p. 122)

Callon's paper—perhaps the first exposition of 'actor network theory'—describes a methodology for studying the social interactions of scientists. His proposal is an extension or generalization of the Strong Programme's symmetry principle. (See §3.1.3) Whereas the Strong Programme proposed that sociologists be agnostic with respect to the truth-falsehood, rationality-irrationality, expert-lay status of the elements of their study, Callon generalized this principle to include all natural-social and human-artificial distinctions.

> Instead of imposing a pre-established grid of analysis upon these, the observer follows the actors in order to identify the manner in which these define and associate the different elements by which they build and explain their world, whether it be social or natural. (Callon 1986, p. 4)

This methodology places actor network theory in a theoretical no man's land as far as the compass rose is concerned. They would be suspicious of the distinctions I have made and do not see themselves as constructivist or rationalist, individualist or collectivist. Being strongly agnostic and abductive they represent the extreme end of descriptivist approaches. Whether their views can be of use to analytic epistemologists is an interesting question but not one which I will be pursuing here mostly for reasons of concision. The distinction between natural and artificial things is discussed in detail in §3.1.5 and §4.2.

2.3 Philosophy of engineering and technology

> Naturalization of epistemology does not jettison the normative and settle for the indiscriminate description of on-going procedures. For me, *normative epistemology is a branch of engineering*. It is the technology of truth-seeking, or, in more cautiously epistemological term, prediction. Like any technology, it makes free use of whatever scientific findings may suit its purpose. It draws upon mathematics in computing standard deviation and probable error and in scouting the gambler's fallacy. It draws upon experimental psychology in exposing perceptual illusions, and upon cognitive psychology in scouting wishful thinking. It draws upon neurology and physics, in a general way, in discounting testimony from occult or parapsychological sources. There is no question here of ultimate value, as in morals; it is a matter of efficacy for an ulterior end, truth or prediction. The normative here, as elsewhere in engineering, becomes descriptive when the terminal parameter is expressed. (Quine, 1986, pp. 664-5. My emphasis)

2.3.1 The current state of philosophy of engineering and technology

This thesis is about knowledge in general and perceptual knowledge in particular but concerns itself more with perceptions acquired using technology than, as it were, purely biological perception (i.e. ordinary unaided human perception). This is for two reasons. The first is that although we do of course use our unaided perception continuously in our everyday lives the most interesting and novel aspects of perception and knowledge acquisition come from the increasing presence and use of technology. Secondly, after undertaking the groundwork required for this thesis it became clear what has been readily apparent to many other researchers: that it is very difficult to draw a line between which particular perceptions are 'natural' and which are 'artificial'; between which tasks are accomplished by the 'skin and skull' and which are accomplished by a co-production between skin, skull, and extra-skeletal schemata – extended cognitive systems (i.e. couplings of human cognitive processes and/or technical artefacts, other persons' cognitive processes, and the performative work of communities). (Clark & Chalmers 1998)

The extent to which technology pervades our experience and behaviour is often underestimated, from enveloping and controlling the environment around us to aiding our cognitive and physical actions to enabling new activities and opening up new fields of inquiry that were previously unavailable or simply did not exist.[31] One of the ways in which technology affects our experience and how we conceptualize our human condition is through these extensions to our existing cognitive capacities. I argue for this position in more detail in §6.1.

The philosophy of technology is a term used to describe an older discipline than the philosophy of engineering. There are differences between the two in terms of subject matter, domain and method but for my purposes here I use them more or less interchangeably. The epistemology of technology concerns what I am here calling technical knowledge (knowledge of, for and by technical artefacts) and this is primarily the possession of engineers and technologists. The subject matter of PET includes definitions of engineering and technology and demarcations of their relationship with science and other bodies of knowledge and practice; the ethics of engineering and technical professions; the nature and ethics of design; the relationship between engineering, technology and society; and much more. Construed more

---

[31] Information technology is an excellent example. For more on the 'enveloping' influence of technology see Floridi 2011b.

broadly, their domain often dovetails with that of the philosophy of science. As I discuss below, many commentators have criticized the boundary often drawn between science and engineering/technology.

2.3.2 The science-engineering relationship

Central to PET is the '*ti ên einai*…?' question: What is science and what is engineering? In 1995, the committee of the American Society of Engineering Education classified the engineering sciences into six categories:

> i. mechanics of solids, including statics, dynamics, and strength of materials;
>
> ii. fluid mechanics;
>
> iii. thermodynamics;
>
> iv. rate mechanisms, including heat, mass, and momentum transfer;
>
> v. electrical theory, including fields, circuits, and electronics; and
>
> vi. nature and property of materials. (in Downey and Lucena, 1995, p. 169)

This demarcationism was far behind the academic research on 'boundary work' which had been critical of demarcating engineering as separate from science.[32] As Downey and Lucena note, it is likely to have been politically motivated: an attempt to construct a legitimate educational framework for engineers beyond the core scientific subjects of mathematics and physics. It had been assumed up until the late 1960s (and is still assumed by many) that the distinction between science and engineering/technology was clear. Science aimed to discover the nature of reality and understand the natural world. Engineering and technology, by contrast, aimed to improve upon reality and create a new, better artificial world. Engineers and technologists applied science and appropriated it as a knowledge base to do things and to allow us to do things. (Auyang 2009b, Bunge 1966, Kline 1999) The renowned aeronautical engineer Theodore von Kármán exemplified this view when he said,

---

[32] The study of how boundaries, demarcations, and divisions between disciplines or professions are constructed, advocated, reinforced, and challenged. See, e.g. Gieryn, T. F. 1983. Boundary-work and the demarcation of science from non-science: strains and interests in professional ideologies of scientists. *American Sociological Review* 48, pp. 781–795.

> Scientists study the world as it is; engineers create the world that has never been.' (Von Kármán 1970, p. 467)

Notably, given these alleged political motivations, engineering has often been described as different from science and more closely aligned with art (for example, with reference to Leonardo da Vinci, craft and design, concept cars, architecture, and so on). The sentiment behind this is that in engineering, problems are given by societal needs and requirements. Some problems are well-defined and readily imply standard, well-known heuristics to solve them. This is sometimes known as 'cookbook engineering'. (Pitt 2001, p. 24-25)

Other problems, however, are well-defined but there is considerable and valid disagreement over the best heuristic(s) to employ in solving them. (Vincenti 1990; Bucciarelli 1996; Franssen & Bucciarelli 2004) That is, there is creativity, innovation, ingenuity, and art involved in how to solve a particular problem or how to design an artefact to best suit the societal needs and requirements. Engineers do utilize concepts and models from mathematics and physics but these sciences provide a range of models that the engineer must then use, together with her experience and judgment, to select the best for the current problem. Another prominent engineer in aeronautics, Walter G. Vincenti, quotes approvingly a British engineer speaking to the Royal Aeronautical Society in 1922,

> Aeroplanes are not designed by science, but by art in spite of some pretence and humbug to the contrary. I do not mean to suggest that engineering can do without science, on the contrary, it stands on scientific foundations, but there is a big gap between scientific research and the engineering product which has to be bridged by the art of the engineer. (Vincenti, 1990, p. 3)

Vincenti's *What engineers know and how they know it* is perhaps the *sine qua non* text in the philosophy of engineering. (Vincenti 1990) Through a detailed history and exploration of US aeronautics from 1908-1953, Vincenti constructs a distinctive account of engineering knowledge and the engineering method. He is critical of the 'applied science' view of engineering that had been prevalent for many decades, namely that engineers typically 'borrow' or appropriate their knowledge from scientists and 'by some occasionally dramatic but probably intellectually uninteresting process,' use this knowledge to build technical artefacts and shape our local environments. (Vincenti 1990, p. 3)

We would likely conclude from this, as Vincenti says, that studying the epistemology of science will be more than enough for understanding the knowledge content of engineering; the former will subsume the latter. Yet we would be mistaken do draw this conclusion. Vincenti argues that there is something unique about engineering knowledge that would warrant interesting study in the epistemology of engineering. This is contained in the engineering method. (§2.3.4) We should not think of engineering knowledge as treasure handed down by science and parasitic upon science.[33] The engineering method is a distinct method for producing knowledge from within the discipline that is not contained within the scientific output.

---

[33] §8.2.2 will discuss different attitudes towards the purpose of knowledge within different communities, professions, and cultures. Not all of these groups will take it as intuitive that knowing and inquiring are primarily aimed at true belief (as argued in, e.g., Sosa 2011). Some of these groups will take knowledge to be a kind of treasure that is acquired much as one acquires new material possessions; some will take knowledge to be functional, i.e. that it will enable them to do things; others still will take the purpose of knowledge to be to better understand the universe; and so on. See §8.2.2.

Vincenti and Edwin Layton subscribe to the thesis that technological knowledge (of which engineers are the primary possessors and, crucially, which requires an education in the engineering subjects, such as those listed above) is knowledge of how to do or make things whereas science has a more basic epistemic aim. (Layton 1974, 1976, 1991) Vincenti invoked Gilbert Ryle's distinction between knowing that (the propositional kind of knowledge with which epistemology is almost entirely concerned) and knowing how to describe the distinction between what engineers and scientists know. (Ryle 1962, p. 441; c.f. Stanley & Williamson, 2001) This seems too crude.

If we are to preserve the distinction, it should be to the extent that there is a distinction between technical knowledge, which tends to enable know-how, and theoretical knowledge, which tends to enable know-that.[34] Secondly, although Vincenti writes in support of Barry

---

[34] This is the distinction between the Greek concepts of *epistêmê* and *technê*; the two kinds of knowledge which allow one to know that things are as they are and to know how to do and make things. It is not always a useful or, indeed, an appropriate distinction as, for instance, engineers make use of ontology and knowledge of underlying mechanistic models as often as scientists make use of their knowledge of how to use instruments, create experiments, and construct.

Barnes' 'interactive model' (which purports to describe the way in which science and engineering are in mutual symbiosis rather than non-mutual parasitosis), this seems to dilute the claim that these are two separate types of knowledge, one of the central claims of his book. I would argue that although the two professions do intermingle and there is certainly no clear line to draw between the two – as if an engineer could not walk across the corridor and pick up and make use of a physics textbook and vice versa – there are good reasons for looking carefully at the particular practices and institutions that produce knowledge in a field like petroleum engineering as distinct and different from those practices in, for example, a molecular biology laboratory. I hope that these differences are made clear in Chs. 4 and 5.

For now, let us turn to how scholars have changed their minds with regard to the relation between science, technology and engineering. Towards the mid-1970s a number of scholars began to criticize the 'applied' model. In 1974, for example, Edwin Layton wrote in a landmark paper that,

> Science and technology have become intermixed. Modern technology involves scientists who "do" technology and technologists who function as scientists… The old view that basic sciences generate all the knowledge which technologists then apply will simply not help in understanding contemporary technology. (Layton 1974, p. 210)

A general consensus emerged in philosophical and sociological circles that simplistic models and generalizations ought to be abandoned in favour of studying an 'intermixed' discipline: *technoscience.*[35] This was, according to Barry Barnes, a 'major reorientation in our thinking about the science-technology relationship'. (Barnes 1982, p. 166) He wrote, referring in part to a traditional attitude in the UK and elsewhere that engineers and technologists were considered to be intellectually or professionally inferior to scientists. In the UK it is common to see people's surprise when someone introduces themselves as an engineer and then describes how they are working on, say, diagnoses of flat feet and pedal movement, or designing robots for handling nuclear material. For more common for the person to assume that a person who introduces themselves as an engineer will be able to help with

---

[35] The term 'technoscience' was coined by Gilbert Hattois to encapsulate this interactive relationship. (Hattois 1984, p. 59-60).

household plumbing or repairing their car. This apparent intellectual snobbery has given the profession a desire to establish itself as as institutionally significant as the various scientific disciplines.

Somewhat counter-intuitively, one of the consequences of the increasing respect given to engineering as a discipline of equal intellectual significance to science in its own right, is the trend to analyse science and technology as converging upon a single point – a new technoscientific discipline which incorporated the reciprocal interaction between both disciplines.

> We recognized science and technology to be on a par with each other. Both sets of practitioners creatively extend and develop their existing culture; but both sides also take up and exploit some part of the culture of the other…. They are in fact enmeshed in a symbiotic relationship. (Barnes 1982, p. 166)

Layton had described the view that, if there is a distinction between science and engineering, that it is social and not based on 'abstract functions of knowing and doing.' (Layton 1974, p. 209) Researchers were criticized for assuming that science and technology were 'well-defined monolithic structures' (Bijker, Hughes, & Pinch 1987, p. 20) and failed to notice the social constructions that gave them that appearance. (Mayr 1976) Sociologists such as Barnes, Wiebe Bijker, and Trevor Pinch suggested that Science and Technology Studies should use social distinctions rather than abstract-theoretical ones. They argued, indeed, that there were no underlying distinctions beyond the particular, contingent distinctions that arise out of social negotiation and organization.

Likewise, Vincenti argues that the engineering method and engineering knowledge are not the same as, or derived uninterestingly from, the scientific method and scientific knowledge. He quotes the British engineer G. F. C. Rogers saying,

> Engineering refers to the practice of organizing the design and construction [and, I would add, operation] of any artifice which transforms the physical world around us to meet some recognized need. (Vincenti 1990, p. 7. Original parenthesis.)[36]

---

[36] For further 'boundary work' done by Vincenti in demarcating the science-engineering relationship see p. 11, 112, 151, 166, 170, and 200-58.

Indeed, towards the end of the book he provides a case study which purports to show that sometimes science plays 'no role' in solving the problem. (Vincenti 1990, p. 193) Vincenti starts from the appealing proposition that, if technology is merely applied science, then there is no unique form of knowledge that could account for the many achievements of technology which are independent of scientific discovery, such as the pyramids of Egypt or the roads of ancient Roma. (Pitt 2001) Vincenti's own case studies from the history of aeronautic engineering appear to back this up. One can make technological discoveries that are not scientific discoveries. Consequently, despite his admiration for the interactive, mutual, technoscientific model of Barnes et al. Vincenti also wants to claim particular territory for engineering knowledge – what has been called the 'epistemic emancipation of technology'. (Houkes 2009, p. 310)

Layton had similarly wanted to position engineering as its own discipline and area of knowledge. He argued that it may be divided into two types:

i. the 'less idealized natural sciences', and

ii. sciences that seek 'to gain a scientific understanding of the behaviour of man-made devices,' such as parts of thermodynamics that advance idealized models of heat engines. (Layton 1974, p. 10)

Layton's concept of a 'less idealized natural science' is perhaps better illustrated by thinking of scope and scale. When we describe a mechanistic model in, say, fluid dynamics, as a true model of the viscosity and flow of a fluid our expression 'true' is indexical to the scope of the investigation. It would also be possible to describe the behaviour of the fluid at a Newtonian or Quantum level but this is not needed for statements in fluid dynamics to be true. Here we assume that the matter has some sort of heterogeneous microstructure even though a model calibrated for a different scope of investigation may reveal that it does not. In engineering, whether something is true depends on the scope of the investigation. This is, at least on some accounts, different from how we normally think of scientific knowledge. If we read our physics textbook and it tells us that the density of a material is equal to its mass divided by its volume then we do not think this is true only given a certain scope of investigation but true simpliciter. Likewise, it we learn that blood cells are produced in mammals by haematopoiesis, then we do not take this to mean that this process produces blood cells given a certain scope of investigation but that the process produces blood cells no matter at what level we wish to look at it. In engineering, the very opposite can be true. What is true of a substance given the methods of fluid dynamics is not true were one to undertake a more atomistic level of inquiry. We will also discuss later how, in some cases, what is true depends on geographical location, contingent operating procedures, and what a particular group of professionals takes to be true.[37] Denying this, generally speaking, is to deny the existence of the body of knowledge Layton describes under (i).

---

[37] Whilst it is, of course, trivially true that scientific theories may also take account of geographical location it is not typically true that they admit of difference across different geographical locations depending on the particular historical and pedagogical differences to be found there. In other words, as I discuss below, whilst there is no such thing as French physics or Russian chemistry, I have observed in my fieldwork that there is such a thing as Scottish and American logging analysis.

By contrast, science is, usually, concerned with general, non-geographical, absolute principles. As Sunny Auyang writes,

> *Science* generally means state of knowing or possessing knowledge that is sufficiently general, clearly conceptualized, carefully reasoned, systematically organized, critically examined, and empirically tested. (Auyang 2009b)

Finally, we should also be careful not to characterize engineering as a heterogeneous discipline. Larry Bucciarelli has described the various design processes used in engineering and demonstrated that no one process is dictated by the nature of the object being designed or the problem to be solved. (Bucciarelli 1996) Central to the philosophy of engineering and technology is the concept of design. Proper function is thought, in part, to be determined by design considerations. Design, in this instance, can mean both the content of a set of places (as in 'the design for a new production logging tool') and the process by which those plans are realized, which typically will involve specifications of the configuration, layout, and dimensions of the artefact; testing of the artefact to see if it fulfils its function either through mathematical analysis or experiment; and modification. (Vincenti 1990, p. 10) Larry

Bucciarelli (1996, 2003, 2009), Edwin Layton (1974), Walter G. Vincenti (1990), and others have placed design at the centre of their philosophy of engineering.

In fact, Bucciarelli and Vincenti both almost exclusively restrict themselves to talking about *design* knowledge, it is such a central component in their work. In *The origins of the turbojet revolution*, Edward Constant makes a distinction (echoing that made in the philosophy of science) between 'normal' and 'radical' technology. (Constant 1980) Normal technology, he states, is what engineers and technologists spend the vast majority of their time doing. It involves practicing and making small improvements to the accepted traditional practices of the community. Radical technology, on the other hand, occurs when technological communities face a problem which cannot be solved using traditional methods and must innovate to produce an entirely new artefact, the configuration of which, and even how it works, is largely unknown to the community.

To this concept of normal and radical technology, Vincenti adds normal and radical design. (Vincenti 1990, p. 10) Normal design adheres to the accepted customs and practices of the community when designing an artefact. The type of artefact, how it functions, and how it should be used, is well understood. When an engineering problem requires the invention of a new kind of artefact, the mechanics and functions of which are little understood, practitioners become involved in radical design. Both design paradigms are multilevel and hierarchical. He describes the levels at which design takes place in aeroplane design and how problems at the higher levels (including the project definition and the general overall design) tend to be 'conceptual and relatively unstructured. (Vincenti 1990, p. 12) This typically includes the translation of military or commercial requirements into some technical problem.

These are socially-sourced requirements and, as such, must be interpreted by the team of engineers into a problem of engineering science. As you can imagine, this translation may be normal (reducing to well-known engineering principles and heuristics) or radical (such as the cases described in Constant 1980 or Bloor 2011). Vincenti is keen to emphasize the effect of social concerns on whatever design approach is taken,

> Engineering knowledge reflects the fact that design does not take place for its own sake and in isolation. Artifactual design is a social activity directed at a practical set of goals intended to serve human beings in some way. As such, it is intimately bound up with economic, military, social, personal, and environmental needs and constraints. (Vincenti 1990, p. 11)

If we accept Vincenti's claims we should be careful not to separate social from non-social elements. In fact the two are inextricably linked.

2.3.3 The myth of the lone engineer/scientist/epistemic agent

The theoretical physicists Freeman Dyson once said, 'A good scientist is a person with original ideas. A good engineer is a person who makes a design that works with as few original ideas as possible.' (Dyson

1979, pt. 1 ch. 10) Dyson here suggests that science when it is at its best creates new ideas, unlike engineering which adapts existing ideas as expediently and frugally as possible. Science is about discovering knowledge, as opposed to the technicians in engineering who apply it in an efficient and expedient manner. There are two things wrong with this picture. The first is the reliance on an interpretation of engineering which does not include technical knowledge. We discussed this in §2.3.2 so there is no need to retrace it here. The second is its individualist bent. The myth of the scientist as a lone man looking at the world around him, drawing up hypotheses, and then testing them against the evidence in a dispassionate and amoral manner bears little relation to reality. As a society, we have convinced ourselves that discovery and knowledge are produced by individual flashes of inspiration in the mind of a single genius. The tale of Friedrich August Kekule von Stadonitz, the 19th century organic chemist who discovered the molecular structure of benzene after dreaming of a snake biting its own tail, exemplifies this idea which is repeated throughout scientific historicizing as well as in popular cultural heroes such as Sherlock Holmes or Dr. Gregory House.

Of course, there is no doubt that science does occasionally leap forward as a result of a 'Eureka' effect but most progression is through collective step-by-step advancement; a tower of ancestors each standing on each other shoulders. (Merton 1965) Not only is the lone scientist rare but it, if it was common, science would progress at a slower rate. Dunbar shows that teams of scientists have a significant advantage over individuals pursuing the same task but only when the members of the group have different areas of expertise. (Dunbar 1993) He suggests that this may be because progress and discovery in science often depends on analogy and application of knowledge from other disciplines to the current problem. Where members of the group all have similar 'analogical reserves' to draw upon, this confers little advantage above and beyond the work of any single member of the group.

Equally, there existed (and does exist) a mythical engineer: again, typically, a man acting alone. This man is personally interested in 'how things work' and making: *homo faber* rather than *homo rationalis.* Billy V. Koen, for instance describes the engineer as a pragmatic individual using a set of heuristics to 'cause the best change in a poorly understood situation within the available resources.' (Koen 2003, p. 28) Whilst this is a good definition of the engineering method, it is a poor definition of the distributed reasoning that most commonly characterizes that method. This involves drawing on the resources and background knowledge of others—the analogical reserves—to solve problems.

This myth is not the fault of scientists or engineers or even our lay understanding about science and engineering. It exists in philosophy as well, suggesting it reveals something much more deeply grained in our understanding of us and our place in the world. As we have seen, part of the motivation behind social epistemology is to challenge epistemologists who draw conclusions based on lone, isolated individuals, single-handedly inquiring into the nature of some fact or a-socially and apolitically going about their business, as if they were always behind the veil of ignorance. (Harsanyi 1955; and Rawls 1971) No doubt these individualist epistemologists are partly motivated by Occamian considerations but it seems that social epistemology, which gives consideration to more complex social organizations and individuals' places in them, has solid grounds to justify its position as an important addition to the individualistic epistemology, particularly in the fields of science and engineering.

2.3.4 The engineering method

Another reason to distinguish engineering from scientific knowledge is by its dependence on scarce resources. Engineers are tasked with 'improving' the world in some way and consequently desire a method that would maximise their chance of success, but also to minimise risk, cost, time, and so on. Demands for these improvements typically come from wider society, governments, commerce, industry, military, and so on, as do the constraints. Consequently, the method is always a sacrificial, compromising, social heuristic: you cannot please everyone all the time but you can please some of them some of the time. Many cases of scientific knowledge proceed in exactly the opposite manner. In answering the question of how the universe began, for example, we do not need an answer tomorrow, or even within decades. Instead, we hope to arrive at better and better answers as our scientific knowledge advances. In other words, science and engineering often differ in terms of their method or heuristic. Here is Billy V. Koen's definition of *heuristic:*

> A *heuristic* is anything that provides a plausible aid or direction in the solution of a problem but is in the final analysis unjustified, incapable of justification, and potentially fallible. (Koen, *Discussion of the Method*, p. 28)

Other engineering heuristics might include ‘allocate resources to the weak link’, ‘safety first!’, ‘improve upon designs incrementally’, ‘rules of thumb’, and so on. It is easy to see why these methods appear to suggest that engineering is very different from science. Further, whereas in science we tend to assume that there is one correct answer for each problem, if there is an answer at all, different engineers may approach a problem in different ways and be able to provide justifications as good as each other for their solutions. Thus, Koen suggests the following definition of the engineering method,

> …the use of heuristics to cause the best change in a poorly understood situation within the available resources. (Koen 2003, p. 28)

These heuristics are inevitably dualistic. They are based partly on knowledge of physical structure and dynamics as well as knowledge of social needs and requirements. (Kroes 2010; and Kroes & Meijers 2006) John Dewey once argued that whilst one can quite properly speak of a national philosophy, the same cannot be said of science: there is no 'Indian neuroscience' or 'Japanese biology'. Science is the same all over the world; or, at least, it ought to be the same even if individual nations or communities will have their own traditions and histories. Whilst we might wish to argue this point, Dewey could not have denied that there are national engineering practices. Each has acquired their own set of heuristics – dependent on school, university, practical experience, geographical conditions, available resources, moral attitudes, and much more - which they use to solve problems.

In my fieldwork research I was taught by and interviewed Scottish, American, Thai, Malaysian, Filipino, and English engineers. Most had worked in a variety of different countries and, in particular, many had undergone training in the United States. It was clear that different countries and communities had their own favoured approaches or 'styles'. Many of these may not be at all easy to study or codify in the way that a science studies researcher might be able to ask a chemist, for example, why they do experiments on particular animals and not others or why they test their equipment in such a way. Estimations, judgments, figures, and so on may be made on the basis of extensive practical experience and, crucially, trial and error, optimization, and the techniques used in design synthesis. (Vincenti 1990, p. 166) There are many methods and no one correct approach in engineering for a problem: the best approach is the one that is likely to bring the best results or, in hindsight, brought about the best results. Another aspect that must be emphasized is that some problems are 'pressing' in a way that some scientific questions—such as how was the Universe formed or why humans have two legs—are not. (Lamberth 1999, p. 57) They call for a decision to be made within a specified time frame and using certain limited resources.

Knowledge attribution and inquiry are inextricably linked. In fact, it has been argued that inquiry is the central focus of knowledge attribution. (Hookway 1999, p. 7) Further, it seems plausible that one could construe all knowledge as the answer to a question and inquiry is the process by which one finds an answer to that question. (Pritchard 2008b, p. 232; Schaffer 2007) Thus, changes in the context of inquiry entail changes in the context of knowledge ascription. When we ascribe knowledge we, at least in part, assert that the attributee can, generally, answer a question. (Schaffer 2005, p. 117)

Now, consider the following example:[38] Suppose I ask an audience of engineers and non-engineers to tell me how many ping-pong balls you could fit in this room. Typically, according to Koen, most non-engineers will not put their hand up to offer an answer. The engineer, at least so the story goes, always will. The quality and nature of the answer depends, in part, on resources, time and the importance or pay-off awarded for a good solution. If we gave the audience 30 seconds

---

[38] The example is taken from a presentation by Koen to the Workshop in Philosophy of Engineering at the Royal Academy of Engineering, London, 10th November 2008, although I believe he took the example from a discussion at the Massachusetts Institute of Technology.

we could expect a quick back-of-the-envelope calculation. If we gave them a couple of minutes and a tape-measure, we might expect them to quickly measure the walls, estimate the dimensions of a ball and derive a figure from the difference. If their lives depended on an accurate answer and they had several months, we might well expect them to fill the room with ping-pong balls and count them or go to even more extensive measures.

It is an anecdote but it carries a point: that an engineering answer, a perfectly correct engineering answer, depends on the resources available to dedicate to that problem and that the initial back-of-the-envelope calculation may be a perfectly good answer relative to that context. The answer the engineer gives is never *the* answer to a problem, but it is her best answer to the problem she is given - all things considered. (Koen 2003, p.61)

This seems to support a contextualist account of knowledge attribution. (See §3.2) For example, in any given situation, I can ask, 'How many ping pong balls could one fit in this room' and attribute knowledge based on context-relative criteria such as the profession of the speaker or the importance of the outcome. We can illustrate this using the high-low framework of some contextualist theories. These hold that there is a high context, under which the conditions for ascribing knowledge are strict and demanding, and a low context, under which the conditions for ascribing knowledge are not so demanding.[39] If the answer I am given is 'About 10,000' and I am in a high-stakes context (i.e. it is important that I am given a precise answer to the nearest 5 balls and I require a precise and demanding justification) I can appropriately infer that my interlocutor does not know how many ping pong balls could fit in this room. If, on the other hand, I am in a low-stakes context (i.e. all I require is a 'ball park figure'

---

[39] Different authors spell this out in terms of high and low 'standards', (Cohen 1987; DeRose 1992, 1995, 1999b, 2002, 2009, 2012; Williams, M. 2001) 'stakes', (Leite 2005; Nagel, J. 2008; Stone 2007) 'thresholds', (Cohen 1998) 'practical interests' (Stanley 2005), as well as contexts of 'alternatives' (Dretske 1970, cf. Schaffer 2001), 'warranted assertability' (Blaauw 2003, Brown 2006, Stone 2007; cf. Baumann 2011, DeRose 2002, Leite 2005, and Pritchard 2005 for critical discussion), and so on. Schaffer (2005) observes that there are fine details to be made between what it is that shifts relative to a context. He further argues that defending the claim that *x* shifts relative to a context must 'illuminate inquiry.' (Shaffer 2005, p. 117) That is, whether we focus on stakes, as I have here, standards, thresholds, practical interests, or whatever else, our account of this variable must connect to the role that knowledge ascriptions have more widely in practices of inquiry. Although it is of course true that one can know something and not be able to answer a particular question (either out of shyness, or muteness, or whatever else), the purpose of ascribing knowledge is, in part, to assert that the attributee can, generally, answer a question.

or rough estimate and a less demanding justification), then it may be appropriate for me to say that my interlocutor does know how many ping pong balls could fit in this room. §3.2 considers the validity of contextualism in this context in detail, its merits as an epistemology, and its detractions.

With all this in mind, then, how does engineering knowledge differ from scientific knowledge? Is it accurate to say that all engineering knowledge consists in the content and application of heuristics? The first thing we can say is that engineering knowledge is does not attempt to produce accurate (in the sense of true, in the sense of corresponding to reality) models of the phenomena it studies. Rather it produces useful models which can aid engineers in accomplishing their ultimate aim, whether that is to extract hydrocarbons from a subterranean rock or to build a gas pipeline and so on. The model is instrumental to the problem-solving process; it is not the product of the inquiry. Further, there is no one-to-one correlation between the process of inquiry and the model as there is in the natural sciences; the engineer may select from a number of available models. The second thing is that what counts as a correct procedure of inquiry depends on a lot of extra-theoretical factors: time, resources, location, personnel, and so on. Likewise, what counts as a good answer depends on the same extra-theoretical factors.

2.3.5 RINSE and CCDSE in philosophy of engineering and technology (PET)

Much of the work in characterizing the epistemology of technology has been done within schools of sociology and history, rather than analytic philosophy, particularly regarding CCDSE. Trevor Pinch and Wiebe Bijker argued in their seminal paper that the constructivism which had been so productive in the sociology of sciences can be fruitfully applied to an empirical analysis of technology. (Pinch & Bijker 1984) Pinch and Bijker tick all the relevant boxes for constructivist, collectivist, descriptivist social epistemology and their work has been prominent in much of the subsequent research. Generally, we can segregate RINSE and CCDSE in PET in an analogous manner to how they were separated in §2.2.[40] CCDSE in PET is suspicious of 'rationalistic' or 'veritistic' concepts used in explanation of why a theory succeeds or in normative prescriptions for scientific and engineering practice. These concepts include logical inference from theory, inference to best explanation, parsimony, commensurability, derivation from true axioms, falsifiability, induction from experience, and so on. CCDSE typically states that a philosopher, sociologist or historian who studies scientists and engineers should look at the practices involved, the behaviour and relationships within scientific and engineering communities, the interactions between people, the settling of disputes, and the resolution of problems. What one should not do is appeal to logical abstractions, idealized visions of the scientific method and Platonized interpretations of what scientists and engineers say, to explain the success of theory, progression, and design.

---

[40] Collections of CCDSE articles in the philosophy of technology can be found in Hackett, et al. 2008; Kaplan 2004; and Scharff & Dusek 2003; and Barnes, Bloor, & Henry 1996. See also Mitcham 1999 who presents a history of PET and constructs his own philosophy of engineering and technology from a Continental philosophical perspective. An analytic approach can be found in Kroes & Meijers 2000. See also Ch. 4 of this thesis.

In the last couple of decades, more attention has been paid to PET in analytic philosophy. These researchers seek to use the methods of analytic epistemology and metaphysics to deliver solutions to outstanding problems in PET. That is, instead of the relativist, sociological/ethnomethodological model presented by CCDSE, they evaluate technological change using the concepts of ontology, rational design, know-that/know-how, and so on. One example of this work is discussed in detail in §4.1.2. Another example is the work, already mentioned, of Walter Vincenti and Edwin Layton. These two authors preserved a distinction between scientific knowledge and technical knowledge. To defend this epistemic emancipation, Layton and Vincenti pointed to examples where engineering and technology have been responsible for expanding our knowledge and outline the cognitive and heuristic differences between an engineer and a scientist.

> Such a view immediately defines the science-technology relation—technology is hierarchically subordinate to science, serving only to deduce the implications of scientific discoveries and give them practical application…. Such a hierarchical model leaves nothing basic to be discussed about the nature of the relationship. A model with such rigidity is bound to have difficulty fitting the complex historical record. (Vincenti 1990, p. 5)

> We might restate the matter by noting that the laws of science refer to nature and the rules of technology refer to human artifice. The function of technological rules is to provide a rational basis for design, not to enable man to understand the universe. The difference is not just one of ideas but of values; "knowing" and "doing" reflect the fundamentally different goals of the communities of science and technology. The thought that embodies the values of technology will relate to active and purposive adaptation of means to some human end, that is, it will relate to design. (Layton 1974, p. 39)

Both Layton and Vincenti demarcate engineering as separate from science and direct their research towards on understanding of those differences, particularly the relation of engineering design and the growth of knowledge to social demands. Ch. 4 and 5 are, in many respects, repeating these differences and marking them out as worthy of epistemological study—specifically, social epistemological study. I argue that there are such things as technical artefacts. This can be demonstrated by showing that we talk about them in different ways. These differences in reference-talk (see §3.1) lead to the possibility of different kinds of knowledge, such as technical knowledge.

Having said that, the relationship between scientific and technical knowledge is rarely clearly distinct—there are engineers and technologists doing science as well as scientists doing engineering and technology—and several authors are skeptical as to whether we should preserve any such distinction at all (e.g. Barnes 1982b; Latour 1987; Latour & Woolgar 1979). It is also worth noting, contra Layton and Vincenti, that particular scientific worldviews have had influence on engineers and technologists and that much engineering and technical knowledge can ultimately be shown to be shaped by underlying scientific ontologies.[41]

---

[41] It is easy to think of historical examples of this, but for some recent examples see Zelić and Stahl (2005) who show that the Irish government's decision to implement an electronic voting system was (mis)guided by an ontological commitment to realism; and Hongladarom (2002b) who demonstrates that technology and engineering in Thailand has been adopted from research and development in other countries and has developed in a different manner due to its ontological commitment to Buddhism rather than a 'Western' modern scientific worldview. See also Reynolds (1976) and Day & Reynolds (2000) who describe the adoption of Western cosmography in Thailand.

2.3.6 Conclusions of literature review

This chapter presented two literature reviews: a survey of work in social epistemology and a review of the philosophy of engineering and technology. In the first section, I discussed the aims and methods of social epistemology, broadly construed. I then presented a novel way of surveying the field through a 'compass rose' metaphor. This provides a useful way of distinguishing between different approaches to social knowledge and the various 'sticking points' at which different approaches diverge. I described three such sticking points and their implications: rationalism vs. constructivism, individualism vs. collectivism, and normativism vs. descriptivism. It was found that most scholars tend to fall into one of two camps: a rationalist, individualist, normative camp; and a constructivist, collectivist, descriptive camp. This is not to say that all social epistemologies must fall into one or the other of these categories (indeed, some represented in Fig. 5 do not). It is to say that these are two prevailing, incommensurable trends in social epistemology.[42]

The third section reviewed the history and current state of the philosophy of engineering and technology (as well as some related disciplines). I discussed the relationship between science and engineering and how attitudes to this relationship have changed over the last half decade. Others have argued that the distinction between science and engineering ought to be abandoned. (Barnes 1982b; Latour 1987; Latour & Woolgar 1979) These authors tend to cite the intermingling of scientists and technologists and the transfer and convergence of knowledge across multiple disciplines to argue against a distinction between science and engineering. However, this is to ignore the valuable work which demonstrates that engineers create a distinctive kind of knowledge that cannot be explained under the terms of history, philosophy, and sociology of science. I intend this thesis to

---

[42] Work which discusses whether two representatives from each of these camps — the Strong Programme in the sociology of knowledge and analytic social epistemology — really are incommensurable is discussed in §3.1.3.

provide one example of this unique kind of knowledge, which I am calling technical knowledge.

Chapter 3

Communitarian and contextualist epistemology

3.1 Knowledge by agreement

3.1.1 Introduction, commitments and politics

Having provided a general survey of social epistemology, this chapter focuses on Martin Kusch's communitarian epistemology as a promising means of navigating between the three sticking points. Kusch explicitly attempts to bridge various opposing sides in SE. I discuss Kusch's communitarian roots in the sociology of knowledge and some important analytical concepts that will be used in following chapters including the ontology of artefacts, collective performance, and reference-talk. In the second section I compare communitarianism with contextualism, a similar but distinct movement in epistemology. The two are found to be incompatible although aspects of contextualism are relevant for illuminating the central theses of epistemological communitarianism.

The epistemological programme advocated in the 2002 book *Knowledge by Agreement* by Martin Kusch commits itself to two principles:

> CE1. The term 'knowledge' and its cognates like 'know' and 'knower,' mark a *social status* – like 'head of department.'
>
> CE2. The social status 'knowledge' is typically granted to, or imposed on, *groups* of people. Knowledge is not just social in that it is a social status; it is also social in that it is typically attributed to groups rather than to individuals (Kusch 2002, p. 1, original emphasis).

In Kusch's view, many (if not all) social epistemologists have departed from the initial radical claims that the movement promised and his book serves as a redress. Kusch views SE as a spectrum of positions which give varying degrees of privilege to social and political explanations in the study of knowledge. (See §2.1.4 and Fig. 2) He notes that the term has been applied to both philosophical and sociological theories and each end of the spectrum represents the methods of each of these disciplines. The most prominent figure in the sociological movement has been Steve Fuller whose prescriptive programme Kusch terms the ‘science policy programme’. The most prominent figure in analytic philosophy has been Alvin Goldman who also has prescriptive ambitions for a ‘complementary programme’. Goldman’s expansionist complementary programme seemed to leave a window open for a radical critique of the tenets of traditional epistemology from within the framework of analytic epistemology (see §2.1.2). Kusch’s epistemological communitarianism (henceforth, e-communitarianism) purports to offer such an approach. It also distinguishes itself from Fuller and Goldman by being an explicitly descriptivist programme and, therefore, having no ambition to prescribe or recommend policy change.

Consider C_E1 and C_E2. The first identifies one crucial subject matter: the shared practices and language of social roles that confer status upon an agent. The second subject matter involves reversing the traditional order of explanation by claiming that ‘knowledge’ is most often granted to communities rather than individuals thus contravening

Goldman's TE1. (§2.1.2) These two tenets demonstrate how difficult it can be to steer clear of political theory when discussing knowledge simpliciter. Despite being a descriptive theory of knowledge and not advocating any particular policy or rearrangement of social organizations (as argued for, e.g., by Goldman), e-communitarianism does have a political element in its origins, its interdisciplinary framework, and its theoretical concern with differentials of social power. Kusch writes at the beginning of the book,

> …epistemology and politics are more closely connected than tradition would have it. To understand knowledge is to understand epistemic communities: and to understand epistemic communities is to understand their social and political structures. (Kusch 2002, p. 2)

Statements such as these are part of a wider trend in epistemology which acknowledges the Hobbesian-Baconian maxim that knowledge is power.[43] This is particularly relevant when discussing knowledge produced in an artificial social organization such as a science laboratory

---

[43] See Fricker (2007) where the author makes the claim that any SE must be political since epistemic interdependence is connected to social power structures. Standpoint and feminist epistemology also display a commitment to tease out the connections between epistemic agents, social organization, and politics. See, e.g., Alcoff 2001; Anderson 1995; and Harding 1991.

or an engineering office. The role that trust, authority, expertise, and hierarchy play in generating knowledge is discussed at various points throughout the case study.

There are two schools of thought that have influenced *Knowledge by agreement* more than any others: communitarianism in political philosophy (henceforth, just communitarianism) and the Strong Programme in sociology of scientific knowledge. I will now turn to each of these in turn.

3.1.2 Communitarianism

Epistemological communitarianism alludes explicitly to political communitarianism, the movement in political theory which opposed John Rawl's claim that the primary function of governments is to distribute resources according to principles of fairness. (Bell 2012) The reader might ask why anyone would oppose such an obviously reasonable proposal. For Kusch, the pertinent insight of this movement that epistemologists ought to heed is that the community is prior to the individual in the order of explanation. Put another way, the task of epistemology is to 'explain why individuals can know only in so far as they are members of epistemic communities' (Kusch 2002, p. 10).

There is more to the affinity than this, however. Alasdair MacIntyre, Michael Sandel and Charles Taylor, three key political communitarians, argued against universalism and individualism and emphasized the role of the community in cementing political and moral values. (Bell 2012) Kusch deploys similar arguments against universal claims in epistemology in the form of global theories of justification as well as foundationalism and coherentism which base universal theories of knowledge on universal features of human psychology. The communitarian dispute with Rawls' 'original position' thought experiment – as well as his description of an 'Archimedean point' from which a person may regard the human condition 'from the perspective of eternity' (Rawls 1971, p. 587) – is comparable to Kusch's criticisms of individualist epistemology and its de-contextualized thought experiments.[44]

By the same token, a project which aims to recapture knowledge as a social institution is the epistemological (and sociological) equivalent of the communitarian project to explain political and moral values via an explanation of specific local contexts in which political communities and institutions are created. In short, the take home message of communitarianism and the sociology of knowledge for epistemology is that knowledge is a social institution created within particular communities in diverse, locally-contingent ways. It is the specific determinations of these systems that should be understood before we can properly understand how individuals engage in epistemic activities: acquiring or generating knowledge; ascribing knowledge to others both within and without their own community; justifying these ascriptions and their own claims to knowledge against the challenges of others; and so on.

---

[44] Note that Rawls later recanted from the import of this thought experiment for his theory of justice, describing it as 'a useful rule of thumb' (Rawls, *Justice as Fairness: A restatement*, p. 97).

3.1.3 The Edinburgh School: the Strong Programme in the sociology of scientific knowledge

The second theoretical influence on e-communitarianism is the Edinburgh School or Strong Programme in the sociology of scientific knowledge (SP). SP is a research area in the sociology of scientific knowledge developed by a group of sociologists at the University of Edinburgh in the 1960s including Barry Barnes, David Bloor, John Henry, and Donald MacKenzie. Kusch believes, and this dissertation follows through some of these claims, that some of the most urgent challenges to real epistemology do not come from disciplines which philosophers most readily and commonly ally themselves with: cognitive science, evolutionary biology, psychology, and so on. These subjects will never give comprehensive answers to epistemological questions because they are predominantly individualistic in their outlook. They can answer questions about an individual's cognitive processes, brain events, and thinking but cannot answer questions of socialization, status, expertise, and so on. Rather they come from the sociology of knowledge, particularly its attempts to provide empirical methods for analysing the function of communities and epistemic teams.[45]

SP originated as an attempt to correct a common fallacy perpetrated by sociologists, historians and philosophers of science. These scholars often gave sociological explanations for the development of scientific theories but reserved such explanations for the unsuccessful ones – eugenics, the Lysenko affair, phrenology, etc. Successful and surviving theories could be explained by appeal to their truth-value, natural facts, rationality, objectivity, logical inference from theory, inference to best explanation, and so on. This was taken by SP to be asymmetrical, given that scientific theories have historically tended to replace others even when substantiated by the rules given above. Sociological explanations ought to be applied across the board in explanations of any theory succession or cessation. Cultural values, self-interest, commitment, education, trust, expertise, education, history, and so on, are always

---

[45] See, e.g., Barnes 1982a; Barnes, Bloor, & Henry 1996; Bloor 1983, 1991; and MacKenzie 2006.

relevant explanatory factors. The myth (often replicated even in influential works on SP such as R. K. Merton's *The Normative Structure of Science* 1979) of the objective, detached scientist ought to be dispelled and replaced by a conception of scientists as human beings: 'highly gregarious, interdependent social primates.' (Gagnier and Dupré 1998, quoted in Barnes 2000, p. ix)

Interestingly, the connections between SP and communitarianism are also politically vexed. The term 'communitarian' has a history in political theory also dating back to the 1960s where researchers used it to replace the 'dubious' term *communist* that would otherwise represent collectivist or community-centric sociologies. R. K. Merton, the early sociologist of science who strongly influenced the formation of SP, defined four universal norms of science: Communism, Universalism, Disinterestedness, and Organized Skepticism (known as CUDOS). (Merton 1979. See Mitroff 1974 and Mulkay 1991 for some compelling counter-arguments.) Merton defined the norm of communism as,

> The common ownership of scientific discoveries, according to which scientists give up intellectual property rights in exchange for recognition and esteem.[46]

This criterion was renamed communitarianism or communalism by successors embarrassed by possible connotations with Leninism. Essentially, Merton believed that scientific knowledge is a public good. Scientists consider, according to Merton, that their epistemic product is free and only ask that they are rewarded with recognition of their achievements. Whether or not Merton believes that this is how scientists *do* feel or rather how they ought to feel is not clear.

Kusch uses the term ‘communitarianism’ to distinguish his approach not from Soviet-era politics or Marxist ideologies but from a particular class of social epistemologists, viz. Goldman’s real epistemologists (§2.1.3). Kusch wishes to take seriously SP’s criticisms of real epistemology and use them to develop a sociophilosophical epistemology. (Kusch 1997) According to this account, real

---

[46] It is interesting to note at this point that Merton identified communism (communitarianism) with common ownership of discoveries (or the collective bearing of knowledge).

epistemologists such as Alvin Goldman are wrong to speak of a social 'dimension' or 'aspect' to knowledge which is deserving of additional research. For real epistemologists, the social is an aspect or feature of knowledge which epistemologists have not much noticed and ought to expand their research into. For SP and e-communitarians, on the other hand, all knowledge is social. That is, there is no knowing subject without a shared context of knowledge ascriptions. Knowledge generally requires learning, education, and teaching, and these comprise socialization. (Durkheim 1956) Methodologically speaking, analysts must explain the social context (the socialization) in order to understand and explain the acquisition and stabilization of knowledge.

Whilst the Strong Programme focused its attention on scientific knowledge, the point is much stronger with regard to engineering knowledge. Successive research has shown how technology and engineering develops and generates knowledge in tandem with social developments and ideas.[47] Walter Vincenti (1990) argues that the

---

[47] See, e.g., Aitken 1985 who argues that the history of technology should be seen as a subject within the history of ideas. He explains the development of technological invention in the history

history of technology cannot be done properly with ignorance of social history and sociology. Barnes, Bloor, et al. would doubtless agree but saw a tougher and subtler challenge from the philosophy of science (and from anti-constructivist practising scientists, such as Alan Sokal, see Sokal & Bricmont 1998).

However, the focus on the sociology of scientific knowledge has, at least in its formative years, paid little attention to the specific challenges for a sociology of engineering knowledge. This is beginning to be addressed within the Strong Programme through Bloor's work with Vincenti (Vincenti & Bloor 2003) and his most recent book comparing British and German approaches to aeronautics during the First World War; (Bloor 2011) and Donald MacKenzie's analysis of the performative models used in economics (MacKenzie 2006, 2007) and military technology. (MacKenzie 1993)[48] I chose to concentrate on

---

of radio by examining the social ideas that emerged at the same time as new innovations were being developed. See also Laudan, R. 1984 who argues that changes in thought cause shifts in the kinds of solution adopted by technologists.

[48] There are many more examples which could have been mentioned from inside and outside SP. A notable example with a tangential approach to SP is Collins & Evans 2008, which examines the 'Hawk-Eye', a device used in sports such as tennis to decide difficult 'calls' and which, they

technology and social epistemology in this thesis in part because of a desire to contribute to this recent work and to approach the challenge of synthesizing SP and SE from a slightly different angle.

SP proposes a radical theory of knowledge, one that is highly controversial even in its own field. (Brown, J. R. 2002; Friedman 1998; Haddock 2004; Laudan, L. 1984; Okasha 2000; Ritsert 1990, p. 101-26; Sokal & Bricmont 1998) So why should we take its challenges to SE seriously? Haddock (2004) proposes a useful way to steer clear of arguments that seem to lose their force the more one considers them. Instead, Haddock argues, we can focus on those arguments from constructivism that do pose a threat to our current methods and understanding in epistemology. In Bloor's manifesto for SP (1991) he set out four principles which the sociology of knowledge ought to adhere to:

---

argue, may cause naïve spectators to 'overestimate the ability of technological devices to resolve disagreement among humans.

1. Causality: All explanations must be causal. That is, they must concern conditions which enable a particular belief to come about. In particular, Bloor is introducing a method for the sociology of scientific knowledge and so he is talking about causal explanations of why a particular scientific belief gains traction. For example, the belief that such and such a genome is responsible for a chosen hereditary medical condition or the belief that Mendelian inheritance correctly describes how characteristics are passed from parent organisms to their offspring.

2. Impartiality: Sociologists should be impartial with regard to the truth and falsity or the (ir)rationality of a belief. This is closely connected to principle 3, viz.

3. Symmetry: They should explain both true and false beliefs in terms of the same kinds of causes. For example, they ought not to explain successful theories by appeal to rationality, necessity, and correspondence to reality, whilst explaining superseded theories by appeal to sociological and ideological commitments on behalf of scientists and other interested parties. In fact, Bloor would like *all* theories to be explained sociologically.

4. Reflexivity: Remember that Bloor is attempting to put the sociology of science on the same strata as the natural sciences. Accordingly, the sociology of science ought to be critically examined according to the three previous principles with which it examines natural sciences. We should, therefore explain any sociological belief (including those of SP) causally, sociologically, and without preconception of its truth or falsity.

The third principle — the principle of symmetry — has received the most attention from advocates and detractors. However, Haddock argues that the symmetry principle does not, given a certain interpretation, come into conflict with much of contemporary epistemology. Indeed, there is much in these principles that can be made to work with a traditional epistemological account of knowledge. Here is the statement of the symmetry thesis from Barnes' and Bloor's 'Relativism, rationalism and the sociology of knowledge,'

> [A]ll beliefs are on a par with one another with respect to the causes of their credibility. It is not that all beliefs are equally true or equally false,

> but that regardless of truth or falsity the fact of their credibility is to be seen as equally problematic. The position we shall defend is that the incidence of all beliefs must be accounted for by finding the specific, local causes of this credibility. This means that regardless of whether the sociologist evaluates a belief as true or rational, or as false and irrational, he must search for the causes of its credibility. In all cases he will ask, for instance, if a belief is part of the routine technical competencies handed down from generation to generation. Is it enjoined by the authorities of the society? Is it transmitted by established institutions of socialisation or supported by accepted agencies of social control? (Barnes & Bloor 1982, p. 23)

Haddock considers what it means to ask for the 'causes' of a belief's credibility. One way of interpreting this request is to ask: 'What causes this proposition (that N holds true) to be worth holding true?' (Haddock 2004, p. 22) A second way of interpreting the request is to ask: 'What causes people to hold this proposition true?' In this passage, Barnes and Bloor seem to be asking both 'seek the causes of its credibility' and 'seek the causes for people judging it credible,' which is of course contradictory. Unless, that is, we assume that Barnes and Bloor are adopting what may be called a non-evaluative understanding of credibility.

To understand this, let us look at Bloor's definition of knowledge as: 'that system of beliefs that a community collectively accepts as knowledge. (Bloor 1991, p. 3)[49] Elsewhere, Barnes adds that there 'is no reason, of course, why philosophers should adopt this conception of knowledge for their epistemological purposes.' (Barnes 2001, p. 268) In other words, Barnes and Bloor are not making the claim that beliefs cannot be both true and justified. Rather, they are offering an instrumental or methodological definition of knowledge which they suggest is useful for sociological purposes.

Haddock suggests that we understand their use of the word 'credibility' similarly. A belief is credible, in the non-evaluative sense, if and only if a given community accepts it as credible. This definition need not conflict with an evaluative definition of credibility. Thus, Haddock shows that we need not think of the Strong Programme as posing a challenge to epistemological (evaluative) accounts of knowledge and

---

[49] A definition which Bloor acknowledges is 'rather different to that of either the layman or the philosopher' given that philosophers usually take knowledge to be, at the least, a true and justified belief. (Bloor 1991, p. 3)

credit but rather that the two are proposing theories of two different kinds of concept.

Whilst I agree with Haddock's approach taken here I do think that SP provides something that epistemologists have failed to provide: a methodology for analysing practical situations and for deciding whether an agent is right to attribute knowledge to another. The problem lies with a difference of opinion over the limits of epistemology. Most epistemologists would say that it is possible to define the conditions for knowledge attributions to be true. As Haddock says, it is often unclear whether SP argues that it is not possible to define truth conditions or whether it is just methodologically impractical. I believe that either statements such as those above are equivocations or that they are shorthand for saying that it is methodologically preferable to avoid any discussion of what is true or not beyond the consensus of the community.

Nevertheless, there remains a conflict between SP and most epistemologists and it is this conflict that I would like to focus on. As described in the opening chapter, the questions raised by Deepwater Horizon are questions that epistemologists are obliged to tackle. They are not merely empirical questions about who knew what and when but about when it is appropriate, correct or true to ascribe knowledge to an agent or a group, if it is ever appropriate, correct or true. However one wishes to expand on the terms appropriate, correct or true there remains a task to be undertaken here. A JTB-themed account of knowledge cannot answer whether the knowledge ascription is true or not as it leaves open the question of whether the belief is true and what is required for it to be justified. The sociological approach advocated by SP has the potential to provide an answer. The penalty is that we must, at least methodologically, assume knowledge attribution to be commensurate with knowledge itself.

3.1.4 The ontology of knowledge: Is knowledge natural, social or artificial?

The Strong Programme, and to a certain extent Kusch, claim that knowledge is socially constructed. In stating this, we should be careful to delineate three senses of epistemological social constructivism: (Schmitt 1994, p. 20-7)

1. The propositional *content* of knowledge is social. (e.g. Rouse 1987)
2. The *conditions* of knowledge are social. (e.g. Latour & Woolgar 1979; 1987; Kusch 2002; Rouse 1987)
3. Knowledge is typically *caused* by social, economic, and political interests. (e.g. Barnes 1977; Bloor 1991)

We can also frame the issue in terms of what kind of thing we consider knowledge to be. Different societies have different attitudes towards what kind of thing knowledge is and what other doxastic attitudes or metaphysical concepts it is cognate with. (see, e.g., Goldman 2004; Hallen 2004; Harding 1998; Hongladarom 2002a; Maffie 1995; Mulder 1996; Weinberg, Nichols & Stich 2001) This will be discussed in greater detail in §8.3.2 of this thesis. Hilary Kornblith – a preservationist, for Goldman, and a member of the complementary programme, for Kusch, makes the argument that knowledge is a natural kind. (Kornblith 2002, esp. p. 61 and 161) Kusch, on the other hand, claims that knowledge is a social institution. These are ontological claims about an epistemological concept. In other words, they are claims about what sort of thing knowledge is.

Let us examine Kusch's account in more detail before it is applied to a case study in Ch. 5. Although SP purports to be naturalistic it should not be confused with attempts such as Kornblith, Karl Popper, and W. V. O. Quine's to 'naturalize' knowledge. (Kornblith 2002; Popper 1972, p. 106-15; Quine 1969) Such naturalizing theories view knowledge as belonging to a group of natural kinds (such as water, H_2O) or artificial kinds (such as works of art or spiders' webs).

Kornblith, Quine, and possibly Popper, are committed to the view that knowledge continues to exist even if we, as a community, cease to refer to knowers, knowledge, and cease to make the relevant communal performative utterances. (Kusch 2002, p. 165) If we deny this, we should make explicit that knowledge is a social kind. (§4.2) SP and e-communitarianism view knowledge as a social kind (the set of things that include, e.g., marriage, money, and professional promotions). Kusch distinguishes these kinds by the nature of the references or 'communal performative utterances' made of them. (See Fig. 6)

Kind	**Examples of things**	**Ontological status if references cease**
Natural	H_2O, spiders, hydrocarbons	Continues to exist
Social	Marriage, money	Does not continue to exist
Artificial	Works of art, gyroscopes	Continues to exist as a natural thing only

Figure 6. Kinds of things and their ontological status.

3.1.5 Barnes on natural and social kinds

Before deciding whether we should accept Kusch's claim that knowledge is a social kind let us look at the history that led him to this point. In 1983, Barry Barnes published an influential paper making explicit SP's ontological commitments. (Barnes 1983) This paper relates the question as a problem of reference.[50] For Barnes, reference concerns the relationship between 'our speech and that which is spoken of.' (Barnes 1983, p. 524) Barnes considers as central to this problem an epistemological question: how can we know that the terms we use match the external world to which they purport to refer? Even the natural sciences seem to ignore any need for a systematic definition of what counts as successful reference for their own terms.

---

[50] When Kusch speaks of 'knowledge' marking a social status, he is also equating the ontological problem with a problem of reference. That is, when 'knowledge' *marks* a social status the term 'knowledge' or the practices that take the attributee to be a knower this term or practice refers to that social status. In such cases, the referencing practice constitutes the social status. §5.2 provides an example of referencing practices and communal performative utterances with the view to providing empirical methods for studying the production of knowledge although I come to different conclusions about the ontological status of that knowledge.

We tend to think of most things we encounter as belonging to kinds: things which may be treated or taken as being similar to other things of the same kind. Some kinds of things, for example minerals or plants, we think of as existing independently of ourselves. We can group elements in the periodic table or construct taxonomies of living beings, and yet in these cases we have come across a classificatory system, in some sense, already in place before we got there or, perhaps, that we were naturally, necessarily or inevitably inclined to so classify. If we stop grouping minerals and plants as kinds – if we stop referring to them as minerals and plants – then they will not pop out of existence. Barnes claims that, for these kinds of things, the reference becomes fixed because each person has a stored pattern which they compare to the object before them and attach a label accordingly.

For example, as I approach what I think looks like a flower I expect certain empirical characteristics. If what I encounter matches my internally stored pattern for 'flower' then I will (typically) identify it as a flower. If it does not, I will put it in the non-flower category. Either of these confirming or disconfirming references adds to the community's proper use of the term 'flower'. Proper use and proper concept application are key theoretical terms. We will later discuss the collectivist view on properness and its relation to proper function, a much-used term in the literature on technical artefacts. (See, e.g., Millikan, 1987 pp. 1-50; Franssen 2006) The proper application of 'flower' will likely differ across lay communities, horticulturalists, and so on, and is liable to change if enough agents change their behavior. The proper application of *'tulipa gesneriana'*, on the other hand, will change only incrementally and slowly, if at all. In general, all these terms are applied by comparing empirical characteristics to a pattern, model, norm or prototype.

The situation cannot be the same for social kind terms. There are no shared empirical characteristics for 'marriage', for example. No superficial pattern recognition could fix proper use. Marriage may be *symbolized* through wedding rings or another physical object but these are not necessary to constitute a marriage (if a person loses their wedding ring, for example, or burns their marriage certificate, the marriage remains intact or, if you like, we may still properly refer to the couple as 'married'). Rather, marriage is constituted (in part) by a series of performative utterances. (Austin 1975) These references could include the initial ceremony, the pronouncement of the minister, the idle chatter of neighbours and colleagues, official documents, and so on.

Social kinds are considerably more nebulous and varied in their properties and behaviour than natural kinds and the variety of social kind term references is a consequence of this. Consensus, concord and agreement between community references, together with the initial instantiation which 'primes' the whole process, forms this stereotypical account of social kind terms. (Barnes 1983, p. 529) If the stereotype for referencing natural kinds is broadly ostensive – identifying a thing by comparison with previous examples – then the stereotype for referencing social kinds is broadly stipulative – defining a thing by stating it as such.

Central to Barnes' account of reference is the 'performative model.' This model is a representation of what constitutes a natural or social kind. (Barnes has a dualist ontology to which Kusch added artificial kinds, see §4.1.4 and §4.2) In cases of natural kinds, references are constituted by a community's iterative use of the term to refer to objects displaying similar empirical qualities. For example, an individual encounters what looks to her like water: it is liquid, transparent, odourless, and tasteless; it is in a bottle labelled 'water'; and so on. She acquires a pattern for 'water' which she can compare to liquids she encounters in the future. If she uses this application in her interactions with others they may sanction, correct or rebuke her usage. It is through successive iterations such as these that the proper usage of the term 'water' comes to be stabilized within the community. What we are told of water and how others use the term affects our own behaviour.

Social kind terms, on the other hand, are not classified by their empirical characteristics. Instead, according to this model, they are entirely constituted by performative utterances which baptise and then either reward or sanction proper use of the name. For instance, take the social institution of money. 'Money' cannot be said to refer to any single physical structure or kind of thing. Rather, it can take many physical forms from metal coins to cheques and signatures to electronic data. In such cases proper use of 'money' depends on a system of use within the community.

Firstly, if we, as a community, stop calling the metal coins and cheques money and switch to an entirely electronic monetary system then the *nature* of money – its ontological status – changes. The rules of proper use change and the coins and paper simply no longer belong to any special kind of thing. Secondly, when an individual uses the term 'money' or 'married' he may look at certain physical markers in the same way that he may identify water or species of animal. However, no one set of markers is enough to fix proper use. For instance, physical markers that someone is married may include a diamond ring on the person's *digitus medicinalis* (the 'ring finger') or a signed marriage certificate. Nevertheless, the absence (or destruction) of any of these signs does not preclude proper use. Thus social kind terms are not dependent on empirical characteristics or perceptual performance.

3.1.6 Authority and trust

In some cases, agents may have been given some authority by the community to make the proper classifications. In others, the authority to classify will be dispersed across the community. By comparison, many philosophers (as well as scientists and lay people) argue that natural scientists enjoy an authority in classifying natural kinds. (Putnam 1975; Bird & Tobin 2008) They maintain that if, for instance, we want to know to what the term 'water' refers we ought to ask a scientist who explains that 'water' refers to a substance with the chemical formula, H_2O. Kusch and Barnes, along with John Dupré and Ian Hacking and others, have argued against such realist essentialism instead insisting that there is no single ontology appropriate to all communities ahistorically. (Barnes 1983; Dupré 1996; Hacking 19990, 1991a, 1991b, 1995; Kusch 1999) Rather, Dupré and Hacking favour a promiscuous realism where interests limit the class of relevant properties of an entity. On their account, no taxonomy is best for deriving generalizations regarding all properties and none is privileged by the world. Consequently, classification becomes a collective enterprise always indexed to a community and their interests. Kusch and Barnes can be said to be, in a certain sense, relativists about kinds in that what kinds there are depends on the interests and behaviour of those making the classifications. (See §4.1.3)

Some terms have technical uses. That is, the proper use of some terms differs across communities. Some terms, such as *digitus medicinalis* is typically only used within a specific community of medical experts. Others are common words which acquire technical uses within a discipline, such as 'information' within computer science. In fact, one might argue that almost all terms are technical terms in that all have proper uses indexed to particular communities. In many cases, some individuals or groups carry more weight in their uses of a term than others. For example, in a University classroom, a professor may have more authority than a student or group of students in deciding how to properly use a technical term.

The above description of natural and social kind terms is a stereotype. In reality, both are part-determined by authorized use. In considering a very simple example of social kind terms, Barnes asks us to think of one individual, *I*, holding unassailable authority to instantiate the proper use of a term 'as if by fiat'. In such a case, he pronounces something, *X*, an *S*, or any kind of a thing, *K*, an *S*, and by this very act, make *X* or *K* an *S*. (Barnes 1983, p. 526) In such a case, an 'S' just is whatever *I* calls an *S*. From the baptism on, whether *X* or *K* is an *S* depends solely on whether *X* or *K* is sufficiently similar to what *I* originally called an *S*.

Of course, there are few occasions where one individual would wield such power. In most communities cognitive authority is dispersed with a 'web of diverse and conflicting usages, with the simple self-reference loop of the first case replaced by spaghetti junction.' (Barnes 1983, p. 526) It requires careful and diligent analysis with particular attention to the specific properties of the community to understand any reference-talk that takes place. Of course, this kind of analysis applies to all social kinds. 'Authority' is also a social status and having authority is entirely dependent on being granted that authority by a community. Sometimes the authoritative status a person holds aligns with an epistemic responsibility. (§8.2.2 and §8.2.3) The person may be considered an authority or expert on a particular matter; they may be relied upon in testimony as a credible witness; or they may hold professional positions in which they have a responsibility to advise others on what is known.

3.2 E-contextualism vs. E-communitarianism

3.2.1 Historical background

In this section I will discuss a comparable theory to communitarianism with the distinction of being both individualist and truth-oriented. Despite their differences, there are elements of this group of theories, bracketed under the term 'epistemological contextualism' (henceforth, just contextualism), which can be preserved under an e-communitarian approach. Contextualism, put simply, is a theory about the truth conditions of knowledge attributions. It was initially proposed as a means of dealing with the problem of radical skepticism but has since been used to deal with other epistemological problems. (Cohen 1988, 1999, 1998, 1999; DeRose 1995) The essence of the idea is that attributions of knowledge are contingent upon a determining context, whether conversational or circumstantial.

Contextualism originated in the Relative Alternatives Theory (RAT) advocated by Fred Dretske during the 1970s and 1980s. (Dretske 1970, 1981) Dretske attempted to provide a more acceptable account of knowledge by replacing justification (that is, the requirement that an agent must have some grounds, or reason for their true belief) with the requirement that the agent is able to eliminate as possibilities all the 'relevant alternatives' to what she believes. In other words, possessing knowledge depends on an agent's capacity to rule out a certain range of alternatives which varies according to what kinds of alternatives are relevant. In subsequent formulations by Dretske and other sympathetic voices, the range of alternatives which are relevant has been held to be sensitive to a variety of factors. Unfortunately, it has proved as difficult to define 'relevancy' as it was to respond to Gettier. (Cohen 1994; Floridi 2008b, 2010a, p. 303)

3.2.2 DeRose and standards

Contextualism is similar to e-communitarianism and the reader might ask why it is not more appropriate to adopt a contextualist approach, thereby preserving a more orthodox individualist stance and not multiplying entities beyond necessity. Let us recall what contextualism asserts, its similarities to e-communitarianism, and whether it can be applied to sociotechnical knowledge. We will then consider what can be saved from the valuable contextualist insights in a communitarian-collectivist epistemology. I will focus on Keith DeRose's theory as it is among the most developed.

Contextualism, as DeRose describes it, is the position that the truth conditions of knowledge ascriptions (i.e. attribution and denial statements of the form '*S* knows that *p*' and '*S* does not know that *p*', etc.) vary according to the conversational context in which they are uttered. What varies, for DeRose, is the epistemic standards *S* must meet (or fail to meet) in order for such a statement to be true. In some contexts *S* must have a true belief that *p* and also be in a very strong epistemic position with respect to *p*. In others, *S* need only meet some lower epistemic standards. This is in direct contrast to Unger's invariantist account which argues that knowledge is invariant across all contexts. (Unger 1975)

Unger advocated a skeptical form of invariantism whereby a single set of standards governing knowledge ascriptions are consistent and highly demanding. This has the disconcerting consequence that much of the 'knowledge' that is routinely attributed to agents is, in fact, knowledge that they do not have. Unger later appeared to retract from an absolute commitment to invariantism, presenting a contextualist account in opposition to the invariantist account of his 'Ignorance'. (Unger 1975. See also Unger 1984) After proposing this account, he came to the relativist conclusion that there is no fact of the matter as to which of these views is correct.

DeRose's account is largely an attempt to resolve the apparent dilemma left unresolved by invariantist responses to the skeptical problem – what Michael Williams refers to as our 'biperspectivism', the intuition that skepticism is compelling under the conditions of philosophical reflection but never able to affect our everyday life wherein it is all but ignored. (Williams 1999; Williamson 2001) Contextualists argue that the truth-value of knowledge ascriptions vary according to a context and DeRose's bank cases are often used as an illustration. (DeRose 1992, p. 913) Here, there are two cases for us to consider. The reader should ask whether in each case DeRose knows that the bank will be open on Saturday morning.

> **Bank Case A**
>
> My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our pay cheques. But as we drive past the bank, we notice that the lines are very long, as they often are on Friday afternoons. Although we generally like to deposit our pay cheques as soon as possible, it is not especially important in this case

that they be deposited right away, so I suggest that we drive straight home and deposit our pay cheques on Saturday morning. My wife says, 'Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays.' I reply, 'No, I know it'll be open. I was just there two weeks ago on Saturday. It's open until noon.'

Bank Case B

My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our pay cheques on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important cheque. If our pay cheques are not deposited into our checking account before Monday morning, the important cheque we wrote will bounce, leaving us in a *very* bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, 'Banks do change their hours. Do you know the bank will be open tomorrow?' Remaining as confident as I was before that the bank will be open then, still, I reply, 'Well, no. I'd better go in and make sure.'

The message we should take home from this is that, although DeRose has precisely the same information, beliefs, and reasons for believing in both cases he is much less inclined to say that he knows in Bank case B and, intuitively, we are inclined to agree. What is the reason for this? 'The context of my utterances in the two cases make it easier for a knowledge attribution to be true in Case A than in Case B.' (DeRose 1992, p. 914) In other words, the truth-value of a knowledge attribution varies according to the context. Those who would agree with this interpretation, Unger calls 'contextualist' and those who disagree he calls 'invariantist'. (Unger 1984) As noted by Kent Bach, DeRose's case is interesting but problematic. (Bach 2005, p. 6. Fn. 6) The case asks us whether DeRose is justified in attributing knowledge to himself: 'Do I know that the bank will be open on Saturday morning?' Bach points out that the standards for evaluating a self-attribution of knowledge such as this can vary for reasons other than the subject's epistemic position, *qua* subject. First-person cases like this 'muddy the waters'. Instead, we should concentrate on third-person cases where the attributer of knowledge is not the subject, such as the airport case provided by Cohen, (Cohen 1999, p. 58)

> **Airport**
>
> Mary and John are at the L. A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds, 'Yes I know – it does stop in Chicago.' It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, 'How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute.' Mary and John agree that Smith doesn't really *know* that the plane will stop in Chicago. They decide to check with the airline agent.

Cohen suggests that examples of this form strongly suggest that ascriptions of knowledge vary according to context. (Cohen 1999, p. 59)

3.2.3 Semantic contextualism

One common accusation against contextualism is that its theories are not epistemological but linguistic, and so should be confined to the philosophy of language and not epistemology. However, this is to misread the contextualist project. Contextualism is concerned with knowledge-attributing and knowledge-denying statements, not knowledge *per se*. Its central tenet is that when we attribute knowledge to someone, what matters is the context in which we use the term 'knowledge'. Pritchard considers semantic contextualism - as put forward by such figures as Cohen, DeRose and Lewis to be the most dominant form of contextualist response to the skeptic. (Pritchard 2002, p. 215) Early contextualist theories did not explicitly express their account in conversational terms but, as DeRose recognises, the primary support for contextualism necessarily comes from 'how "knows" and its cognates are utilized in non-philosophical conversation.' (DeRose 2002, p. 168) In other words, it appears almost unavoidable that any contextualist theory will depend on conversational usage of the concept of knowledge.

For any comprehensive theory of knowledge, I would suggest, the meaning of 'knowledge' – and precisely whether its character and/or content are variable – must be articulated. Without such a definition we have an incomplete picture of how statements, assertions, and knowledge attributions are *used*. (Wittgenstein 1953, para. 43) However, this articulation will be revealed in our linguistic accounts of knowledge and it seems that the only accessible meaning that can be given to 'knowledge' is with respect to knowledge attributions. There is no other way to access 'knowledge' other than through knowledge attributions. One does not experience 'knowledge' in the same way that one experiences pain, visual stimuli, and immediate experience. I cannot report the sensation of knowing. Instead, our access to it comes through our articulation of the concept, our use of it in assertion and attribution.

DeRose acknowledges the objection that his conception of 'contextualism' is a theory of language not epistemology but does not see this as necessarily deleterious to the importance of the theory. (DeRose 1999b, p. 188) DeRose's response to the charge that his is a linguistic theory seems accurate but weak (although he hints that others could be made). He argues that, if it is a theory about language, then this does not detract from the significant importance it will have on any theory about knowledge. DeRose points out that those who work on the problem of free will and determinism, for example, should be very interested in the issue of what it means to call an action 'free'.

However, they would not reduce their theory merely to equivocations over linguistic meaning. DeRose seems to accept a lesser role for his theory as a guide for bigger epistemological questions but that seems to underplay the relevance of the meaning of terms to the study of concepts. He accepts the accusation that all his theory can tell us is how we use words and he may then use these conclusions to shed light on our theory of the real thing. But unless we think that knowledge is something that exists *out there* in the world independently of society and our human language, there does not seem to be reason to think there is much more to shed light upon. Instead, we can present an entirely linguistic account of knowledge and, at the same time, understand it as a functional concept within particular communities. Analytically – that is, when it comes to empirical investigation – one cannot properly detach “knowledge” from its conversational context.

For Jonathan Schaffer, the central dispute is over the interpretation of “knowledge,” specifically whether it is considered an absolute or contextually variable term. Schaffer considers the issue a straight dispute between the skeptic – who argues that ‘knowledge’ is an absolute term – and the contextualist – who argues that it is a relationally absolute term. (Schaffer 2004) The task for epistemologists is therefore to determine which of these provides the better hypothesis. Peter Unger defines “knowledge” as belonging to a category of ‘absolute concepts’. (Unger 1975) This includes terms such as ‘flatness’ or ‘emptiness’ which are ‘absolute’ in the sense that they do not admit of degree: an object is either flat or it is not. It does not make sense, for instance, to talk of a person knowing that the sky is blue *better* than someone else. Admittedly there is a sense in which we might say that someone knows better than someone else. We might say, for example, ‘Nobody knew the seas around the Cape of Good Hope better than Ferdinand Magellan.’ Here we are making an assertion about someone's general knowledge on a subject. We are stating that they have more, or more quality, knowledge on the subject than anyone else. This is a common use of 'knowledge' but it is sufficiently exceptional for us not to be distracted by it here. If pushed, we might say that this is a shorthand for expressing that someone has more knowledge about something or that they have knowledge which is of a higher quality.

Unger points out that nothing is ever absolutely flat or empty just as one never obtains absolute knowledge (i.e. certainty). A road may be called flat but it still contains bumps and ridges. A fridge may be called empty but it still contains water droplets and plastic trays. Dretske (1981a) and Lewis (1979) both argue that Unger's alleged 'absolute terms' should properly be understood as 'relationally absolute terms.' That is, that knowledge is absolute in relation to a range of contexts. Thus, absolute knowledge is obtainable in one context but there is no real context-independent absolute context. Dretske feels justified in his rejection of Unger's classification as 'to be empty is to be devoid of all relevant things.' (Dretske 1981b, p. 367)

In his attempt to answer the skeptic, DeRose constructs a linguistic account of knowledge where the term 'knowledge' is to be understood as an *indexical*. Broadly, this stipulates that the term has two aspects: a fixed character and a non-fixed content. (Kaplan 1989) For example, the word 'I' has a fixed character (or sense or meaning) in that in each mention of it, the word refers to the speaker. These are the rules or conventions that govern the use of the word. The content (or 'propositional component expressed by' or 'referent') of the word, however, is non-fixed in that it may fluctuate depending on who the speaker is. Similarly, according to DeRose, 'knows' has a fixed character – true belief – and a non-fixed content depending on the context expressed by the following: *S is able to track the truth of* p *through a set of relevant possible worlds*. By extension, the meaning of complex lexical terms, such as 'S knows that it is raining' are also context-sensitive. DeRose claims that the truth-conditions for all knowledge-attributing and knowledge-denying sentences can vary according to the context in which they are used. He maintains that what varies is the epistemic standards that S is required to meet (or fail to meet) in order for such a statement to be true. (DeRose 1999a)

3.2.4 Meaning and reference

We have seen in this chapter that e-communitarianism rejects all forms of globalism about justification; foundationalism; and the 'universal pretensions' of Rawlsian agents and thought experiments. In this sense, it can be described as a relativist programme. We have also seen that Kusch adopts a *semantic* approach to epistemology: he is concerned with the meaning of the terms 'knowledge', 'knows', 'knower', etc. In §3.1 'knows' was claimed to refer to the social status of an agent. In §3.2.3 and §3.2.4, we discussed a theory of knowledge attribution where 'knows' is considered an indexical. There are however, once again, subtle differences in which Kusch appears to distance himself from the orthodox SP account. In this section I argue that we should consider the use of 'knows' (and cognates) in language as relevant to our study of it.

First I would like to introduce a particular approach to meaning. Here is Alice Carroll's character Humpty Dumpty's take on meaning,

> 'I don't know what you mean by "glory",' Alice said.
> Humpty Dumpty smiled contemptuously. 'Of course you don't – till I tell you. I meant "there's a nice knock-down argument for you!"'
> 'But "glory" doesn't mean "a nice knock-down argument",' Alice objected.
> 'When I use a word,' Humpty Dumpty said, in rather a scornful tone, 'it means just what I choose it to mean – neither more nor less.'
> 'The question is,' said Alice, 'whether you can make words mean so many different things.'
> 'The question is,' said Humpty Dumpty, 'which is to be master – that's all.' (Carroll 2012, p. 123)

Most people, I am assuming, find Humpty Dumpty's pronouncements absurd. We might formulate a widely-held intuition about meaning thus,

> The meaning of a word does not depend on what I decide, at this moment, that it means.

If I decide now that 'lion' refers to, say, tigers I have not changed the meaning of 'lion' one iota. My use of the word would be viewed by others as incorrect: perhaps I am wrong about what 'lion' *means*, or perhaps they might think that I am trying to deceive them. But what makes such use wrong or deceptive? Extensional semantics is the view that a term's *extension* - the 'set of things of which it is true' (Barnes 1982, p. 31) - constitutes its meaning. On this account, there exists, in advance of any particular application of a term, an extension, which determines whether or not that application is correct. The extension of 'lion' is all the lions in the world that have been, are, and will ever be. Note that this view is compatible with social as well as non-social theories of meaning.

What it does commit one to is the view that, once the extension of a term is fixed - whether individually or collectively - then it follows from this that some applications will be correct and some will be incorrect. Once the meaning of 'lion', say, is fixed, this divides the entire infinite universe of entities that have been or may be encountered into two sets: entities that may be correctly referred to as lions and entities that may not be so referred; green entities and entities that are not green; solutes and substances that are not soluble. In other words, it divides the world into a set of entities of which the application of the term is true, and a set of which it is false. This is not to say that the meaning of words cannot change. Of course, they can and do. It is, however, to say that at any particular given time there is a determinate, fixed meaning.

Whilst being an intuitively plausible theory of meaning, it leads to some undesirable conclusions. Suppose that the meaning of 'knows' is determinate in this way. That is, at any time, *t*, there is a set of propositions of which it is true to say that *S* knows, and a set of propositions of which it is false to say that *S* knows. If we think of knowledge as a social kind this is problematic. It would mean that our social institution is structured in such a way that it neatly divides all propositions into those that are known and those that are not known. Considering that there are a near-infinite number of propositions which may be known, this seems wholly implausible. It would, at the very least, be extremely convenient if we had evolved an institution that could do this. It is much more likely that whether it is appropriate to say that someone knows something is not determined in advance but is a feature of their utterance within a particular context that comes into play when the assertion is made.

If we agree with the semantic contextualist approach we can apply this to our understanding of knowledge. Hence, knowledge is not a fixed, absolute term which, once its character is determined, cannot be altered. Instead it is a term with specific meanings that have evolved within particular communities. This explains why knowledge may have a meaning in engineering quite different from that used in everyday life and why it can have a different meaning in one community that is not compatible with another.

Consider the following case,

> **New Tribe Case**
>
> Kent is an anthropologist who has been staying with a newly discovered tribe in the Indian Ocean. He has been observing their behaviour and practices and thinks he may have seen the tribe performing religious rituals and ceremonies. He sends a message back to his research team that, 'I now know that the tribe practice religious rituals and ceremony.'

Let us suppose that the meaning of religious ritual and religious ceremony is something that is constructed by the community rather than from outside pressure of some sort. That is, religious rituals, for example, do not exist *out there* in the world prior to our performing and deciding that such and such behaviour would constitute a religious ritual. It does not seem at all fanciful in this case to suggest that Kent could not possibly know in advance of his observations whether the tribe practised religious ritual, not just because he had not yet observed any putative ritualistic behaviour but because the *meaning* of 'religious ritual' within that community is being produced by the behaviour itself and by others taking that behaviour to be a religious ritual. If we have indeterminacy in the case of religious ritual – a social kind – then we also have it in the case of knowledge attribution. If that does not sound plausible, consider another (real) case,

> **Odd Mammal Case**
>
> It is the 18th century and a team of British zoologists are presented with a new animal which they have never seen before. It looks extremely odd – something between a duck and an otter. It looked like a mammal with

> its beaver-like tail and otter-like feet, but it also lay eggs and had a bill like a duck. The zoologists are convinced the animal is an elaborate hoax created, perhaps, by Chinese taxidermists who are well-known for expertly stitching together animal parts.

It was later discovered, of course, that what we now know as a platypus is real and not a taxidermic chimera. Imagine the first scientist to christen the platypus. It does not make sense to think that he knew that the animal was a platypus. Nor does it seem justified for the second person, who was told by the first scientist that the animal is a platypus, to assert that they know that the animal is a platypus. After all, precisely what 'platypus' refers to is contentious and in the midst of formation. By comparison, one might think of how many ascriptions of the word 'quark' to particular phenomena created by hadrons had to be made before we could correctly say that anyone knew that such and such a phenomena was the effect of quarks.

Whether one thinks of knowledge as a social kind – such as a religious ceremony – or a natural kind – such as a platypus and a quark – it would seem that there is always an initial stage of indeterminacy around whether we are correct to use such and such a term to refer to such and such an object or event. If it becomes more stable and definite precisely what the term refers to then this is because of the hundreds and thousands of previous references that combine to make up proper use of the term. It would not be possible for a single individual utterance to overturn these thousands of uses and this is why the use has the appearance of determinacy.

The same applies to ascriptions of terms to objects as with ascriptions of knowledge to subjects. 'Knowledge', in ordinary settings such as those in the bank cases and the airport case, is a term whose proper use has been formed over countless utterances and so this gives the appearance that there should be determinate rules in place before any particular utterance that determine whether, say, DeRose and Kent are correct in their knowledge ascriptions. (§3.25) In fact, it is not determinate but highly-determinable. That is, one can say with a high degree of confidence what the conditions are given the large bank of previous instances contributing to proper use, even though this assertion cannot ever by true since for it to be true would require one to know in advance what the utterance was going to be and what effect that utterance would have on the proper use of the term.[51]

3.2.5 Conclusions on communitarian and contextualist epistemology

Gettier's paper purported to demonstrate that a widely-accepted definition of propositional knowledge is susceptible to a range of counter-examples. It had been assumed by many that propositional knowledge required three conditions to be met: a truth-condition, a belief-condition, and a justification-condition. Accordingly, knowledge requires that an agent possesses a belief in the proposition, reasons or adequate justification for believing the proposition, and the proposition must be true. Indeed, it does seem odd to characterize knowledge as being false, not believed, and lacking adequate

---

[51] Remember that some single utterances can dramatically change the proper use of a term. Some historians, for example, attribute Marco Polo's use of the word 'Madagascar' to refer to the African island rather than a part of the mainland to have changed the general meaning of the term. (Evans 1973; Kripke 1980)

justification. However, §3.1 discusses an account of knowledge that is not dependent on knowledge being true or, rather, given Haddock's refinements, does not require us to answer that question within the approach I am taking here.

In the first section of this chapter we discussed Martin Kusch's e-communitarianism, its historical roots in the Strong Programme in the sociology of knowledge and in political communitarianism. I presented a number of issues arising from the book that are relevant to my thesis: the ontology of knowledge or what kind of thing knowledge is; and the distinction between natural, social, and artificial kinds. The second section compared e-communitarianism with contextualism. I discussed the role of semantics in contextualism and its putative solutions to the skeptical problem.

We can conclude by saying that semantic theories are critical to epistemology and not a side-issue that might illuminate the 'true nature' of knowledge. Knowledge is, after all, not obviously a phenomenon that exists in the world but is a feature of language: we attribute it to others, claim it for ourselves, and argue over who has it and what it is that is known. Semantic theories are useful for this kind of phenomena as it provides a method for analyzing the use of terms and their place in social structures. In the next chapter I will look more closely at one kind of social structure or, to put it more precisely, a kind of sociotechnical structure. That is, structures and institutions which generate knowledge through the machinations of people and technical artefacts. The semantic approach is very useful for explaining not just how we can have various kinds of knowledge in different communities but also how the proper *function* of a tool can similarly evolve (see §4.2.7 and §5.2). In Ch. 4 I discuss what kinds of thing technical artefacts are and suggest that a communitarian approach will be the most fruitful in illuminating their various aspects.

Chapter 4

The collectivist and dual nature accounts of technical artefacts

4.1 The continuity of philosophy, science, and engineering

> There are not two (or five or seven) fundamental ontological categories, rather the act of categorization itself is always interest relative. (Searle 2002, p. 59)

> How does it happen that a properly endowed natural scientist comes to concern himself with epistemology? Is there no more valuable work in his specialty? I hear many of my colleagues saying, and I sense it from many more, that they feel this way. I cannot share this sentiment… Concepts that have proven useful in ordering things easily achieve such an authority over us that we forget their earthly origins and accept them as unalterable givens. Thus they come to be stamped as "necessities of thought," "a priori givens," etc. The path of scientific advance is often made impassable for a long time through such errors. For that reason, it is by no means an idle game if we become practiced in analyzing the long commonplace concepts and exhibiting those circumstances upon which their justification and usefulness depend, how they have grown up, individually, out of the givens of experience. (Einstein 1916, p. 101-102)

It is little more than a platitude to note the ubiquity of technology in everyday life.[52] What is perhaps even more platitudinous would be to note the ubiquity of technical artefacts in knowledge production. In fact, almost every area of inquiry that we use today – particularly inquiry that expands our common knowledge – is sociotechnical. That is, it involves the generation of collective technical knowledge. This is knowledge produced not with individual unaided human minds but with technically-enhanced human minds. And yet epistemology, for the very most part, studies only the unaided minds and treats the technical aspect as unimportant, unnoticed, or worse, uninteresting.

This is a chapter about the kinds of things that populate our world and how they affect our knowledge attributions. In particular, I argue that technical artefacts are a kind of thing in the same way that natural things such as water or vitamins are generally taken to be things that populate our world. In other words, technical artefacts are 'genuine

---

[52] Parts of this chapter are adapted from Kerr 2013.

substances'. (Baker 2008, p. 3-4). There are a number of theories which have attempted, with varying degrees of success, to establish technical artefacts as ontologically significant. In this chapter I focus on two competing theories in the metaphysics of technical artefacts: the Dual Nature of Technical Artefacts (DNA) and the Collectivist Account of Technical Artefacts (CAT). I argue on behalf of the latter being the most fruitful of the two theories although the two have several elements in common as well as differences.

Here is a fairly typical statement from the philosophical literature on natural kinds.

> Scientific disciplines divide the particulars they study into *kinds* and theorize about those kinds. To say that a kind is *natural* is to say that it corresponds to a grouping or ordering that does not depend on humans. We tend to assume that science is successful in revealing these kinds; it is a corollary of scientific realism that when all goes well the classifications and taxonomies employed by science correspond to the real kinds in nature. (Bird & Tobin 2008)

Different scientific disciplines may be accorded different privileges with respect to dictating what kinds are real. Chemistry may be taken to supply paradigmatic examples of natural kinds – chemical elements, compounds, and so on. Physics and astronomy may be second in line providing such kinds as electrons, quarks, and red dwarves. Taxonomies in biology were once taken to be paradigmatic but are now generally thought not to represent natural kinds.[53] Nevertheless, it would be inconceivable that, were science to discover a fault in its own classificatory systems, that this would not force a change in philosophical ontology. It is inconceivable, for example, that metaphysicians would continue their 'essentialist' talk about species following the advent of Darwin's theory of evolution. (Dupré 1993, p. 3; Okasha 2002, p. 191)

Similarly, quantum physics forced changes in certain metaphysical theories such as determinism (de Koninck 2008) and the philosophy of space and time has been brought through paradigm shifts with

---

[53] Seminal works by Kripke (1980) and Putnam (1975) refer to biological kinds as natural kinds although more recent work has challenged this. (Dupré 1993; Sober 1994, p. 163)

Newton, Mach, and then Einstein. (Whitrow 1980) The continuity of philosophy and science – the idea that philosophy ought to heed the results of empirical investigation in the natural sciences – has become orthodox, albeit a controversial orthodoxy, in contemporary analytic philosophy of science. Following this line of thought, philosophers of technology might say that if the natural sciences should inform the ontology of natural kinds, the engineering sciences should inform the ontology of artificial kinds.[54] That is to say, the privilege which scientists have enjoyed when it comes to classifying what Mario Bunge (1977) and David Bloor (2007) have metaphorically called the 'furniture of the world' – its basic constituents; what classificatory systems refer to – should also be accorded to engineers when it comes to classifying the actual *furniture* of the world – its made, or otherwise artificial, objects. If we accepts this thesis, then the purposes of this paper should be clear: to ascertain what kinds of entities engineers do posit with a view to including them in our ontological framework.

---

[54] The view that engineers are in a position of authority with respect to the nature and our knowledge of artificial kinds (or, alternatively, technical artefacts) has been proffered by Garbacz 2012; Layton (1974); and Vincenti (1990).

Despite what I have said above, that there is a distinction at all between natural and artificial things has recently been called into question. (Baker 2008; Pitt 2011; Preston 2008; Thomasson 2008) Others have argued for the stronger position that not only are artificial things not significantly distinct from natural things but they do not really exist at all. (Hoffman & Rosenkrantz 1997, p. 173) Lynne Rudder Baker cites digital organisms, robo-rats, bacterial batteries, and 'search-and-destroy' viruses as technological advances that blur the ontological distinction between natural and artificial kinds. (Baker 2008, p. 2)

However, as has been pointed out elsewhere, the blurred line between the natural and the artificial is not a consequence of modern technology, as these state-of-the-art examples might indicate, but has been present for at least as long as technology itself: as long as, for instance, the domestication of cattle or the agricultural cultivation of land and crops. (Preston 2008) Synthetically-created chemical compounds, genetically-modified foodstuffs, and many other things could cause us problems if we were to draw a fine line between natural and artificial things. What this suggests is that the distinction between the natural and the artificial has been problematic for a long time. Baker predicts that the distinction that is drawn between natural and artificial kinds on the basis of their mind-dependence or mind-independence (that is, in some sense, that the former do not depend on human minds whereas the latter invariably do) will continue to be eroded by future technological advances. (Baker 2004, p. 15; 2008, p. 5)

One normative constraint that is placed on putative distinctions which is often elided by philosophers interested in formal definitions is that they are useful. In other words, classifications are, on this understanding, interest relative. (Searle 2002, p. 59) We carve nature at various joints to suit our purposes of inquiry, knowledge expansion, and, in the case of engineering, to reshape our environment to better suit our needs. They are also profession-relative. That is, as Ian Hacking puts it, 'some kinds are essential to some crafts.' (Hacking 1991a, p. 123) Hacking's argument is that we should not expect one particular taxa, which has been formed historically to suit particular theoretical and experimental needs in one subject, to also be an excellent model for other fields of inquiry. None of this should be read as implying that the physics taxa or the biological taxa are the real taxa whilst the engineering taxa are not of, in Baker's terms, 'genuine substances'. (Baker 2008, p. 3-4) Rather, we place certain authority on expert classifiers and these professional taxonomies ought to be continuous with, or at least inform philosophical ontology. In the following sections I will describe a theory of classificatory practices that is particularly suited to empirically analyze reference practices to technical artefacts within the engineering professions and the results of which should be taken into account in discussions of the ontology of technical artefacts.

Instead of being considered worthy of attention in itself, it has often been assumed that technology aids various kinds of knowledge acquisition in a supporting role only. As we saw with the 'applied science' model of engineering in §2.3.2, technology and engineering are often not considered distinct fields of inquiry and areas of epistemological import in the way that science has been. That is, the technology itself has taken the role of the archetypal 'lab assistant' whose contribution is to bring supplies to the real 'brains' or hand the scientists what they need in order to produce knowledge. The job could not be done without the technician, but what she does is transparent and uninteresting. I wish to demonstrate that this picture is as inaccurate for the role that technical artefacts play in our epistemic practices as it is for the role that technicians play in scientific practices. Artefacts are not merely conduits for knowledge to simply pass through but shape knowledge and it is only through looking at what the artefacts are doing and what they are that we can understand the knowledge that is produced with them.

4.1.1 The ontology of technical artefacts

There are several ways in which we might build an ontology of technical artefacts. We could list examples of what is meant by ‘technical artefact’ – scissors, computers, microscopes, pianos. One might suggest that these are instruments or tools for accomplishing some goal; that they are means to ends. Perhaps this is a good definition of a technical artefact. The problem is that it appears to open the net too widely as a great deal of things can be said to be means to ends or tools for accomplishing goals. My preparatory school entrance exam was a means to entering a University programme. However, if I were to call it a 'tool' I would be speaking figuratively. If I walk into the woods and use a tree branch to overturn a large stone we could say that my using the branch was a means to an end: overturning the stone. However, again, it sounds odd or metaphorical to call the branch a 'tool' or a 'technical artefact'. It is, at most, being used *as* a tool but is not itself a tool. Tree branches are natural objects that exist independently of us whilst tools are, typically, objects which we create.

Crows, which remarkably are better at using natural objects as tools than most primates, will bend wire into hooks to access food and Woodpecker finches will use sticks to wheedle out a grub. (Bluff et al. 2010; Shumaker et al. 2011) However it sounds odd to describe the length of wire used by the crow as a 'tool' in anything but a metaphorical sense. (See §4.2.6) There are a number of intuitive reasons why these examples do not constitute tool use. At first glance, in order for something to be a tool it must have been designed for a specific purpose by a designer with a community of users in mind who use the tool in an appropriate way. It must also have some physical presence and be structured in such and such a way within limits. However, these reasons are based on intuitions and a more systematic, empirical approach is needed.

4.1.2 The dual nature of technical artefacts

The Dual Nature Thesis of Technical Artefacts (DNA) is a thesis about the epistemology and ontology of technical artefacts. This research comprises one of the few systematic, analytic approaches to technical artefacts. (See, e.g., Kroes & Meijers 2002, 2006; Houkes & Meijers 2006; and Houkes & Vermaas 2004) DNA stipulates that technical artefacts have two components: a structural, physical component and an intentional, social component. The guiding intuition is that we routinely describe the world in terms of two basic conceptualizations,

- a world of physical objects interacting through causal connections, which we describe by means of the concepts of the natural sciences (e.g. position, velocity, spatial dimension, and other physical magnitudes).
- a world of agents, which we describe by means of intentional concepts (e.g. thoughts, wishes, belief, will, and effort).

There is, of course, a long-standing problem in philosophy over how these two descriptions relate to each other which DNA does not directly aim to resolve.[55] The research project started as a joint programme between three Dutch technical universities (Delft University of Technology, Eindhoven University of Technology, and the University of Twente), and three U.S. universities (University of

---

[55] It parallels, for instance, the mind-body problem in the philosophy of mind.

Buffalo, Massachusetts Institute of Technology, and Virginia Tech) at the beginning of the third millennium. It has since produced a large amount of theoretical and empirical research into engineering practice and described an analytic approach to the philosophy of technology and engineering which is sorely lacking. This work includes treatises on the nature of design (Houkes, Vermaas, Dorst & de Vries 2002); function and planning in engineering (Vermaas & Houkes 2002, 2003, and 2004; Houkes 2006; Franssen 2006; Vermaas 2006); research into engineering knowledge and practice; and an ontological framework which respects the place and significance of technical artefacts (see the special issue of *Studies in the History and Philosophy of Science* 2006 vol. 37, especially Houkes & Meijers 2006 and L. R. Baker 2006).

Following from this research, we might propose the following five-fold necessary criteria for a technical artefact:

a) Physical: has a structural (physical) component;

b) Social: has an intentional (social) component;

c) Design: is purposefully designed;

d) Function: is designed for a correct, proper or appropriate function; and

e) Normative: its use is constrained by normative factors

In summary, technical artefacts are physical objects that are described by physical concepts (e.g. the transistor has a length of 10mm) and by intentional concepts (e.g. the transistor has the function of amplifying electronic signals). These conceptualizations are indispensable for fully describing technical artefacts. If an artefact is described by only physical concepts, it is usually unclear what functions it has, and if it is described only functionally, it is usually unclear what physical properties it has. The artefact must also be designed for a particular function depending on normative considerations (e.g. the electric kettle is designed for boiling water for domestic use and has an overload cut-out switch for reasons of safety).

It should be clear from this that technical artefacts are not simply physical objects. They are not known purely through 'reading off' their empirical characteristics. (See §3.1.5) One cannot know, for example, what a turbine engine is simply by listening to a description of its physical characteristics.[56] This would suggest that artefacts are a kind of social thing. (Barnes 1983) However, neither are artefacts known purely through social, normative, functional properties since the realization of their function depends crucially on their physical structure (unlike, for example, money, which serves the same function in a number of different physical forms).

4.1.3 DNA and technical knowledge

Houkes outlines the DNA approach to knowledge of technical artefacts (technical knowledge). (Houkes 2006) He claims that

---

[56] Although it is not possible to work this out simply on the basis of a description of physical characteristics, one may be able to work out the social, functional and design aspect of an artefact through an understanding of particular cultures, histories, needs, requirements, predilections, the function of other similar artefacts, and so on. This is done frequently in archaeology and reverse engineering as discussed in Vaesen and van Amerongen 2008.

knowledge of the function of an artefact is not provided exclusively by beliefs about its physical characteristics but is primarily provided by what he calls 'know-how related to its use.' (Houkes 2006, p. 102) Second, knowledge of the proper function of an artefact is 'primarily based on testimony and a social division of labour with respect to rational artefact use.' (Houkes 2006, p. 102) Knowledge of artefact functions includes, for example, knowing that a tin opener is for opening tins; a computer technician's knowledge of how to recover damaged data from a hard-disk drive; the proposition 'S knows which wire is earth,' and so on. He states that knowledge of artefact functions may be expressed thus,

> [T]he best evidence for the claim that an artefact can be used to *f*, or has the possible function to *f*, is that you have use know-how, that is, that you know how to use it to *f* successfully. (Houkes 2006, p. 103)

Use know-how has three distinctive components:

i. Not only observation of physical characteristics, including dispositions.

ii. 'Procedural knowledge' that a sequence of actions leads to the realization of a goal.

iii. 'Operational knowledge' (skills) needed to take these actions.

In summary, use know-how is knowledge of functions and other non-physical properties as well as the relevant physical characteristics; knowledge that executing a use plan for the artefact in a certain range of circumstances and employing a set of skills and auxiliary items realizes the goal state of the plan; and/or operational knowledge or the requisite skill to perform the actions in the use plan.[57] Procedural and operational knowledge are often complementary. Typically, use know-how is based on both and only in special cases on either one alone. Houkes wishes to make a general epistemic claim about how we acquire knowledge but for our purposes here it is enough to note that in almost all cases users acquire use know-how from sources beyond

---

[57] That is, a 'use-plan.' (Houkes & Vermaas, 2004) 'That a bicycle can be used for cycling may to some extent be explained by appealing to the mechanical principles on which its use plan is based, but actually showing someone how to ride a bicycle may be a necessity rather than a merely convenient alternative.' (Houkes 2006, p. 105) It has also been shown that Galileo was able to prevent others from replicating his experiments by withholding information about how to carry out the experiment and that replication of Boyle's experiments with the air-pump depended on the replicant directly witnessing the experiment and not in his published accounts. (Biagioli 2000; and Shapin & Schaffer 1985, p. 225-82)

their personal experience and set of beliefs. It would be an enormous amount of personal work to consistently do otherwise. Hence, social knowledge is the norm when it comes to use know-how.

Much of the work done by DNA researchers involves describing the apparent tensions between our knowledge of the empirical characteristics of an object and our social use know-how. However, these problems are created from conceptualizing this dual nature as analytically separate. This misrepresents how we come to acquire use know-how. We do not acquire use know-how, for instance, from first encountering the physical structure and then learning or deducing the artefact's function or use plan. Rather, knowledge of both are acquired co-temporally. Compare this with the way that a child will not first learn that an arm is a long cylindrical part of their body and then learn how to use it to pick up objects but will learn both the physical features of the arm and how it is used together. In the next section I introduce what may be seen as a tweak to the general DNA approach which I've called the Collectivist Account of Technical Artefacts (CAT). This account aims also to deliver epistemological and ontological respect for technical artefacts but does not require the dualism of the DNA.

We have discussed how DNA conceives of the world as given by two kinds of description: one using physical concepts and one using social or personal concepts. Most readily, this duality calls to mind the Cartesian dualism of mind and body. This correlation between how we conceive of ourselves and how we conceive of our technology should not be surprising. Recent work in cognitive science indicates that complex tool use is accomplished in macaque monkey brains by treating the tool as simply an extension of existing limbs and capacities. (Rizzolatti et al. 2002; Johnson-Frey 2004) Thus, when we use a tool, it may be that our brains learn that the tool is a kind of arm or hand which is connected to us and belongs to our bodies. In the study, monkeys were given three tasks involving grasping for food: one where a set of pliers is used, one where a set of reverse pliers is used—pliers which require the user to 'open' their hand grasp in order to grasp with the pliers, and one where only the hand is used). In all three tasks, the same sequence of neuron-firings was observed within the monkey brains. This suggests that the way we conceptualize technical artefacts is often very similar to the way we conceptualize our own bodies. If, further, we acknowledge what is now common knowledge amongst cognitive scientists – that minds are embodied things – we might conclude that the project to separate physical and intentional descriptions is somewhat moot. A better account would include both these descriptions within an extended system of tool-plus-body-plus-mind. This is what CAT aims to do.

4.1.4 The collectivist account of technical artefacts

Before outlining CAT, we must answer a recent objection to the category of 'artificial' as marking a distinctive kind of thing. As stated earlier in this chapter, recent work on the ontology of technical artefacts has staked a claim for the 'shrinking difference between artefacts and natural objects.' (Baker 2008; Preston 2008) Criticisms of long-standing distinctions are commonplace in philosophy and often have argumentative support. However, in this case, and specifically in engineering, I wish to argue that the distinction does have heuristic value and it is this value that is often omitted from philosophical arguments which find counter-examples to the distinction and, on this basis, conclude that the distinction is invalid. In the following sections I outline the collectivist account; its distinction between natural, social, and artificial kinds; its merits and demerits. In Ch. 5 I show that this distinction is very useful within petroleum engineering. Whether or not there are counter-examples in some domains is not a knock-down argument to whether or not there are contexts in which the distinction is valid, properly used, and heuristically valuable. The case study demonstrates how – beyond the intuitive way in which we might divide natural, social, and artificial classes – the uses of 'tool' are quite specific to particular kinds of objects. This warrants classifying them as particular kinds of things which cannot be reduced to either natural or social categories or subsumed beneath one un-blurred but less useful metaphysics.

Instead I draw on the collectivist epistemologies and ontologies of Martin Kusch and the Strong Programme outlined in Ch. 3 to describe the dynamic and inseparable manner in which physical and social components combine. Instead of two separate components which inscrutably and inexplicably hold together, there is a single description which constructs or co-produces both components co-temporally and co-dependently. The project to somehow separately analyse the two components is perhaps the one objection I have to the DNA approach. Once this is jettisoned we will have a more robust and flexible model of technical artefacts. CAT also places greater emphasis on the social, institutional and historical aspects of technical artefacts and, with this in mind, §5.1.1, §5.1.2, and §7.2.1 provide historical accounts of why tools developed in the form they did and how this changed knowledge production.

In §2.2.2 I said that this chapter will argue on behalf of social constructivism about ontology. DNA is not a constructivist approach. In the following I will outline a method for analysing ontological categories as socially-constructed whilst preserving the valuable insights of DNA. First of all, we ought to consider what kinds of objects we can have knowledge of. It is now widely accepted that there is a conventional aspect to how we classify objects. The 'arbitrary' decisions we make to classify this element as a gas and this as a metal is partly 'interest-relative' or 'constructed' (Searle, 2002; Smith & Welty 2001) When thinking about what kinds of objects exist in the world it is easy to suppose that these objects belong necessarily to a particular group: hydrogen is necessarily a kind of gas as opposed to a kind of metal; fruit flies necessarily belong to a group called 'invertebrates' as opposed to frogs who necessarily are 'vertebrates'. That is, features of structures in nature determine as a matter of necessity that hydrogen is a gas. However, over time our supposedly necessary classifications diversify and change. New and more precise distinctions are made and occasionally the whole system of classification changes. This variation is a feature not just of our lay classifications of things (such as fences, games, and fruit) but of the most precise scientific classifications. It is often hard to imagine that scientific classifications are subject to change as we change our interpretation of phenomena but all could theoretically change if the paradigm in which they are postulated changes.[58]

---

[58] See Kuhn 1962. All scientific classifications, if they have been around for long enough,

Classifications can only be made sense of within a structure of classification and we call this ontology. By studying rare or extinct languages, researchers have illustrated the remarkable diversity in ontologies. One example is the Australian aboriginal language, Dyirbal. Whenever a Dyribal speaker uses a noun in a sentence, the noun must be preceded by a variant of one of four words: bayi, balan, balam, bala. These four grammatical categories organize the Dyirbal speaker's world into four ontological groups:

- Bayi: men, kangaroos, opossums, bats, most snakes, most fishes, some birds, most insects, the moon, storms, rainbows, boomerangs, some spears, etc.
- Balan: women, anything connected with water or fire, bandicoots, dogs, platypus, echidna, some snakes, some fishes, most birds, fire-

---

change and the definition of what set of things they describe mutates just as the definition of what set of things are described by 'game' changes. The gravitational force, for instance, has changed its boundaries and referents significantly from the early formulations of Galileo, through Robert Hooke, Isaac Newton, to Albert Einstein. The point is even more stark if one includes the physics and astronomy of Aristotle, Plato, Aristarchos (perhaps the first to claim that the Earth orbits the Sun), Ptolemy, and so on.

flies, scorpions, crickets, the stars, shields, some spears, some trees, etc.

- Balam: all edible fruit and the plants that bear them, tubers, ferns, honey, cigarettes, wine, cake, etc.
- Bala: parts of the body, meat, bees, wind, yam-sticks, some spears, most trees, grass, mud, stones, noises, language, etc. (Lakoff 1987, p. 92)

As Ian Hacking notes, it is difficult for us to know how to interpret the 'etc.'s in passages such as these, although it would be fair to assume that the categories seem wholly natural to native speakers. (Lange 2007, p. 382) This is because we belong to a culture which has its own, different, ontologies which likewise feel wholly natural to us. The more we are used to seeing a category used in a definite way, the more it seems like a feature of reality and not merely one of our classificatory system. Familiarity breeds convention, one might say. Hacking concludes that no single taxonomy is privileged by the world (i.e. the world does not offer or provide a substantial normative constraint on our classificatory systems), but different kinds are natural for different purposes. (Hacking 1991a, 1991b) With different purposes in mind, we must take an interest in predicting different facts; recognize different contrast classes in explanations; and so on. Different classifications are best for induction and explanation.

That no classification is best for all purposes is no mere contingent matter. There is no such thing as *all* properties, Hacking says, so there is no such thing as a taxonomy's being best for arriving at reliable generalizations regarding all properties. Interests must limit the relevant properties before there can be a best taxonomy. What distinguishes natural kinds is that they characterize kinds about which numerous generalizations can be made. (Boyd 1991, Hacking 1991a) This is not an anti-realist position though it may be, in Dupré's terms, a "promiscuous realism". (Dupré 1996) Further, you do not need an entire grammatical structure of language to create a distinctive ontology. As §5.2 demonstrates, professions such as petroleum engineering have created their own taxonomy of technical artefacts that may or may not be commensurable with taxonomies in the natural sciences.

One problem with the promiscuous and relativist accounts is the possibility of an explosion in the sheer volume of kinds. If ontology is always indexed to a community or discipline, then communities and disciplines may potentially 'invent' kinds without much constraint from non-human factors. This would be an undesirable result for philosophers interested in identifying the basic constituents of the world. One way to rein in this unconstrained promiscuity is through ascribing authority to certain communities or agents. Another is to abandon the search altogether. CAT suggests that analysis can be conducted at the level of the community and the influence of non-collective factors will be felt in the practices of that community. Issues such as authority and consensus are relevant but issues such as whether classifications truly reflect the world are not. This is not to suggest that there can be no cross-pollination from one community to another (for example, from the classifications of natural scientists to laypersons) but it is to suggest that there is no ontology that adequately fits all communities. CAT stipulates that these ontologies are created as 'working oppositions' which are heuristically valuable. (§4.2.5) Whether or not they reflect the 'joints' of nature is moot.

4.1.5 Reference and referencing practices

CAT analyses technical artefacts from a sociophilosophical perspective and takes its inspiration and methodology from theoretical and empirical research in SP. The application of SP’s theory of reference and meaning finitism to technical artefacts is drawn from Martin Kusch's tripartite ontology of natural, social, and artificial kinds and Pablo Schyfter’s critique of DNA. (Kusch 1999, Schyfter 2009) It focuses on practices of reference or what I will call ‘reference-talk’. This is based on a methodological commitment that studying ontology is best done by studying reference-talk which contributes to classification. As Bloor describes, SP should not be concerned with the ‘furniture of the world’ conceived of as basic constituents entirely separate from sociological practices such as science and engineering. (Bloor 2007)

Absolutist ontologies – those that attempt to define the objects of the world separate from sociology – are methodologically unsound in that they can show no systematic method from progressing from intuitive hypothesis to the claim that the hypothesis is true absolutely. CAT purports, therefore, to offer a better methodology for approaching ontological questions; one which approaches the classifications of communities as having as stable and 'foundational' a justification as any ontology can ever have. Larry Bucciarelli and Peter Kroes have recently demonstrated how careful study of the rhetorical output of engineering can provide fruitful insights to those interested in classifications of entities both inside and outside of the profession. (Bucciarelli 2009; Kroes 2010) Recall Barnes' distinction between natural kind and social kind terms. (§3.1.5) The distinction Kusch makes between three kind terms (natural, social and artificial) will be discussed in detail in §4.2. Kusch suggests that these three kind terms are distinguishable by virtue of the type of references made in conjunction with their use.

To provide a brief summary, consider what would happen to the classification 'rock' if we were to cease to use it to refer to objects as we currently do. The answer is that 'rocks' (although not so called) will continue to exist however we choose to refer to them. Their basic properties are not changed by a change in our referencing practices. These are natural kind terms. On the other hand, a social kind term such as 'marriage' depends entirely on our continuing to refer to it in a particular way. If we all ceased to treat the couple next door as 'married' then they would quite simply cease to be married. Likewise, we can destroy certain physical markers of marriage – wedding rings, signed documents – and yet the marriage endures. Artificial kind terms are of particular interest to Kusch and CAT researchers because they appear to contain qualities of both these types of kind term. In one sense, a pair of scissors does not vanish if we stop talking about it (the object will still exist in space-time). But in another sense it no longer seems to have many of the basic properties it once had and needs to have in order to be a pair of scissors. Something is lost when we stop treating an artefact as an artefact and CAT researchers believe that we can gain insight into what is lost by examining the community of users that take it to be an artefact.

The distinction is designed to be intuitive but empirically accessible. Its intuitiveness is exemplified if one considers how frequently the N/S/A distinction appears in innocuous and philosophically light situations: Most of us are familiar with the children's game 'rock, paper, scissors'. Here natural, social, and artificial objects are taken to be opposed to each other (opposition here means that one object 'defeats' the other in the context of the game). Another example is how we tend to think of the oil and gas industry. Here, natural kinds are extracted from the Earth using artificial kinds to satisfy social kinds (needs). In our case we compare acts of taking-to-be with acts of reference. Just how one studies reference-talk empirically (and even what an act of reference is), however, has been less explicit and so §5.2 offers one example in the form of an analysis of the textual output of a particular community and the kinds of references made to a key class of objects in this industry: tools.

There is, of course, more to reference than writing or talking about something (there is gesture, behaviour, signs, signals, and so on). Much of the practical or tacit knowledge involved in using and physically manipulating a tool, for example, is not encoded even in detailed instruction manuals but learnt, as it were, on-the-job. Further, a collectivist account of reference is not so much focused on the references of individuals but with how these compare with others within the community and with how incompatibilities are ironed out and agreement over usage is settled. Remember that correct application consists in matching a pattern to a previously established prototype.

The practice of explicitly classifying objects is familiar to many engineers. Many companies today employ 'engineering ontologists' who are given the specific task of constructing an, often quite blunt, system for organizing and ranking objects such as tools within the company. These are often arranged according to a type and given serial numbers or names which reflect a designed hierarchy. Similarly many company manuals and reference documents will contain ordered tables of, for instance, what each tool is called and what it does. It is evident that such practices have great practical value for the efficient workings of a company. The collectivist account hopes to develop engineering ontologies based on sociophilosophical work rather than those created in systems engineering. (See Garbacz 2012; Vermaas & Garbacz 2009)

4.2 The (N/S/A) kind terms distinction

4.2.1 Artificial kinds

Let us now consider the N/S/A kind terms distinction developed by Kusch and Barnes in more detail. Barnes distinguishes natural and social kind terms by virtue of the ways in which they are referenced. (Barnes 1983) Reference here is conceived quite broadly as the talk or behaviour of agents taking an entity to be a certain kind of thing. Subsequently, Martin Kusch added a third kind term – artificial – whose nature lies between the natural and social, combining as it does types of reference from both kinds. It would appear that artificial objects are a peculiar type of object in the world of petroleum engineering and that there is something specific to them which cannot be reduced to either natural or social kind terms. The arguments for this are presented in §5.2.

Recall what was said about Barnes' natural and social kind terms and the non-necessity (contingency) of classification. (§3.1.5, §3.1.7, §4.1.4,

§4.1.5) Kusch (1997; 1999) presented an augmented version of Barnes' account that added artificial kind terms to the natural and social. The reason for doing this is that artificial kind terms (terms that refer mostly to made things) are not wholly identified by their empirical characteristics but nor can they be reduced to an entirely social system of fiat and consensus. The referent of a social kind term, for Kusch, is entirely constituted by the references themselves. Talk (and behaviour which takes the referent, R, *as an R)* actually creates the referent and so, in the final analysis, the referent is the talk (and behaviour) itself.

> The reference… is, as it were, 'exhausted' by the self-reference. (Kusch 1999, p. 245)

Consequently, if the references cease, the referent drops out of the ontological framework. Natural kind terms, on the other hand, have an 'alter-reference'. Here the term refers away from the talk towards something in the physical world; things that exist independently of the talk. There are still central patterns, models, norms, prototypes, exemplars, and so on, for what constitutes, say, an igneous rock, and so even natural kind terms have a self-referential collectivist component. However, the pattern, etc. is not the thing itself.

Artificial kind terms are odd cases in this framework. 'Scissors', for example, does not just refer to talk (and behaviour) about scissors. However, nor can they be properly identified solely through reading off their empirical characteristics. The term has an alter-reference to the physical composition and structure of the object. For example, the reference to a gyroscope is partly a reference to a system or community in which the gyroscope functions and partly a reference to the physical materials that make it up. However, if all we had were these physical materials with no function or instrumental component to back it up, then we would not, strictly speaking, have gyroscopes (even though we would have physical, natural objects which look identical).

One might call to mind the aftermath of the knowledge 'cull' imagined by Alasdair MacIntyre in *After virtue.* (MacIntyre 1984) MacIntyre imagines a post-apocalyptic world in which all scientific artefacts and texts are completely or partly destroyed and all that is left of a previous scientific culture (ours) is fragments and remnants. None of the survivors' descendants attempting to piece together the forgotten knowledge understands what the things are *for*. It is this element that is missing in an account that describes artefacts without reference to a community of users. In order to be more than just a collection of useless materials, a technical artefact requires a performative social institution referring to the artefact *qua* artefact. Part of the task of collectivist studies is to identify just what 'referring to an artefact *qua* artefact' is: does it involve design, purpose, functionality, instrumentality, normativity? What do the referencing practices of the community reveal about these questions?

4.2.2 A shrinking/blurry distinction?

As I pointed out at the start of this chapter, the blurred line between the natural and the artificial is not a consequence of modern technology. However, it also seems a persistent distinction across a wide timeframe. There are countless references to this in Western culture from Hippocrates' distinction between physics and techné, to Plato and Aristotle's definition of the mimetic arts, to modern science and Bacon's image of the natural philosopher's 'rape of nature'. In the same way, the natural/social divide pervades much of our contemporary thinking but it is easy to spot tensions: human psychology is a rich example with its tug-of-war between so-called natural instincts and social conditioning. In petroleum engineering we can picture a tripartite world of oil rigs (artificial) extracting hydrocarbons (natural) from the Earth (natural) by means of sociotechnical systems (artificial and social) to satisfy social needs and demands.

On the other hand, it is not difficult to think of plausible counter-examples to a sharp dividing line between natural, social and artificial kind terms: a rock being used as a crude hammering device; the artificial canals of Amsterdam; engineered biological organisms and genetically-modified crops. And we have recently witnessed the creation of a "synthetic" cell.[59] It is clear that, if there is a divide at all, it is a blurred one. Blurriness is only worrying, however, if one rejects the picture of classification given above. Since these classifications are constructed by our communities, it would be more surprising if they did divide the world perfectly into two hemispheres. Blurriness is a likely condition of all non-logical oppositions. With this in mind, it is important to remember that, although such tensions and imperfections may exist, the distinctions and boundaries we invent serve or have served a purpose. Recall Hacking's platitude that we make concepts for a reason (at least in the first place). If there is nothing new about these problematic cases there is less reason to suppose that a once useful distinction has fallen into desuetude. Most entities, states, events and processes do not pose a problem for the N/S/A framework and, although some challenges may prove fatal to the framework in the future, they retain their current heuristic value.

---

[59] The New York Times, 21st May 2010, p. 17.

4.2.3 Re-engineering the distinction as a working opposition

We should not read the N/S/A distinction as a report on facts about how the world is but on facts about how we have often framed how we see the world. This is not something that cannot be overturned but is what I have termed a 'working opposition'. It is important to remember that these distinctions are not the way they are because they *have* to be but because we have built them (probably unwittingly) to be incompatible. The world does not come 'pre-carved', as Plato had it. At least, it requires some effort on our part to decide where to draw the boundaries and the question of whether or not it is pre-carved serves no immediate practical use. The reason why the British discoverers of the platypus had such difficulty believing that such an animal could exist was because they could not make it fit into their current taxonomy for classifying animals. This taxonomy had worked superbly for every animal encountered up to that point and so, they reasoned, this alien creature must surely be an artificial invention. Just as we decide which kinds of things qualify as mammals and which qualify as fish, every other object in our world is (metaphorically) placed into a particular box with its own criteria for entry. Invariably, these boxes are multiple with an object going in either one box or another depending on the reasons for classification. Plants are living organisms which produce energy through photosynthesis *whereas* animals are living organisms which are motile; elements are molecular compounds whereas elementary particles are sub-atomic particles that do not have a substructure.[60]

The way in which I propose we look at such distinctions is not as logical-binary oppositions but as *working oppositions*. The activity of making distinctions is not meant to preclude counter-examples. It is meant to be a useful way to split up referents that are spoken about in very different terms. We should not expect all references to be reducible to entirely natural, social or artificial descriptions. Working oppositions do not comprise an airtight ontology but a pragmatic solution to picking objects out of an increasingly populous world – as we should expect from the product of communities working over time

---

[60] These particular distinctions I have invented based on my own limited knowledge and, no doubt, there are more and less precise distinctions that could be made.

to reiterate and evolve classifications. Nevertheless, I hope to demonstrate that there is uniformity in artificial kind term use which is markedly opposed to typical natural and social kind references.

4.2.4 Virtues and vices of the CAT of the N/S/A distinction

The Kusch-Barnes approach is perhaps more esoteric than the methods used by engineering companies but has a number of virtues. For one thing, it explains very well why and how referents change over time and classifications become unworkable. If Marco Polo used the term 'Madagascar' to refer to the island and not, as was normal then, part of the African mainland, then he would likely have been ignored or corrected by his contemporaries.[61] This single out-of-place use would not change the proper application of the term. However, over time, others did begin to adopt this application and so the reference changed. One vice that may be lodged against the account is that, by

---

[61] The example of Madagascar as a case of reference-switching is taken from Evans. (Evans 1973) Evans' suggestion is that when Marco Polo used the term to refer to the island we now call Madagascar, he intended to use it as it had been used by others, not introduce a new usage. The lesson I illustrate here is that once Polo did use it to refer to the island, and this stuck, the proper usage subsequently switches. (See also Kripke 1982)

focusing on referencing practices, it is overly superficial whereas traditional metaphysicians have looked for something more robust and fundamental.

Perhaps Kusch or Barnes would respond that this search is wrong-headed but, even if it is not, the accusation of superficiality is, I submit, unjustified. For instance, a traditional realist metaphysician might state, as Hilary Putnam did, that the ‘essence’ of water is H_2O. That is, the true referent is not based on identifying a string of empirical characteristics but it is the molecular composition of water, H_2O. However, there seems to be no principled distinction here. We might ask how the molecular composition is identified. It is through various tests all of which are empirical investigations. There is no principled method for distinguishing between ‘is wet’, ‘flows in lakes’, or ‘is tasteless, odourless and transparent’ and passing a scientific test for having the molecular composition H_2O or, as is more common, for identifying water in a scientific setting, testing neutral under a pH test or boiling at 100°C at sea level. In each of these cases, proper application of the reference is determined by comparing a set of empirical characteristics and behaviours with previous exemplars. At this level of abstraction the analysis of lay and scientific reference can be the same.

4.2.5 Robinson Crusoe's 'tool'

The common counter-example to collectivist approaches such as CAT is a variation on a theme which I will call 'Robinson Crusoe' examples. The relevant variety here is 'Robinson Crusoe's tool'. Consider the following remark from Amie Thomasson,

> [I]t makes perfect sense to suppose that a solitary human could create a knife, but not a government or money. Thus artefacts don't seem to be essentially *social* objects at all. (Thomasson 2007, p. 52. Original emphasis.)

Suppose Thomasson's solitary human was trying to create the first ever knife. Is it still intuitive to say that this person can do that but not create the first ever government or money? It seems that Thomasson's statement only makes sense if one first supposes that there exists a community of tool-users (knife-users, etc.) and that this solitary human is creating a tool of a kind which is recognized and established within the community; a kind of which it makes sense to say that it is an artefact only within the system that defines what counts as an artefact.

> **Crusoe's tool**
>
> Robinson Crusoe is a man stranded on a desert island. He was stranded there as a young child and all he has ever known are natural objects – trees, plants, fruit, sand, and so on. Nevertheless, he uses what he has around him to fashion a rudimentary axe, knife and various other useful things. He uses them to accomplish various tasks, and makes other similar objects, gradually improving them each time.

Has Crusoe created tools? Some people will say, 'Yes, of course he has created tools. As you say, he has created axes and knives. What are these if they are not tools?' Note that I said he has created a 'rudimentary axe.' Alternatively, I might say that he has created a proto-axe or an axe-like object. This would minimize any preconception about whether or not he has created a tool. Now let us consider what it would mean to say that he has created a tool. Would we also have to say that you can create a tool even if no one else recognizes it as a tool? Perhaps we can say that they *would* recognize it were they to see it. But both DNA and CAT hold that one cannot recognize a tool just by looking at it. One must also know the context in which it is a tool. That seems to be missing here since only Crusoe knows of the object's existence. Beth Preston has argued that tools are typically identifiable regardless of our witnessing anyone using the tool,

> [U]nlike naturally occurring objects, manufactured tools are designed for a specific use, and normally have a standardized form. More importantly, they are identifiable as tools—indeed often as specific kinds of tools apart from *any* actual occasions of use. (Preston 1998, p. 6. Original emphasis.)

Whilst this may often be the case it seems anthropologically-flawed. The only reason why we would ever be able to identify what a tool is for would be on the basis of either a) witnessing actual use, or b) prior knowledge as to the proper function of similar tools. (Vaesen and

Amerongen 2008; and Dennett 1990) Consider a variation on Crusoe's tool,

> **Crusoe and the coconut**
>
> Robinson Crusoe is lonely on the island and decides to create a friend. He draws a face onto a coconut, places it on a stick, and begins to talk to it. Since Crusoe has been on the island since he was a young child he has never learnt to speak English and so he makes up his own noises to talk to (or, rather, at, for it cannot respond) the coconut.

The scenario is the same as Crusoe's tool although in this case the question is over whether Crusoe has created a language rather than a tool. I think in this case we are unlikely to say that Crusoe has created a language, especially those of us who are familiar with Wittgenstein's later work. As Wittgenstein said in his *Philosophical Investigations*,

> The words of this [private] language are to refer to what can be known only to the speaker; to his immediate, private, sensations. So another cannot understand the language. (Wittgenstein 1953, para. 243)

Just as Crusoe cannot create a language by himself (since a language requires communication) he cannot create a tool by himself (since tools require proper functions). Crusoe may construct something that looks like a tool, functions as a tool, and is intended to fulfil a specific function. Crusoe cases are rare but they highlight a specific issue that is underplayed in Thomasson's and Preston's account. We require a background of tool use in order to know what the proper function of a particular kind of tool is and, once established, that background knowledge can be extremely powerful. However, none of that implies that we can separate our analysis of what the proper function of a tool is from the actual occasions of use. The individual causal history of the artefact and the non-individual normal use history, to borrow Preston's phraseology, do not 'come apart' but are both vital to any understanding of that particular tool and its function within a community. This point is more evident when we consider not simple or basic tools such as hammers and screwdrivers but with complex technological artefacts such as aeroplane engines or hadron colliders since the background knowledge of tools of that kind may be quite limited and weak (since there are fewer and their occurrence is limited to specialized zones). Crusoe's makeshift tool cannot be properly referred to as a tool because it exists in isolation from a community, and is consequently without proper function.

Consider the idea of tool-appropriation. Appropriation occurs when an object (either purely physical or artificial) is used *as* a tool. For example, suppose that I am unfamiliar with surgeon's scalpels and, on finding one in a colleague's desk drawer, I use it as a letter-opener. In so doing, it seems we should say that I have not instantly changed the scalpel into a letter-opener but am using a scalpel *as* a letter-opener. In other terms, I am taking the scalpel to be a letter-opener; referring to it as such. But this is just one instance among a whole classificatory system and, what is more, I have no particular authority in the community when it comes to decided what surgeon's scalpels are for.

Consider instead a more mundane example where you walk into the forest and construct a rudimentary lever using a branch. The collectivist account analyses this as using a branch *as* a lever, not as creating a lever (since it is an individual action). This is an extension of the principle discussed earlier that technical artefacts must be purposefully designed for a particular function or functions. It is not enough to pick up a palm leaf and use it to fan yourself for that palm leaf to stop being a natural kind and become a technical artefact (i.e. an artificial kind). For this to be the case, social structures of proper use would have to be implausibly flexible and quick to pop in and out of existence. Secondly, the function (of fanning) is incidental to its physical structure; it was not purposefully designed for this function.

Further, referring back again to the five criteria outlined in §4.1.2, there are no normative constraints on the palm leaf. No one is saying that the palm leaf *ought* to be used to fan. There might be occasions where someone would do this but the whole point of appropriation is that there is not this complex normative structure already in place. It is a fleeting moment of innovation. If, on the other hand, there exists a community who routinely appropriates palm leaves for the function of fanning; who amends their physical characteristics (or, at least, executes some design decision upon them); and who collectively says 'You should use palm leaves for fanning'; then the cut palm leaves may change their ontological status from natural to artificial kind; from simple leaf to tool. There is undoubtedly a grey area concerning this transition and whether it is right to call the object natural or artificial is a contingent, not an absolute, matter. Any analysis of the transition has to pay close attention to the reference-talk of the community, its design choices, and the development and use of the putative tool.

4.2.6 Tool use in animals

We may gain insight into the nature of tools if we do not confine our analysis to how humans use and describe tools and ask instead whether animals create and use tools. The question was first given serious attention in the 1960s when behaviour was observed in some non-human primates that resembled tool use: using a stick to sweep food, using bones to crack nuts, washing fruit and vegetables, and so on. (Hall 1963; Kortlandt & Kooij 1963; Vevers & Weiner 1963) Today much more has been learned about the capacity of some non-human animals to use objects in their environment to accomplish tasks.[62] Is this mere appropriation or is it tool use? Much of the debate has been around whether or not animals propagate culture – pass on knowledge to others – or whether these uses are spontaneous. Given the account of tool use that CAT provides researchers are right to focus on this aspect. Consider the following cases of animals using objects in a way which has been described as tool use,

- Boxer crabs pick up poisonous anemones to wield off attackers; (Karplus et al. 1998)
- Woodpecker finches and New Caledonian crows use sticks to winkle grubs from rotting wood. In fact, the latter travel with a tool kit of

---

[62] See, for example, the huge collection of research collected in McFarland & Bösser 1993, and Shumaker et al. 2011.

implements proven to be useful; (Bluff et al. 2010; Shumaker et al. 2011; Tebbich & Bshary 2004)

- Ground squirrels kick sand into the faces of venomous snakes to deter attacks; (Shumaker et al. 2011, p. ix)
- Macaque monkeys use pliers to grasp food; (Rizzolatti et al. 2002)
- Mice set out markers to serve as a reference point in exploration; (Stopka & Macdonald 2003)
- Gombe chimpanzees use captured meat as a political tool, withholding it from rivals and doling it out to allies. (Nishida et al. 1991)
- Beavers build dams so that a pool is created suitable for den construction;
- Egyptian vultures open eggs by dropping them onto stones from a height. (Shumaker et al. 2011, p. 37)

I have selected these examples more or less at random from various studies on animal behaviour.[63] How many of these examples would we

---

[63] Beck 1980 provides a large list of studies from the time that purported to show tool use in animals. Shumaker et al. 2011 list 53 thought-provoking examples. (Shumaker et al. 2011, p. 1-2)

say are examples of tool use? The sources given all describe or suggest the possibility that each behaviour is a form of tool use but is this philosophically sound? Animal researchers have developed many definitions of what they mean by tool use,

> [T]he use of physical objects other than the animals' own body or appendages as a means to extend the physical influence realized by the animal. (Jones & Kamil 1973)

> [T]he use of an external object as a functional extension of mouth or beak, hand or claw, in the attainment of an immediate goal. (Van Lawick-Goodall 1970, p. 195)

> [T]he active external manipulation of a moveable or structurally modified inanimate environmental object, not internally manufactured for this use, which, when oriented effectively, alters more efficiently the form, position, or condition of another object, another organism, or the user itself. (Pierce 1986, p. 96)

> [T]he external employment of an unattached environmental object to alter more efficiently the form, position, or condition of another object, another organism, or the user itself, when the user holds or carries the tool during or just prior to use and is responsible for the proper and effective orientation of the tool. (Beck 1980, p. 10)

Jones and Kamil's definition is very broad. If we were to accept a definition such as this we would also have to countenance a chimpanzee climbing a tree to reach fruit as tool use (i.e. using the 'tree tool' to reach fruit). The tree is not part of the animal's body or appendages (it is not clear whether Jones and Kamil wish to make a distinction here) but it extends the physical influence realized by the animal. If we want the definition to be this broad then it is so meaningless that it becomes a matter of animal behaviour in general; behaviour that inevitably interacts with non-animal objects at some point. Van Lawick-Goodall's definition may also be liable to admit too much a lot depends on what is meant by a 'functional extension'. Would it include, for example, the last example on our list? McFarland and Bösser argue that, under van Lawick-Goodall's definition, tool use would not include the vulture dropping an egg onto stones – since the stone is not an extension of the vulture's body. (McFarland & Bösser 1993, p. 195) However, Egyptian vultures are also known to carry stones into the air and drop them onto nests or pick stones up in their beaks and throw them at eggs. These behaviours would count as tool use because the stone is a functional extension of the vulture's body. The research in Rizzolatti et al. (2002) suggests also that extensions of existing appendages and capacities are closely linked to brain function when using putative tools.

Pierce and Beck's definition comes closest to the CAT definition although they do exclude the design element. If we want a definition of tool use that applies to both human and non-human animals – and I think there are sound reasons for wanting so – then we ought to include some element of design choice. This could involve altering the physical or social structure of the object in some form or it could involve simply choosing the particular object as *good* for a particular job. This is the normative condition described in §4.1.2. For example, CAT may be able to accommodate the behaviour of the Indonesian octopuses which collect and adapt shells for use as hiding places as a form of tool use. (Finn, Tergenza, & Norman 2009) The behaviour was reported as being the first evidence of tool use in an invertebrate animal.[64] It was discovered that the octopus *Amphioctopus marginatus* would select halved coconut shells from the sea floor, empty them of detritus, carry them for up to 20m, and assemble two shells together in a clam shape to make a defensive fortress.

---

[64] Henderson, M. 2009. Indonesia's veined octopus 'stilt walks' to collect coconut shells.' *The Times*, 15th December 2009.

It is in fact well-known that octopuses of various genuses will use external objects as a form of shelter, this behaviour is much more complex and worthy of the name "tool use". The way the animal selects appropriate materials for the job, carries them to an appropriate location, and then assembles them, is well-accommodated under our definition of tool design, manufacture, and use. Animal behaviour researchers often define tool use as being the activity of keeping and maintaining an object for future use. (Finn, Tregenza, & Norman 2009; Mulcahy & Call 2006) I would suggest that when we talk about design as a central component of tool manufacture and use, the kind of behaviour displayed by *A. marginatus* is, if the reader will excuse the pun, on the margins of a definition of tool use.

CAT can accommodate tool use amongst non-human animals in the same way as tool use amongst humans. Whether DNA can so accommodate tool use amongst non-human animals is not clear.[65] The

---

[65] Although Ted Cavanagh does say the following: Design describes intellectual activity that differs across disciplines. This chapter argues for differentiation into engineering, architecture, or other types of design before any general conceptualization. Studies about the 'dual nature of

examples in the literature are all drawn from human tools and focus on elements of design and normative constraints that could only plausibly be said to apply to humans. (See §4.1.2, §4.1.3) There are sound reasons for wanting an account that accommodates both groups, besides parsimony and the lack of a clear historical or evolutionary separation between the two groups. We may, for example, want to ask whether the animal behaviour studies referenced above have any implications for human behaviour. Certainly they should, but we cannot draw many conclusions without an understanding of tool use that is pan-specific.

4.2.7 Conclusions on the collectivist and dual nature accounts of technical artefacts

This chapter focused on accounts of technical artefacts, tools, and tool use. I outlined an existing theory (DNA) and a fledgling one (CAT).

---

artefacts' concern engineering design. The transferability of philosophical concepts from these studies to other fields of design is questionable. (Cavanagh 2008, p. 301)

From this discussion we can create the following conditions for being a tool,

1. Must use both structural and social terms in order to be comprehensively described;
2. Is designed (or chosen) to meet a need;
3. Has a proper function;
4. Has proper conditions of use;
5. Is not appropriated as part of a single unorthodox instance of use but is part of a systematic structure of tool use of that kind which determines 1-4.

Tool use is simply the use of one of these objects in an attempt to attain the end specified in the design plan. In Ch. 5 we will see whether these conditions fit those of the reference-talk in petroleum engineering by conducting an empirical study. In Ch. 7 and 8 I discuss whether this definition leaves open the possibility of tool use in artificial systems, extended epistemic systems, and groups.

Chapter 5

Case study observations and a textual analysis of reference to 'tools' in directional drilling and production logging analysis

5.1 Introduction: The case study disciplines

> Like trying to measure the force of a fly lifting its leg. (Donald H. van Steenwyk, engineer) [66]

> Nothing is known. (Engineer BK during interview on production logging analysis)

This chapter provides an empirical basis for CAT. The results of a fieldwork study indicate that technical knowledge is distinguishable from other types of knowledge and that the definition of tool use given in §4.2.7 is, partly, justified. In the first section I provide a non-

---

[66] Van Steenwyk died during the writing of my dissertation. He was the owner and CEO of the company with which I conducted my fieldwork research. I will always be grateful for that opportunity to gain an insight into the workings of an engineering office. This quote captures the idea of precision but also uncertainty and verisimilitude involved in gyroscopic surveying.

technical description of two ‘epistemic cultures’ of petroleum engineering: borehole surveying (including directional drilling and wellbore navigation); and well logging analysis (specifically production logging analysis). I present a history of these fields and the evolution of their technical artefacts. Finally, I provide an illustration of what is done on an oil rig and how engineers problem-solve. This resonates with the outline of the engineering method given in §2.3.3 and §2.3.4. The second section provides a statistical analysis of textbooks in the two epistemic cultures with a view to analysing reference-talk therein. This analysis gives a possible definition of ‘tool’ and what is taken-to-be a tool within the epistemic culture. The final, third, section diagnoses what this definition means for technical knowledge.

Wellbore navigation, directional drilling, and borehole surveying are some of the specialist fields involved in drilling the well towards a predetermined target. Production logging provides techniques and heuristics for optimizing production by measuring various properties of fluids in the well. In short, these two disciplines use tools to measure properties of a well (be it the orientation and inclination of the well or the properties of fluids such as velocity, density, temperature, and so on). The artificial kinds such as the pipe or an oilrig are of interest probably only to those specifically tasked with manufacturing, maintaining and using them (unless, of course, they malfunction). As with most social kinds, they are not the primary focus of most textual output from the disciplines of wellbore navigation and production logging. The artificial kinds most engineers have to have functional knowledge of in these disciplines are tools. More specifically they are measuring instruments, sensors, and so on, which provide data for drilling engineers and analysts. This can then be made into meaningful information.

5.1.1 A brief history of borehole surveying technology

Today, the task of borehole surveying is to direct a wellbore towards a specific subterranean target area (typically an area of about $50m^2$ or a volume of $50m^3$) which is predicted to contain a hydrocarbon-bearing reservoir. Most of the time the drilling team will not aim to penetrate the target directly ('as the crow flies', so to speak) but will have to avoid other wells and obstacles ('fish') as well as consider the rate of build acceptable to the pipe (it is not possible to drill a perfectly vertical well). This is done through modern sciences of geology and engineering, however its epistemic roots extend back much further. According to the Greek historian Herodotus, humans have been exploiting the Earth's oil reserves for over four thousand years. Herodotus wrote that Babylonians were using bitumen (a substance found in crude petroleum) to build walls and towers. (Macaulay 1890, Bk. 1 Para. 179) The earliest oil wells were drilled in China in 347 AD.[67] These wells were drilled using bamboo poles and plunged to an impressive depth of about 800ft (240m). In the following millennium petroleum was used for numerous purposes—to evaporate brine and produce salt; to produce kerosene for lighting and heating or simply to burn; for paving roads; and for military weapon.

However, up until the middle of the 19th century, it was assumed that crude oil was something which seeped in small amounts from coal. A Christian missionary named Laurent-Marie-Joseph Imbert (now Saint

---

[67] From Geo-Help Inc. 2011. History of the world petroleum industry (key dates). http://www.geohelp.net/world.html.

Imbert) had returned from China with news of their drilling techniques but the information could not be put to use. It is often stated in history books that the first commercial oil well was drilled by Edwin Drake in 1859 but in fact the first appeared in Poland six years previous and a second in Romania in 1857. (Howarth 1997, p. 12) It was the Russian empire, rather than North America which led world crude oil production at the end of the 19th century. (Arkiner & Aldis 2004)

Surveying actually dates back to Euclidean geometry and some archaeologists suggest that ancient Egyptians used surveying methods. (Paulson 2005) Herodotus recounts how surveying techniques were used to settle boundaries for tax rate purposes. (Macaulay 1890, Bk. II Para. 109) As an industry, however, borehole surveying did not develop much until the late 19th century. The first tool developed for this purpose was a simple wax cylinder used for the first time in the African diamond mines in 1882. The engineer's experience of drilling in those days would, of course, have been much different from the contemporary industry.

And yet, the epistemological questions central to the enterprise—by how many degrees is the hole oriented; what is the depth of the hole; where is the target area in relation to the current drill position; what is the likelihood of penetrating the target area; what is the current inclination of the pipe?—have not much changed. To measure inclination in 1882, engineers would lower the cylinder to the bottom of the hole and allow the wax to solidify at a particular inclination.[68] Once the cylinder was removed from the hole, engineers could simply read off the degree of inclination. We will come to appreciate, however, if it was not already apparent, that every new technology brings limitations along with the new possibilities and knowledge it opens up. For one thing, the wax would not solidify properly at lower depths (due to the high temperature) and this would affect its structure as it was removed from the hole to be examined. Acid was put in the wax solution which would etch a line around the circumference of the cylinder but even this was a sticking plaster over a deeper wound: a

---

[68] Nothing more than the simple laws of gravity are exploited, although the philosopher in me wants to ask if, like Descartes, they wondered whether the wax is the same as it comes out as it was when it went in. Indeed, in many ways, it is not: It is now a communicator, a potential source of knowledge.

more reliable, more mechanical, more precise instrument for measuring hole orientation was required.

5.1.2 Modern surveying techniques

The first answer to this was the development of clockwork instruments in 1899. These were specifically designed for measuring the orientation of core samples from the wellbore. Then, in 1910, devices began to appear which used the exciting developments in photography at the time. One device would record traces on a smoked glass plate. A later device took photographs of compass plates. Despite the advantages of these devices over the very crude early methods, they lacked mechanical precision. Many tools were subsequently developed in the 1920s and 1930s to plug this gap. Some of these tools are still in use today, albeit in an updated form. The TOTCO tool is the more basic of these. It consists of a pendulum with a stylus attached and a paper disc which the stylus penetrates each time a survey is taken. This paper is marked with concentric rings indicating hole inclinations. Obviously, the paper must be then removed from the hole before it can be read by engineers.

The Teledrift tool uses mud-pulse telemetry to communicate its reports. Like the TOTCO tool, it uses a pendulum but this is connected to a plunger. Inclination is derived from the number of restrictions through which the plunger passes. Telemetric systems of communication will be of interest when I later discuss the temporality of data communication as well as the textuality of these systems. The former involves a distinction between continuous versus single-survey reports; and near-instantaneous communication versus data stored in the unit's memory and recovered at a later time. The later involves an argument I will make, following the work of several contemporary philosophers of technology, that this communication may best be interpreted as a form of writing, marking, inscription, or narrative. This is linked to the point about temporality which makes this more precisely the writing of a history.

These three tools – the wax cylinder, TOTCO and Teledrift – are called 'inclination only tools'. That is, they measure only inclination and not, for example, orientation. They also produce physical representations of the data they aim to represent (although it is debatable that the pressure pulses produced by mud-pulse telemetry should be so included).

The final significant development in surveying techniques is the use of gyroscopes. Gyroscopes were first developed as a pedagogic aid and demonstration tool at the start of the 19th century. By the 1860s, electric motors allowed the gyroscope to spin for long periods of time and could be used in navigation on ships and, later, on aircraft. Since then they have been used in racing cars, motorbikes, gyrocompasses, as stabilisers on mono-trains and ships, robotics, computer mice, and much more besides. In the last few years we have seen gyroscopes extend into the world of consumer electronics and are used in devices such as virtual headsets, the iPhone and the Nintendo Wii. Modern gyroscopes in directional drilling are expensive and highly-sensitive instruments, far removed from their pedagogical ancestors.[69] This gyro spins at 4200rpm and senses a single second of movement of rotation

---

[69] Estimates are between $250,000 and $850,000. There are currently as few as 600 tools in the world.

of the Earth on its axis. This works by exploiting the reactive torque caused by a spinning Earth. The gyro tool will be held stationary (i.e. at a set angle) in a well and the rotation of the Earth is sensed as a torque. The magnitude of this torque is a measure of the alignment of the gyro spin axis with True North.[70] Quite how precise the instrument is depends on what we might call social factors—cost, mobilization, contracts, etc.—and this has a knock-on effect on what knowledge can be gained from the instrument.

5.1.3 An overview of well (production) logging analysis

A second source of knowledge I will examine is well logging analysis. Essentially, this is the science and engineering used in ascertaining the nature and behaviour of fluids at particular depths or 'zones' of the well. One reason for doing this is to acquire information about the surrounding rock composition. This need not be done using PLT and

---

[70] In fact, only the horizontal component of the Earth's spin vector is relevant. This dependency on latitude means that different corroborating data applies in different regions and the accuracy of north-seeking gyros diminishes near the North and South poles.

until recent history it was predominantly accomplished by examining 'cuttings', rock removed from the drilled hole. Geological tests of cuttings would be an example from PL of what I am calling direct instrumental information. Other reasons are to analyse the productivity or injectivity or different zones, diagnose problem wells, and to monitor the results of a stimulation (e.g. hydraulic fracturing treatments) or completion (bringing a wellbore into production once drilling operations have been concluded).

Production logging is the term used for well logging done during production or injection. Production logging analysis is the analysis of data acquired from configurations of tools sent down the wellbore to monitor and record the properties of fluids. This includes spinners (fans which indicate fluid velocity), temperature sensors, and sensors for resistance, density, and so on. The records from these instruments configured within the tool are recorded in a well log such as that pictured in fig. 7.

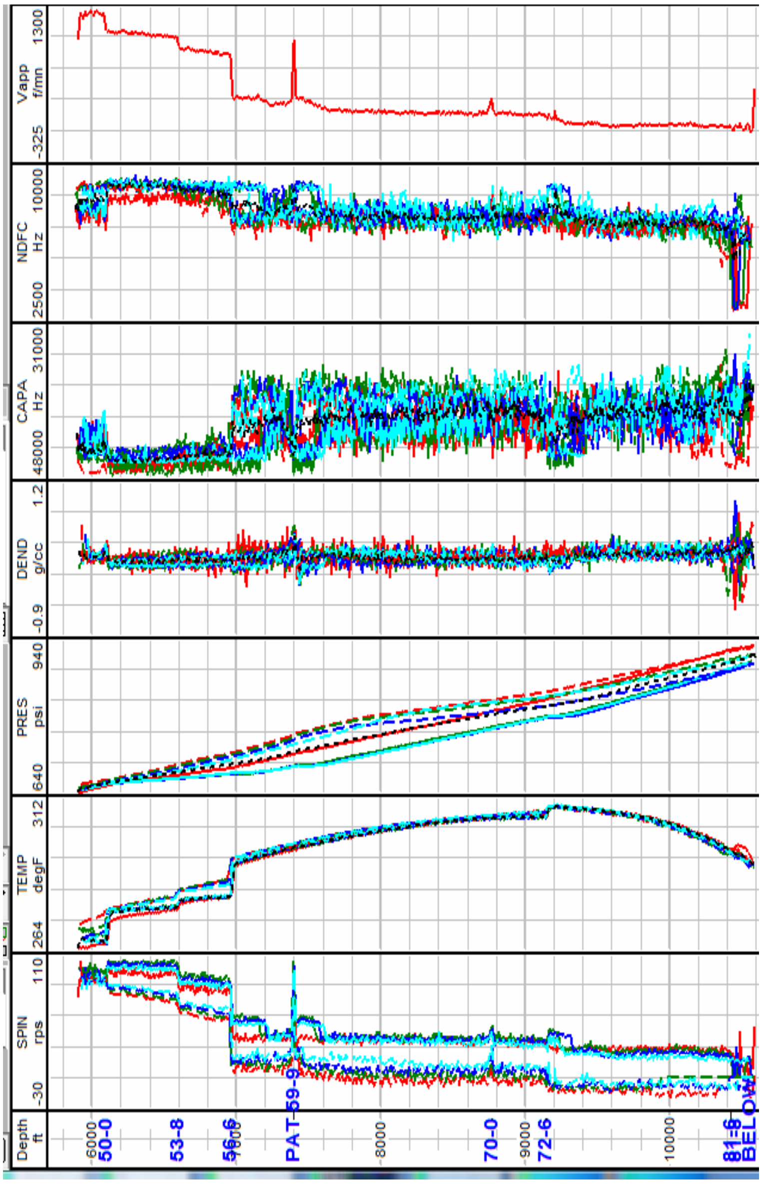

Figure 7 A production log from a well in Thailand

5.1.4 Case study observations from an engineering office and Sirikit oil field, Thailand

As described in §2.3.3, just as social epistemology has criticized the privilege and focus given to individual epistemic agents in epistemology, social research in the philosophy of science and PET has criticized the 'myth of the lone scientist/engineer'. It is rarely the case in scientific laboratories that individuals define the limits of their research to their own experiments and results and this is especially true of engineers. Aside from the obvious multi-personal cooperation that is involved in drilling and extraction on oil and gas rigs, today much of the work of analysing production rates and planning future developments takes place in large office blocks with hundreds, sometimes thousands, of employees. Every company is structured in its own way, although there are of course many near-universal practices and standard *modus operandi*. The company I worked with to produce the case studies and research in this thesis is a multinational company with 17 offices in the U.S. and Canada and further 17 throughout the

Eastern hemisphere. The first foray into South-East Asia by the company was in 1978 when an office opened in the Phillipines (48 employees, 26 field engineers).

It was not until two decades later that the company moved into Thailand and established a head office in Bangkok and warehouse and shipping in the coastal districts of Sattahip and Songkhla (42 employees, 23 field). Finally, a smaller Indonesian office was set-up in 2004 (22 employees, 11 field). A small office in China was opened in 2009 (4 employees, 2 field). One company 'unique selling point' is that it employees predominantly local staff with a small number of managers and more experienced engineers coming from abroad (primarily the U.S., U.K. and Australia). The split is approximately 95% local staff to 5% foreign. Field engineers are exclusively local (and locally trained excepting a number of compulsory training programmes in the US). There are small offices and staff-houses in Kuala Lumpur, Malaysia and Seria, Brunei. The company also sends employees and equipment to Cambodia, Japan, Korea, Myanmar, Papua New Guinea and Vietnam. Clients are split equally between those requiring logging services and those involved in drilling and exploration. These clients include Chevron, PTTEP, EDC, Shell, Total, Petronas Carigali and Schlumberger. Minor clients include Pearl Oil, Salamander, CEC, Pan Orient, APICO, amongst others. The reader will recognize some of these as major multinational corporations and others have a smaller global profile or are local, national companies. Contracts are typically six months to three years with renewals at term or retenders. This latter information is important as, in contradistinction to other industries, companies in the oil and gas industry have a strong motivation to work cooperatively with their competitors on a regular basis as contracts are quite long and it is necessary to keep good relationships with others. This may be contrasted, for instance, with some scientific research laboratories who are often motivated against sharing information and knowledge and competing with other laboratories for research funding, personnel, and so on. (See, e.g., Collins 2001 and de Roure & Frey 2007)

Before we begin our analysis of the finer details of these areas, and how decisions are arrived at and judgments made, it would be instructive to consider, albeit in very general terms, their place in the industry as a whole, its style of business, character and temperament, as well as the petroleum engineer streotype, his ways of working and habits. This, above all, is a conservative industry, protective of its knowledge (rather than, say, academic scientific institutions which may be seen to be quite magnanimous in the spread of information), stuck in its ways, traditional. Units are measured by the same tokens with which they have always been measured, known as Standard Industry Units (SIU), for example, measuring fluid quantities in barrels per day rather than metric litres. Much like philosophy or religion, things take a long time to change. Even on a daily basis, decisions are rarely made in haste and careful, reflective caution is encouraged (partly because safety concerns are made 'paramount', partly because the results of bad decisions may be so costly, and partly because time is not as pressing as it might be in, say, the London Stock Exchange).

Practical restrictions add to this impression of an anachronistic world: many instruments and equipment till recently ran on clockwork mechanisms though modern electronics has replaced many of these systems from 80 years ago. No electricity or other sources of ignition can be near an onshore location hence the liking of mechanical instruments as below; consequently there are no cell phones and as few vehicles as practically possible. A typical engineer is, perhaps counter-intuitively, reluctant to call upon technology to solve every problem that might occur. He—for the profession is still predominantly male—carries a pencil and some paper at all times and likes to sketch out schematics, formula or graphs to describe and work through problems or just to gain an impression of the situation. Mental calculations are not strictly necessary in the digital age but valued because simple errors are easy to miss if one does not know the formula and likely range of results a calculation ought to produce.

5.1.5 Corroboration of evidence

PLA is an excellent example of the multi-heuristic character to technical knowledge and great emphasis is placed on the importance of corroboration or comparison of results, especially with results from other tools. (BP Amoco Upstream Technology Group 1999, p. 88)

Generally, engineers seek to maintain consistency of results between all tools (within the limits of an approved error model). It might seem from a cursory examination of fig 5.1 that the data contained therein would lead a skilled analyst inexorably towards one conclusion. By comparison, if I were an expert in reading tax documents I could relatively easily and conclusively discover the total taxable profits for a chosen corporation. There is only one method for doing this and there would be little doubt about the figure I arrived at. Similarly, in some aspects of petroleum engineering, such as gyroscopic surveying (discussed above and in §7.2.2) the results recorded from the tool are precise, relatively accurate, and it is not a question of expert judgment how to read them.

In PLA, the methods used to interpret well logs are various and non-standardized. Some analysts will favour the results of the 'down-pass' (i.e. taken whilst the tool is travelling down the wellbore) whilst others will favour 'up-pass' data). Some analysts will place a lot of trust in their 'intuitive' judgment or in raw, unprocessed data whilst others will rely on the processes and algorithms computed by proprietary PLA software. I observed analysts using a wide variety of individual techniques for analysing data, particularly when identifying 'good' data. For instance, there is no quantified standard for deleting outlying data points in a correlation curve. Some analysts will have a tendency to leave most data points in even if they are far away from the curve whilst others will remove points which are far away and even ones which are closer than others from the original curve, thus transforming the curve. Often, explanations and rationalizations for why a data point was removed appeal not just to the raw data available for those particular records but to prior knowledge about the well and other data records. Other explanations conceded that it was questionable whether the data point should have been removed and that I could leave it in if I wished (whilst relying presumably not merely on whim but on prior experience). This was referred to by one participant as the 'fudge factor'.

In general, it is good practice to corroborate one's results with other data. Therefore, if an analyst concludes from the temperature log (the log from the part of the tool measuring velocity) that there is *x* oil 'thieving' at zone 5 (i.e. leaking from inside the wellbore to the reservoir either through a perforation or through a hole in the well casing at a particular depth), then she should seek to corroborate this figure, *x*, with another piece of data. However, there is no particular single data set with which one can always corroborate one's analysis. One must rely on judgment and experience to decide whether the spinner log or the derived density is a reliable indicator of fluid properties and there is little in the way of a generalized law for supporting those judgments.

In other words, PLA is an instantiation of Hudson's Wittgensteinian 'river-bed' analogy: the idea that there are no universal rules regarding what information counts as foundational but only rules that some information is foundational in certain cases and that what is foundational (read: axiomatic) in some cases may be a corollary in others. Recall that Hudson argued against the idea that some propositions are invariably fundamental by appeal to Wittgenstein's anti-skeptical analysis of belief. (Hudson 1986, p. 124; and Plant 2005, p. 54) Wittgenstein had observed that there is nothing in our experience that allows us to know how we should interpret our experience (whilst making the anti-skeptical argument that one only asks skeptical questions once one has grounds for doubt).

> If a blind man were to ask me 'Have you got two hands?' I should make sure by looking. If I were to have any doubt of it, then I don't know why I should trust my eyes. For why shouldn't I test my eyes by looking to find out whether I see my two hands? What is to be tested by what? (Wittgenstein 1979, §125)

> If the shopkeeper wanted to investigate each of his apples without any reason, in order to play safe, why doesn't he have to investigate the investigation? …Here all psychological terms merely lead us away from the main thing. (Wittgenstein 1979, §459)

Hudson suggests that we should give up the search for universal foundational propositions (such as Descartes' cogito or that there is an external world) whilst retaining the idea that some propositions are foundational in some cases. Thus, some propositions are like the shifting sands of a riverbed: they are derived from other axiomatic principles or based on other knowledge. Others resemble the immovable bedrock beneath: they are foundational and support the sands above. However, even at the deepest bedrock level, there is a certain amount of shift and that which was bedrock may become sand. (Hudson 1986, p. 125) This renders the distinction almost meaningless but we shall preserve the distinction on case-by-case bases. Similarly, PLA is full of 'bedrock' and 'sandy' data. In one case, participant SK had to analyse the reliability of the well log data based on the surface water (i.e. water production) data. This is quite a routine process but involves a certain amount of judgment in how one handles the various data records. The analyst decided to make the temperature and density curves match the results for the surface water. However, about a week later the office received a report which stated that the surface water value was wrong. Consequently, the temperature and density curves had been corroborated and adjusted to fit false data and were no longer reliable themselves.

There are, however, exceptions to the 'bedrock' analogy and we ought not presume that *all* analysis involves this kind of corroboratory work. Data from gyroscopic tools, for example, must also be corroborated but are always corroborated with the same data. Gyro data, such as that acquired through MWD (measurement while drilling) must pass a 'comparison test' with external data. (BP Amoco Upstream Technology Group 1999, p. 147) Nevertheless, it will often be necessary still to compare the results with results from independent surveys from other tools. Any error in the gyro reference survey will likely be propagated down the remainder of the well and this error will not be detectable from a comparison with previous corrected data.

The gyro tool was made up at the surface and a time delay set (it will start taking surveys at a predetermined time). It was dropped down the drill string prior to a trip (i.e. the PLT survey). In high angle holes, the tool could have been pumped to the bottom but this was not necessary. The tool landed on a baffle plate (a 'TOTCO ring') in the bottom hole assembly (BHA). As the drill string is tripped out of the hole, the survey engineer keeps a tally of bit depth against time, which will later be used to assign a measured depth to each timed survey. When the tool is recovered at surface it is plugged into the engineer's surface computer, which reads the tool's memory and computes the survey. To give the reader an impression of the distances involved, the well at Sirikit is an injection well approximately 6-7000ft deep (about half as deep as a typical offshore well). There are six valves or 'Christmas trees' which lead to points in the reservoir over a distance of a few hundred yards squared. The reservoir itself is approximately one square mile.

One engineer stated one of the working principles of the team: 'If something can go wrong it will.' During the trip, one significant problem occurred. The PLT is configured with a number of tools and sensors for recording various measurements. One of these tools, called a 'spinner', measures velocity. The spinner looks like a fan and, unsurprisingly, spins as fluid passes through it. It will take two surveys: one when being dropped down the drill string (the 'down-pass') and one when being recovered (the 'up-pass'). On this occasion it was the spinner that jammed the PLT in the wellbore and it became difficult to extract it. When recovered, some discussion ensued as to what caused the problem. A number of possibilities were considered such as sand or heavy oil blocking the blades. The spinner was tested by laying it on the ground. The first solution was to physically force the sides of the spinner in so that it took up less space in the narrow pipe. After this failed, the 'collar' of the spinner was moved up. The engineers discussed whether this would prevent the spinner from spinning due to it being too tight. Finally, a solution was proposed where the blades were replaced with shorter ones. This involved a judgment call that the shorter blades could accurately record fluid velocities at various points. It was decided that, since the well had quite a strong injection point (approx.. 5,000ft) they could 'get away with shorter blades'.

This problem and its solutions demonstrated a number of things about the engineering method for these engineers. Rather than following codes of practice or well-established procedures, problems such as these require ingenuity and creativity regarding the physical nature of the problem as well as consideration of what would produce reliable data. The combination of these two factors, with no set procedure for action, means that there is a significant reliance on the engineers' experience and their cooperation to handle the problem and talk each other into a solution.

5.2 Textual analysis of referencing practices

5.2.1 What do the texts say about engineering knowledge?

So far, there has been little constructive interaction between epistemology and engineering, in stark contrast to the multiple interfaces between epistemology and the natural sciences. The reasons for this seem based largely in institutional and historical contingency rather than academic utility. A quick search in the *Blackwell Companion to Epistemology* (Oxford: Blackwell, 1993) reveals no mention of 'engineering', 'technology', 'technical', 'tool', or even 'design'. 'Function', another crucial engineering term does appear but those sections contain no discussion of the use of the term in engineering or even with those philosophers who have drawn on engineering resources (see §4.1.4). The *Concise Routledge Encyclopedia of Philosophy* (London: Routledge, 2000) contains four entries on 'technology' and 'engineering' although only in relation to ethics. If philosophers strive for completeness in a general account of knowledge it does not seem adequate that they do not, typically, mention or think about such a ubiquitous influence on our behaviour and inquiry. Unsurprisingly, a search of several hundred engineering textbooks revealed no mention at all of 'epistemology' and mentioned 'philosophy' only in the colloquial, non-technical sense i.e.:

> The philosophy of LDM [low-dimensional modeling] is to develop very simple physics-based simulation models that contain sufficient information for implementing flow control and insight into the dynamics of the large eddies (coherent structures).
>
> This is the philosophy behind the time-delay networks and the so-called delayed decision making approach.

Significantly, for us, the words 'knowledge' and 'know(s)' were used frequently:

> Accurate knowledge of the course of a wellbore is necessary.

> With sufficient local knowledge results may be quantified at higher water hold-ups but otherwise it should only be used as a qualitative device.

Knowledge is clearly central to what engineers do but they use the word vaguely and to mean different things. Walter Vincenti, who we have spoken about before wrote,

> 'Engineers use knowledge primarily to design, produce, and operate artifacts…. Scientists, by contrast, use knowledge primarily to generate more knowledge.' (Vincenti 1990)

Here Vincenti is making a distinction between engineering knowledge and scientific knowledge. We have previously seen that such distinctions are problematic (§2.3). Nevertheless, there is a kind of knowledge, whether it is unique to any one profession, which primarily aims to 'design, produce and operate artefacts.' To substantiate Vincenti's claim, I analysed some textbooks from my fieldwork to assess how they intend their readers to use technology and engineering to generate knowledge.

5.2.2 The method

The best way to explain CAT's method is through example. Five large texts related to these techniques and professions were selected: two generic textbooks, one on well logging (Darling 2005) and one on production logging (Scientific Drilling International, Inc. 2010); one handbook for surveyors and directional drillers at BP Amoco (1999); and two operations manuals for specific tools used by Scientific Drilling International, Inc. (Kaulback 2009; Hawkinson 2000). It is not my purpose here to parse the complex hierarchies of authority present in the communities under investigation but I will assume that a widely-read textbook or manual will have been accorded a significant role in classifying the objects of its subject (although never an absolute one). (See §3.1.7) Each reference to 'tool' was selected and defined according to what else was mentioned in that sentence. In consultation with engineers and production logging analysts at the companies acknowledged at the end of the thesis, these mentions were broken down into five categories: Informational, proper function, reliability, societal, and other. 'Informational' was broken down into three further categories: measurement, transmission and interpretation. Each of these three was thought to be related in that they referred to the capacity of a tool to provide information about the natural (and sometimes artificial or social) world.

Category I. 'Measurement' mentions are those which refer to the capacity of the tool to measure some property. 'Transmission' mentions refer to how, why, in what form or for what purpose (and so on) the tool transmits its information. 'Interpretation' mentions include those which describe how an analyst or engineer receives, interprets, analyzes and uses this information, formulae used, and so on. For example,[71]

> The Dresser Atlas Spectralog tool measures the counting rates in a number of "windows", each of which spans a certain energy band (Serra 1984, p. 114).

This would be counted twice: once as a tool-reference to what the tool *measures*; once to the proper function (see below) of the tool. The following is an example of transmission,

> Tool will take a set of data, and display output on screen (Hawkinson 2000, p. 16).

---

[71] (Serra 1984) is not one of the texts analyzed.

Category II. 'Proper function' is a more straightforward category relating to what the tool is *for* and how to use it in order to achieve that function. In other words, the 'proper function' of the tool. We also decided to include in these references to how the tool works, references to the structure, calibration, or alternative features of the tool. For example,[72]

> The Teledrift tool comprises of a pendulum that moves along a series of graduated stop shoulders and a signaling plunger at the top that traverses a series of annular restrictions to produce pressure pulses in the mud flow (Scientific Drilling International, Inc. 2009, p. 6).

An example of a reference to an alternative calibration or feature would be,

> A second mode of operation, known as High Speed, allows the tool to survey on-the-fly once inclination reaches 15 degrees or more (Hawkinson 2000, p. 6).

Category III. 'Reliability' is, perhaps surprisingly to laypeople and of particular interest to reliabilist epistemologists (possibly the most

---

[72] Scientific Drilling International, Inc. 2009 is not one of the texts analyzed.

common breed in epistemology), a frequent and persistent concern on a par with the previous types of reference. We include here a broad definition of reliability comprising statistical and probabilistic uncertainty in data provided by the tool as well as the advantages and disadvantages of employing particular tools over others in different geographical locations, rock formations and situations. The following would be an example of a tool-reference to advantages a tool has over another and to the uncertainty or error possibilities in the measurement.

> The Schlumberger Natural Gamma ray Spectrometry tool uses five windows, making fuller use of the information in the spectrum so as to reduce the statistical uncertainty on the analysis of Th, U and K [thorium, uranium and potassium]" (Serra 1984, p. 114).

An example of advantages/disadvantages may include references to current recommendations or histories of the tool.

> …[T]he FDP [fluid density differential pressure] is maintenance intensive and involves the use of mercury, so this tool is being phased out & will be replaced the fluid density delta-p (FDD) tool. (*sic*) (Scientific Drilling International, Inc. 2010, p. 10).

Category VI. Whilst accepting that, from an SP and CAT perspective, it is methodologically unsound to disentangle social interests from ontological and epistemological claims, we reserved a category for what we have called 'societal' references. This is taken from a plausible insight into the petroleum engineering world that certain issues are isolated and often precedent to other concerns. Safety and risk are two such issues which are given high prominence in the education, instruction and day-to-day life of engineers in this field. This is unsurprising especially when one considers the catastrophic, sometimes fatal, effects accidents, as witnessed in the recent Deepwater Horizon blowout.[73] Reference to hiring costs of equipment, the organization and hierarchy of personnel, who to speak to in order to obtain a particular tool, and so on, where also termed 'societal'. For example,

> All non-essential personnel will not be allowed to be in the direct vicinity of the source as it is being loaded into the tool. (Kaulback 2009, p. 130).

---

[73] A timeline of which incident is currently available online at http://www.offshore-technology.com/features/feature84446/. 2010. Deepwater Horizon: A Timeline of Events. *Offshore Technology* (Net Resources International).

> As can be seen from the diagram below the “hottest” point at the surface of the tool gives 2. 10μSv/hr so it would not be possible to reach anything like a classified workers exposure limit in a week! (Scientific Drilling International, Inc. 2010, p. 15)

Category V. ‘Other’ references include those to naming conventions, references to other documents, and so on.

References that purely served the format of the text (for example, in indexes or tables of contents) and those in discussions of what follows or of the layout of the chapter were excluded. It was not clear at the time how these five categories would be proportioned but it was clear that the first four would cover the vast majority of tool-references in the five texts.

5.2.3 Analysis of results

The word 'tool' features throughout these texts as a catch-all term for technical artefacts which exclude semi-permanent artefacts such as rigs, pipes, and so on. In total 1708 tool-references were found in the five texts.[74] The results reveal one part of the big picture of how production loggers and directional drillers (and surveyors, wellbore navigators, etc.) refer to tools. It should be noted that technical operations manuals are likely to contain many mentions of the word 'tool' because it will contain repeat and precise instructions for what to do with the tool and each of its configurations or variants. One textbook (Kaulback 2009) was unusual among the selection in having almost 60% of its total tool-references refer to how to use the tool or how the tool works – in other words, the proper function of the tool. The other four textbooks displayed a statistically even spread across three types of reference: informational, functional and reliability. 3-11% of references to a 'tool' were of societal or other type.

---

[74] The results are compiled in Table #.1 in the Appendix.

If we exclude the operations manual mentioned above (which I suggest skews the significance of each reference by repeating very similar mentions), 32-38% of references were informational and 22-31% were to proper function. Finally, 24-37% of tool-references are to do with the uncertainty, error, or 'corruption' of the data provided by instruments. The average proportion in that case calculates as follows. Approximately one third (34%) of tool-references to what the tool measures/communicates or how these communications are interpreted. Approximately one third (32%) to the uncertainty of data, etc. Approximately one quarter (26%) to proper function. The remainder (8%) to 'societal' and other things.

5.2.4 Extrapolating results to a conceptual sketch of technical artefacts

These results indicate that the texts use three basic types of reference in a full description: they are functional instruments that measure certain properties within a certain range of accuracy. These three elements seem indispensable for describing a 'tool' to a sufficient degree of detail. So far the case study does not, in itself, support a collectivist account. In fact, all case studies such as this are able to do is build up a body of research which can be given a collectivist analysis and from which analysis fruitful insights may be drawn. Future research may be able to generalize the results of this study or perhaps adapt the criteria in order to make it more generalizable.

5.2.5 Proper function

It would seem that one of the key distinguishing features of our reference to tools and reference to other kinds of things is that tools have proper functions. How does CAT conceive of this notion? Before we turn to a consideration of proper function of artefacts let us first discuss the normative aspect of tool use. As we shall see, the former requires the latter. The concept of normativity is a recurring theme in Strong Programme research and Schyfter presents it as a key component in how we conceptualize technical artefacts. (Schyfter

2009) Normativity is here conceived as statements of properness, appropriateness, correctness, etc. and their antonyms.

David Bloor has previously made an argument for such judgments being community-dependent. According to Bloor, if we wish to judge the properness, for example, of an instance of use we must look at the standards of the relevant community. Bloor argued, based on Ludwig Wittgenstein's private language argument, that the distinction between 'is right' and 'seems right to me' could only be made sense of with a community to provide what is right (Bloor 1997).[75] For example, assuming no direct access to a judgmental God, we can only make sense of what a 'proper function' or a 'mistake' is by asking what our peers would agree is proper or mistaken. Consequently, the community is prior to the individual in the order of explanation of proper function and a collectivist account is required.[76] MacIntyre made the same point

---

[75] For a discussion of Wittgenstein's argument and its connection to the Robinson Crusoe example discussed later in this paper see, for example, Blackburn (1984), Misak (1995, p. 54), Wright (1986, p. 235).

[76] Here I appropriate Kusch's (2002) claim that the existence of knowledge is dependent on the existence of communities. Consequently, in order to know the proper function of an artefact within a community, one must explain the normative standards of the community. For an extended discussion of knowledge of artefact function see Houkes (2006).

when asserting the need to index the notion of 'mistakes' to a culture (1967). MacIntyre was responding to Peter Winch, a significant influence on the Strong Programme, and on Bloor in particular. (Winch 1958, 1964) Normative judgments are indispensable to any textbook or manual on a tool. Sometimes the normative judgment is explicit:

> From surface until an inclination of 10°-15° is reached, surveys *must* be taken with the tool stationary. (BP Amoco 1999, p. 155. Emphasis added)

Other times it is implicit and hidden:

> The scope is then sighted up on the pre-determined reference point and confirmation given to the surveyor in the wireline unit that the tool is aligned (give thumbs up). (Hawkinson 2000, p. 75)

Although this appears like a pure description of what happens, in fact we can see that it contains a normative imperative equivalent to, "Sight up the scope on the pre-determined reference point…!" Any imperative or injunction of this kind is a recommendation of what ought to be done and we cannot make sense of such judgments individualistically but must consider the standards of the community. A description of a technical artefact is rarely just a description. In most cases it is also a judgment on what the artefact is supposed to do or, to put it another way, what it is for. The collectivist enterprise is tasked with analysing these judgments, how they arise and how they operate. That is not to say that CAT must make normative claims about how the artefacts ought to be used, just as SP researchers typically claim not to make presuppositions about the scientific knowledge under investigation. It is to say that the practice of making normative judgments itself cannot be absent from a comprehensive analysis. The same cannot be said of natural kinds: there is no such thing as a 'good stone' per se. A stone can only be so referred if it has been appropriated for a tool-like purpose e.g. 'a good stone for grinding food'.

In the Cratylus, Plato argues (390b-d) that the user of an artefact knows it better than its maker. Although the answer to whether Plato was right probably varies depending on what kind of artefact and what kind of design and manufacturing process you have in mind, for most artefacts, it seems he is just wrong. Designers and manufacturers must be able to second-guess the user's intentions for an artefact and the user must, in most cases, have some idea of the blueprint or 'use-plan' (Houkes & Vermaas 2004) in order to use the artefact properly or *qua* artefact. Use is central to any study of technical artefacts. Consequently, I am sympathetic with Plato's argument, despite what is said above but would like to extend the group of 'users' to those designers and makers who engage in using the artefact when they imagine how their consumers will behave. Thus, they contribute to the community use-plan, that is how the proper function of the artefact is perceived by the relevant community of designers, manufacturers and users.

5.2.6 Conclusions

Despite the above counter-intuitive consequence, CAT provides a rigorous methodology for approaching the ontology of tools. Non-collectivist accounts may also succeed by detailed analysis of individual behaviours but this seems an unnecessarily myopic approach akin to using molecular chemical models to analyse the dynamics of fluids in the engineering sciences. Just as the latter has developed macroscopic models for the analysis of structures and properties that suit that scale of analysis, so should sociophilosophers take the community as the primary unit of analysis when investigating ontologies of natural, social and artificial kinds. In this chapter I have presented the framework for a collectivist analysis of N/S/A kinds. I have focused on technical artefacts as the locus for an interesting study into this framework and described one element of how that study can be conducted: a textual analysis of referencing practices in petroleum engineering. I have discussed criticisms of Barnes and Kusch's tripartite ontological framework and problems for the collectivist approach (which, unsurprisingly, are comparable to the problems faced by collectivist approaches in epistemology, moral philosophy, and elsewhere). I have suggested that we look at distinctions, oppositions and dichotomies as working oppositions rather than logical binaries and look to the heuristic value they have for those who employ those distinctions. I have described how one community of petroleum engineers refers to tools and given reasons why understanding the normative judgments of the community are vital for understanding these referencing practices. The outcome supports further research using the collectivist account to analyse classifications of artificial kinds.

5.3 Technical knowledge

5.3.1 Do you know how your radio works?

§5.2 described how engineers have created a distinctive ontology regarding technical artefacts which may or may not be commensurable with a natural scientific taxonomy. We also discussed and investigated Vincenti's claim that engineering knowledge aims primarily to 'design, produce and operate artefacts'. This chapter describes the implications of this for a contextual social account of knowledge. In other words, it describes how different ontologies lead to different kinds of knowledge.

Do you know, for instance, how your radio, a technical artefact, works? What would be required in order for you to know this? Presumably, the answer to this varies according to the context in which the question is asked. Suppose that you are visiting an elderly relative who is quite technophobic and does not know, for example, how to use a mobile phone or personal computer. You might bring them a radio so that they can listen to it when they are otherwise unoccupied. You might ask them, 'Do you know how the radio works?' Here, you are not asking if they know how electromagnetic waves are transmitted through the air by a radio station and the antenna on your radio picks up the signal by tuning its current to the same frequency as the radio waves sent out by the station. The mechanics or physics, as it were, of how the radio works. Rather, you are asking if the elderly person knows how to use and control the radio in order to pick up the particular channel they wish to listen to; if they know how to adjust the volume; and if they know how to store particular channels in the radio's memory. The same question, asked in a school electronics classroom might elicit a response similar to that given above.

5.3.2 Knowledge of artificial kinds: “Not a library, but a museum”

We have seen that several authors claim that artificial kinds are rightfully considered as a separate kind of thing from natural and social kinds given that we refer to them in a different manner. Specifically, artificial kind terms refer both to the alter-reference, the physical, structural, corporeal item out there and to the self-referential talk that constitutes them. Intuitively, there would be no such thing as a hammer if there were not both a) a hammer-shaped object with the structural composition of a hammer out there in the world, and b) a community part of whose referential practices is to talk about hammers. A second peculiarity about artificial kinds is that they provide an epistemic value that natural and social kinds typically do not. Rocks and marriages, significant as they are to the world that we live in, seem epistemically ‘inert’. That is, the knowledge that we have due to their existence, is simply knowledge about them: knowledge that limestone is a sedimentary rock, knowing how to get married, knowledge that Fred and Wilma are married, knowledge of geology, and so on. They do provide new knowledge that we would not have were it not for their presence.

However, this is different from the way that, say, a microscope provides new knowledge. The new knowledge in this latter case is not merely knowledge about microscopes but is a new area of knowledge that we did not have available to us before, which we call microscopy: knowledge of all the microscopic things which are now visible to us and which we were not sure before whether they existed. Surprisingly, even within the specialist disciplines of epistemology and philosophy of science, very little attention has been paid to this distinct role of technical artefacts in the acquisition of knowledge.

The purpose of this chapter has been to construct systematically a framework for epistemic systems that includes human agents, of course, but also instruments, tools, technical artefacts. Davis Baird has developed what he calls a 'materialist epistemology for instrumentation'. (Baird 2004). This is conceived of as another kind of knowledge alongside the traditional propositional kind, a language-centred kind. The problem as he understands it is that the study of instruments in the history and philosophy of science has been blocked by our interest in language. Language has been assumed to be the main mode of communication and thus if we wish to reflect on the history and philosophy of science, we must look to publications, primary sources, textbooks, and so on. Instruments are not linguistic objects and so their role in science is demoted to a machinic, servile capacity.

Of course, described in such terms, this is extremely reductive. Instruments and their history has had a formative role on the particular knowledge we have acquired. The subtitle for this chapter comes from Mike Mahoney's evocative proposal to replace the history and philosophy of science as a history of literature with a history of things. (Mahoney 1999) Inspired by Mahoney, Baird's approach is to articulate a history and philosophy by taking apart instruments and exposing the tacit and visual kind of knowledge which he calls 'thing knowledge'. According to Baird, technical artefacts are 'constitutive of scientific knowledge in a manner different from theory, and not simple “instrumental to” theory.’ (Baird 1999, p. 1) The collectivist approach also considers ‘thing knowledge’ to be a particular kind of knowledge but I do not wish to say, with Baird, that it is non-propositional. On the contrary, the tools described in the case study are fundamentally data-centric (see §5.1). They provide analysts with information which they then turn into knowledge claims. Instead of abandoning traditional propositional epistemology, I hope to preserve that same format but bring technology into that account in a formative role. In particular, I wish to say that the analysts develop perceptual skills which they can use to interpret the data log images that the tools provide.

A functional account of tools may lead us to adopt a functional account of the knowledge of tools. Bird (2010) develops precisely such an epistemology in more general terms. He seeks to provide a social analogue for individual knowledge by analysing social organizations as analogous to an individual organism.

5.3.3 Reliability and error

One third of the references in the textual analysis are to reliability issues. Knorr Cetina notes that at CERN, what is done primarily depends on problems which need to be resolved (be it problems of technology or problems of reliability). Similarly, work in the company on the ground and at the office is often about solving particular problems as they arise.

> Much of the well positioning process is concerned with the avoidance of mistakes. Put another way, it ensures the delivery of objectives which most people involved in the upstream business take for granted. (BP Amoco Upstream Technology Group 1999, p. 19)

This quote encapsulates the emphasis on error. Analysts are counselled against relying on data as accurate. They must always corroborate, compare, contrast, correct, calibrate, and coordinate. This in itself suggests a different kind of knowledge from the kind we are used to in other areas of inquiry. Compare the petroleum engineer reading a well log with, say, a lawyer pronouncing on a matter of law. In the latter case, the lawyer can say, in most instances, with a high degree of confidence, that her assertion is accurate. If she needs to check, she can look up the relevant statute or case law. If she has doubts, she can ask a colleague. At some point, they will come to an assertion that very closely resembles the truth. This is because truths in law are matters of convention. They have been stipulated by the community and have been stipulated in such a way as to make accurate interpretations likely, in so far as this is possible. Any debates can be quite easily settled by referring to relevant texts. Truth, insofar as it may said to come into the picture here, is decidedly pragmatic. The aim of inquiry is not to acquire unassailable facts but to get knowledge *good enough that it works*.

5.3.4 Conclusions of a textual analysis of reference to 'tools' in directional drilling and production logging analysis

Based on the theoretical discussion in chapter 4 and the results of §5.2 we have the following summary of technical knowledge:

1. Technical knowledge is a distinctive kind of knowledge.
2. It includes use know-how as well as knowledge of how to design, produce, and operate artefacts.
3. This knowledge is a central component of engineering and provides a distinction between the knowledge possessed by engineers and that possessed by natural scientists (although there is undoubtedly a great deal of cross-over between the disciplines).
4. Technical artefacts have a proper function. Comprehensive knowledge requires knowledge both of this proper function and its empirical characteristics.
5. Proper function is constrained by normative considerations which are collectively and performatively sustained.
6. Technical artefacts are an artificial kind.
7. Artificial kinds are social institutions (due to self-reference)

8. Technical knowledge (knowledge of technical artefacts) is knowledge of a social institution.
9. Proper function is known, in the main, through testimony.
10. This substantiates Kusch's claim at least in this case that the community is prior to the individual in the order of explanation.
11. Engineering knowledge is social.

If the findings of the textual analysis alone hold more widely, Crusoe's construction cannot be accurately referred to as a tool as we cannot refer to it in similar enough terms. We cannot compare Crusoe's use with an accepted code of use (whether written as instructions or embedded in a community of users); we cannot know if it has achieved successful results as the success conditions are community-relative. Instead, because Crusoe has no community of his own by which to make these comparisons, the collectivist analyst has to compare Crusoe's usage to that of an existing type of artefact in another community to make sense of the artefact *qua* artefact. In other words, we must say 'Crusoe is using a knife' and adopt Crusoe temporarily into our own community in order to make sense of what he is doing. The counter-intuitive conclusion is that Crusoe can only make tools if we posit him as part of another community. The advantage here is that the category of tool does not collapse into anything a person happens to use to fulfill a purpose, a conclusion that would be incompatible with the outcome of my analysis of tools used in engineering.

Chapter 6

Seeing (and knowing) with technical artefacts

6.1 Extended epistemic systems in action

6.1.1 Introduction to extended epistemic systems

> Insofar as theoretical epistemology gets naturalised into a chapter of theoretical science, so normative epistemology gets naturalised into a chapter of engineering: the technology of anticipating sensory stimulation. (Quine 1990, p. 19)

In previous chapters I have emphasized the need for a social analysis of knowledge. I argued that we should examine the textual reference-talk within communities to understand what is meant by a variety of concepts. This provides an empirical method for analysing knowledge attributions. However, this alone would not be enough to characterize the epistemic practices of engineers. It would not capture the reference to the non-social world that their epistemic practices inevitably involve.

As we saw in Ch. 4, technology is not a purely social matter. What is required is a combination of social and physical components akin to that outlined in the DNA and CAT approaches. The next two chapters attempt to give due attention to the non-social (or perhaps I should say less social) elements of knowledge in petroleum engineering. In this chapter I look at the processes of perception involved in analysing well logs and other problem-solving and in Ch. 7 I consider whether the best fit for the analysing this knowledge would be an informational approach.

I focus here on perceptual knowledge. It may be the case that what I say about the sociotechnical nature of perceptual knowledge is generalizable to other kinds of knowledge but the strongest case I can find is that of perceptual knowledge. I introduce the philosophy of engineering problem of unobservables and discuss issues that arise from this problem. I discuss perceptual knowledge in epistemology and cognitive science and assess the merits of a model based on Bayesian inference to production logging analysis and gyroscopic surveying. I further argue that knowledge of this kind is acquired by an extended epistemic system. Extended epistemic systems are, put simply, human agents plus *x*.[77] This *x* could be a technical artefact that provides information and extends ordinary human epistemic capacities (such as an altimeter, microscope, or pair of spectacles); or it could be another human agent (such as in cases of testimony or joint inquiry). In these cases, the human agent alone does not produce knowledge, nor does the *x* produce knowledge but the system as a whole co-produces knowledge. I then explore in more detail the sociotechnical structure of one institution and its practices. In this chapter I will look at three ways in which information is carried by instruments in production logging analysis and wellbore navigation, two niche fields of petroleum engineering. I consider how fruitfully each can be seen as sources of perceptual knowledge.

Such is the extent to which technology encroaches upon the processes we think of as entirely mental that the closer one analyses what is

---

[77] The idea that certain cognitive processes extend beyond the boundaries of the human skull has been discussed extensively in research in extended mind theory (Clark 2000, 2001; Clark & Chalmers 1998) and in relation to epistemic processes in Goldberg 2007, 2012; Hetherington 2012; Marsh & Onof 2008; Pritchard 2010b; Roberts 2012; and Vaesen 2011a.

occurring in, let us say, an engineering analyst's professional work, the less clear the natural-artificial distinction becomes. Indeed, when one introduces the element of socialization which is undeniably relevant to most instances of knowledge it is not possible to draw this fine distinction between skin-and-skull perceivers and extended social corporate agents. If the reader is not convinced of this point at present I hope they will be more sympathetic to why I have taken this approach by the end of the chapter. To illustrate this claim, consider a typical case of perception: an individual agent sees that an object is blue. She does this without the aid of any technology (she does not use, say, a microscope or colourimeter to acquire knowledge of the object's colour).[78] In another case the agent is still acting alone but wearing spectacles or contact lenses. Clearly this does not normally affect the epistemic status of her belief.

In another case she is underwater in a swimming pool and seeing through an artificial solution of water, chlorine and perhaps some

---

[78] A colourimeter is a device that measures the absorbance of wavelengths of light by a solution and hence can be used to determine colour.

other compounds such as cyanuric acid or bromide. Again, it would be churlish to say that now she does not perceive the colour directly but her perception is mediated through the artificial substance in some way which does not occur when there is only air between the eyes and the object. In yet another case the agent is looking through a video recorder at the object. Yes, this is technological extension of a time but to suggest that her epistemic status with respect to the colour of the object is in any significant sense compromised or dependent on the technology seems tenuous. As a counter-point, apparent cases of perception where technology is involved are similarly opaque. In many cases, the processes involved are substantially similar.

As shall be demonstrated in 6.2, it looks increasingly unlikely that sound, stable, non-stipulative, non-arbitrary grounds for distinguishing between a natural organism's eyes and cognitive processes and the processes of an extended organism such as a microscopist or a data analyst (meaning the combination of a person and an instrument or instruments of technology) can be provided. In §6.1.5 I look at the particular technologies in my fieldwork research, how they are used, and how the transmit information to the analyst. Before doing so, however, there is another step that must be made. Even if both extended epistemic systems and individual agents perceive in comparable ways, there remains a difference, namely, the apparent immediateness of the latter as compared with the more elongated, separated transmissions of the former. This is the question I now turn to.

6.1.2 Is perception immediate or inferential?

I have phrased the question above as I have because it may mark a difference in approach between some analytic philosophical theories of perception and a growing research programme in cognitive science. I would like to suggest that the distinction between what we phenomenologically describe as 'immediate' experiences and the more deliberative reflective experiences is tenuous and we should treat all perceptual experiences as belonging to the one spectrum of inferential cognitive processes. Some epistemologists have argued that perception can be immediate and have appealed to the phenomenological experience of seeing what something is without any noticeable reflective awareness. (McDermid 2001, Millar 2000) Alan Millar, for instance, writes that,

> [Perceptual knowledge] is phenomenologically immediate, in the sense that it is not acquired via inference from prior assumptions. If you know perceptually that something is a bottle of milk then it simply strikes you that this is so on seeing it. (Millar 2000, p. 73)

Millar adds that this non-inferential kind of knowledge is not limited to that which is perceptually manifest either (that is, facts concerning the way things appear). Millar expands on this position elsewhere thus,[79]

> What makes sense of the perceptualist model is a conception of *perceptual knowledge as non-inferential knowledge acquired by suitably equipped subjects from what they perceive*. Perceptual knowledge is the kind of knowledge we acquire when we tell from its look that a bird is a magpie or that a flower is an orchid. It is non-inferential knowledge in that it is not acquired by reasoning from prior assumptions. It is phenomenologically immediate in that what is known simply strikes the subject as being so. (Millar 2008a, p. 591. Original emphasis.)

---

[79] To put the following quote in context, Millar is here arguing that knowledge of other minds is a kind of perceptual knowledge. See also Millar 2008b.

I wish to argue in this chapter that Millar's position mischaracterizes perceptual knowledge. Not only is it incorrect to say that perceptual knowledge is non-inferential but it is unduly restrictive as inferential perceptual processes can confer the same epistemic credit upon beliefs as more 'immediate' processes such as the everyday perceptual cases Millar cites.

In §2.2.2 I alluded to the acceptance of constructivism within the biological study of the eye. I remarked that perceptions have been shown to be constructs in the sense that they are fabricated by the perceiver's perceptual apparatus in various ways and not invariable representations of visual stimuli. Let us now expand on that starting point. It would be naïve to assume that we are all equally good perceivers. For example, some people are excellent at recognizing faces; who are sometimes referred to as 'super-recognizers'. (Russell, Duchaine, & Nakayama 2009) Others have cognitive disorders that prevent them from recognizing the faces of even their own family. Researchers have discovered that people will 'hone in' on different clues in order to match a face to a name. Sufferers of prosopagnosia lack most of these clue-identifying abilities and some have learnt to cope by, for instance, remembering the colour of someone's shirt or the kind of shoes they wear. None of this variation in the way the eye and the brain construct images puts into jeopardy their epistemological standing and we frequently and uncontroversially attribute knowledge on the basis of beliefs formed on the basis of perception. Now, constructivism about 'higher order' cognitive processes does not easily fall out of this fact. Constructivism about knowledge and social constructivism do not have the same empirical scientific backing that constructivist perception and psychology do.

We know a great deal about how the eye perceives. Some of the more elementary details of this I have outlined in the following chapter. We accept that merely because these tasks are inferential, this does not put their epistemic standing in jeopardy. There is generally no problem with attributing knowledge on the basis that the knower has seen the thing in question (indeed, some cultures will only attribute knowledge on the basis that the speaker has seen the thing in question). However, when we make the move from saying that an eye perceiving is unproblematically constructive to a whole person inquiring, the thesis becomes much harder to accept.

One of the key findings of my fieldwork research is that the more familiarity, skill, and prior knowledge we have in a particular perceptual task, the more it begins to feel like an immediate experience. Our minds, I suggest, are so familiar and practised in perceiving common everyday objects that we never engage in any reflection upon whether or not what we are doing is perceptual. Millar's argument and mine are both based on intuitive or phenomenal grounds: for Millar, it *seems* like we have immediate perception; and for me I note the phenomenological experiences of analysts working in the field.

Granted, it takes a great deal of additional skill and practice to do what I observed many data analysts such as the logging analysts in an engineering office doing: that is, to feel like those data analyses are also perceptual. When we examine the differences between what our eyes are doing and what our eyes-plus-glasses or eyes-plus-microscopes or eyes-plus-computers are doing it may no longer be surprising that sufficient practice in any one of these tasks will lead to very similar phenomenological experiences. That means that most cases of seeing ordinary objects will appear to require no cognitive effort at all. We just *see* that the tool is on the table. Some authors conflate this phenomenological 'immediacy' with an epistemological privilege or priority. However, there does not appear to be anything in the timeliness of an experience that confers upon it a particular epistemic privilege. By way of example, we do not think that the first belief that comes to mind when considering a mathematical equation is more reliable or more likely to be true than one that is formed later.

Jerry Fodor hints at the inferential character of perception when he writes that,

> Perception must involve hypothesis formation and confirmation because the organism must somehow manage to infer the appropriate task-relevant description of the environment *from* its physical description together with whatever background information about the structure of the environment it has available. (Fodor 1975, p. 50)

The connection between this characterization of perception and the characterization of PLA I have given is brought into relief in the following statement when Fodor compares perception to learning,

> Perception is essentially a matter of problem solving, where the form of the problem is to predict the character of future sensory experience given the character of past and current sensations as data. Conceived this way, models of perception have the same general structure as models of concept learning: one needs a canonical form for the representation of the data, one needs a source of hypotheses for the extrapolation of the data, and one needs a confirmation metric to select among the hypotheses. (Fodor 1975, p. 42)

§6.2.2 details a cognitive scientific interpretation of perception which echoes Fodor's account above. It is not known whether this Bayesian approach is a true reflection of our cognitive faculties but it provides an example of how an inferential process might work. (Wolpert & Ghahramani 2006) At the very least there is a case to answer here as to why we ought to accept the intuition-based arguments in favour of the immediateness of perception over an empirically-based argument that perception is inferential. It is likely that some form of inferential perception such as this does more accurately model the visual system and, more importantly, will allow us to make better predictions.

6.1.3 The engineering problem of unobservables

Production logging analysis is a range of practices for acquiring knowledge about the properties and composition of fluids at various depths of a well. This is done primarily to derive the following:

i. The percentages of hydrocarbons (oil, gas and water) at different depths or 'zones,'

ii. Which zones are producing (i.e. have one or more type of fluid entering the zone from perforations made in the pipe or by other means) and in what composition,

iii. Which zones are thieving (i.e. have one or more type of fluid leaving the zone).

Nowadays, this is most frequently a sophisticated procedure involving teams of analysts discussing amongst themselves well-logs such as that pictured in fig. 7. One aspect of this discipline which is epistemologically interesting is that it concerns knowledge about things which are outside our ordinary unaided field of vision but also look similar to perceptual knowledge as described above. Once one has acquired a certain skill in perceiving, discriminating and interpreting images such it is common to describe ones experiences in perceptual terms. The problem of unobservables has been a longstanding research project in the philosophy of science. (van Fraasen 1980; Muller & van Fraasen 2007) I would like to introduce here a new problem that is related but creates its own philosophical issues and questions. The problem of unobservables in the philosophy of engineering is distinctive in that it arises not because entities, events or states are imperceptibly small but because they are deep under the ground. The principle problem is the same: how can we see something which is outside an ordinary unaided human field of vision?

However, unlike in philosophy of science, no one doubts the existence of, for example, a sedimentary rock formation or a volume of water or considers such entities mere theoretical postulates. It is, at least in principle, logically possible to perceive the rock composition outside zone 6. It would even be technologically conceivable even though modern cameras would be expensive and ineffective (light would tend to reflect off bubbles in the fluid and produce blurry images, and so on). The practical solution is also the same in engineering as in science. We are given the opportunity to acquire knowledge about these unobservables because technical artefacts stand-in for human perceptual capacities and enable knowledge of entities, events, and states which would be otherwise unavailable with ordinary unaided human visual capacities. Since we are talking about things which cannot be perceived (which is more general than things which cannot be observed) we might more properly call such knowledge, knowledge of imperceivables.

6.1.4 The Enactive Torch

I wish to submit that technology bridges this epistemological gap between that which is inaccessible to us and knowledge of it, and thus that technology is a formative element in a comprehensive epistemology of this type of knowledge. Take for example a device developed in England known as the Enactive Torch. (Froese et al. 2012) The Enactive Torch is an ultrasonic sensor with a vibro-tactile motor that generates feedback to someone holding the handle, the intensity of which varies according to the sensor's proximity to an object (like a cross between a rumblepad and a bat). Blindfolded participants are asked to use the Torch to perceive the position of a nearby object. The researchers aim towards developing what they call a 'phenomenological pragmatics of enactive perception' which involves considering the strategies developed by users of this device to navigate their environment or detect nearby objects as either extensions of existing or constitutions of new percepts of perceiving or new ways of perceiving. You can imagine that once users become skilled and practiced at manipulating the device that they can accomplish quite adept and sophisticated tasks easily. In this case the instrument produces a signal that carries the information that replaces the information sought by the eyes and hands in ordinary unaided perception.

The developers of the Enactive Torch describe three kinds of strategy adopted by users to accomplish their tasks: cognitive, intuitive and unknown. The cognitive strategy involves something like panning the room until they detect an object. Unknown strategies are where the user was unable to provide any coherent account of their method; these strategies were not successful. The intuitive strategies are those where the user described their process in terms of 'feeling their way around' or words to that effect. In these cases, the torch has become a new set of eyes or, alternatively, a new perceptual apparatus. In such cases the instrument has begun to disappear as an addition or supplement. As Polanyi poetically put it:

> We may say that when we learn [a] probe, or a tool, and thus make ourselves aware of these things as we are of our body, we *interiorize* these things and *make ourselves dwell in them*. (Polanyi 1962, p. 61)

6.1.5 Three sources of perception in petroleum engineering

Production logging analysis is a messy business. PLTs provide us with well-logs which are open to interpretation in a way that reading a compass card inscription is not. There are many values coming in and they often do not easily fit into a single interpretation. One cannot isolate any one of these values and draw reliable conclusions. Put two analysts in a room with the same well-log and there is no guarantee they will answer questions in the same way (particularly if they have differing levels of experience) and it is very likely that they will differ over the specifics.[80]

As such, an issue arises over how judgments are made. Issues are common where a training manual cannot provide the answer, a software package will not provide reliable results, there is no articulable

---

[80] The effects of standardization and regulation of industry practices, training exercises, together with professional and social influences, instead of optimizing strategies, in practice can hide the lack of clarity of interpretation.

algorithm for solving the problem, and analysts disagree. A judgment has to be made. Occasionally, an analyst will refer to this practice in perceptual terms: ‘I can just see that zone 6 is thieving water’. Certainly a lot of background processing is going on in his mind before he makes that kind of assertion involving corroborating data, past experiences, and so on but nevertheless it *feels* like a perceptual activity. Some of us will know the example from the film The Matrix where a computer programmer is viewing what looks to us to be a screen of green binary code but in his mind seeing a woman in a red dress walk through a busy metropolis. With familiarity and practice comes the immediateness of experience that Millar describes. And instantaneousness at least *feels* like perception.[81]

Now I wish to return in the light of this discussion to photographs and the camera that produce them. First let us be clear about the perceptual nature of the camera. If we are looking down the barrel of a telescope we would, I think, have no problem in saying that we see the Moon

---

[81] I am making a claim here about how knowledge is perceived by the community and so how knowledge acquisition feels is significant.

eclipse the Sun, for instance. There is no reason to think this simple reflection of light and configuration of lenses should be a barrier to genuine perception, any more than is a pair of contact lenses. A digital camera, however, may be different. Common sense intuits two differences:

1. The image produced appears fuzzier, grainier, you might say 'less real' having been translated into the 'blocky' texture of digital pixels.

2. If we are not watching the event in real-time but as a recording, this fact may lead us to conclude that we are not really seeing the event.

Does either of these intuitive differences provide a clear distinction between seeing with a telescope and not-seeing with a camera? Take the first: Does the fuzziness of a photograph mean that it is not seen? Imagine you are in a swimming pool with your eyes open. You see a crab crawl along the floor of the pool and into a toy castle. Now, in the second case the water has been drained from the pool. Again you see a crab crawling along the floor and into a castle. In the second case, the image you receive is much clearer than in the former case. The situation is similar to what would be seen if a camera were to be dropped into an oil well (and is the reason why this is not used in production logging as a means for seeing what is going on down there).[82]

However, as long as the image is still interpretable, as it clearly is in the case of underwater perception or acclimatizing oneself to an unfamiliar perceptual situation (such as a trainee cricket umpire might have to do to spot the minuscule actions and events that an experienced umpire would), then we can confidently call this a case of seeing. One drawback of the film-based system is that when the drill-pipe is bent there is reactive torque. A photograph taken under these conditions appears blurry, like the photographs someone might take of a jogger in a park. Secondly, one needs to keep taking photographs to check the

---

[82] Oil and water mixed together are not easy to see through: not for human eyes and not for cameras. Light does not penetrate the flowing emulsion and reflects off the bubbles. In any case, cameras are temperature sensitive and would not cope well with subsurface conditions.

drilling is proceeding in the correct direction. This is plainly impractical and readings such as these are made by resolvers on the instrument (compass face) that makes an electrical signal that is memorised or sent uphole rather than a picture to be photographed. For example, the magnetic single-shot is one of the earliest camera-based surveying tools developed by Eastman-Christensen. The instrument barrel has a small photographic disc and camera unit which can take a photograph of an angle-measuring unit. The camera captures an image of a pendulum superimposed over a compass card.

What about the issue with recorded images? One can look through the viewer of a digital camera and see an event in real time in much the same way as one looks through it to see a recording of the event. Does the difference that this event happened in the past preclude it from counting as perception? Presumably not, since all acts of perception are temporally delayed. If we are seeing a lion chase and catch a gazelle the excitement of the live event causes us to agree eagerly that we are seeing the lion chase and catch the gazelle. But we are not in fact seeing the event at exactly the moment it occurs. All acts of perception must be received and interpreted by the perceiver's visual apparatus and neural cortex. It is near-instantaneous but it is not quite and some animals will perceive it that bit quicker (others slower).

The reports produced by surveying instruments are necessarily delayed. What we receive are images of what has happened rather than what is happening. Would that be reason enough to deny that it is seeing? On the basis of the previous arguments the answer must be no. An appeal to ordinary language, though never conclusive, provides further support. Certainly, the following interchange seems perfectly natural:

> A: Did you see the football match last night?
>
> B: Yes, I saw it this morning on the later showing.

Would it be appropriate here for A to respond 'So you didn't see the match'? It seems not. Thus, temporality is not central to perception. Then what is? I would argue one part of that answer would include the processing of information. This, after all, is what our eyes and brains do every waking moment. It is also what electron microscopes do when they receive input stimuli and process them to produce a meaningful image. And it is what gyroscopes do when they sense the rotation of the Earth and produce from this a measurement of their orientation.

To return to our issue of perception knowledge in logging analysis, my submission was that there are stable, non-stipulative, non-arbitrary grounds for distinguishing between perception in the 'ordinary' case and in the 'analysis' case. This is important from an epistemological standpoint as perception is generally taken to be a primary source of knowledge together with memory, innate ideas, introspection, testimony, perhaps, and so on. That is not to say that each of these are equally reliable—in fact, it is quite difficult to ascertain the reliability of logging analysis for reasons that will become clear—but that they are sufficient justifications in and of themselves for belief. In other words, a reliable perceptual system is an epistemic virtue and a perceiver can be held epistemically responsible for beliefs formed as a result of perceptual processes. This conclusion will have consequences for the second half of this section.

What is being perceived, however, does differ from ordinary cases of seeing bottles on tables, and so on. We can 'see' using film-based systems such as these but they do not record images, only values. There is a problem of how we can see subterranean and submarine events and entities and whether this is different from, for instance, seeing with a camera or seeing with a microscope. I call this the engineering problem of unobservables. The problem of unobservables usually relates to a problem in the philosophy of science about the ontological status of microscopic entities such as electrons or quarks or Higgs Bosons. Given that we cannot see these entities directly – as we see what Austin called 'medium-sized dry goods' – then can we be justified in saying that they do really exist in the same sense as fridges and persons exist or should we say that they are at best theoretical postulates?

Now, there is no question that the hydrocarbons and rock formations, and so on, do really exist. And yet we cannot make, with the same confidence, assertions about their nature as we do with regard to bottles on tables or the colour of the sea. The solutions we use to improve the reliability of our assertions, however, are the same: technological. With regard to microscopic entities, we use, naturally, microscopes and such like to see the entities (or at least, the traces they leave). With regard to subterranean entities, we use gyroscopes, sensors and other measuring tools. In logging analysis, for example, it would be possible to send cameras down to record and transmit real-time images of the surrounding rock formations or of the fluids themselves. But it would also be impractical, expensive and ineffective. Likewise, we used to examine physical cuttings from the rock formation recovered from the drilling process i.e. directly perceive the rock particles. However, this is now largely superseded by the more useful methods of logging analysis.

Let us split this technological history into three parts. We have in the first instance what we might call direct examples of perceptual surveying. These involved perhaps simply looking at samples of the rock formation or, slightly more obtusely, a physical wax form of the inclination drawn out of the well. Any knowledge derived from these methods is unarguably perceptual. Secondly, we have methods one step abstracted from the direct methods. I will call these inscription methods. They involve a tool which makes a physical mark (a stylus point, pen or photograph) referring to a state of the well. Thirdly and lastly, we have what I will call inferential methods. These involve digital data records which are then interpreted and inferred as referring to states of the well. These inferential methods are the latest in a progression towards more abstracted and more sophisticated techniques for acquiring knowledge about oil wells but, as I have argued, we can still analyse them as perceptually-sourced, if the agents involved possess sufficient perceptual and professional capacities.[83]

---

[83] The history as I have written it pertains directly to the oil and gas industry. However, the reader should be aware that these methods are put to use in other industries such as drilling underground pipelines, tapping geothermal resources, exploration of other kinds of mineral deposits (such as coal, phosphates, and potassium), drilling underground chambers for nuclear tests, and others. Whilst we are interested in how sociotechnical systems in the oil industry shaped the development of this technology we ought to keep in mind that it was not the only relevant industry and that other actors have played a significant role.

Suppose that you accept that the users of the Enactive Torch acquire a new way of perceiving or a new percept of a way of perceiving or extend an existing capacity. Compare this with the following three cases in petroleum engineering. I will present three putative sources of perceptual knowledge which I will call, for want of better terms in each case direct, inscribed, and inferred. These sources follow roughly chronological order of their introduction to the industry.

Category I Direct. We have already seen some direct forms of instrumental information in Ch. 5. In this chapter I discussed the early tools of borehole surveying and drilling such as the wax cylinders used in African diamond mines. Similarly, geological methods which examine rock cuttings are called here 'direct'. Both involve 'face-to-face' encounters with a physical (not a representational) product of the rock formation. I would include it in what I am calling direct instrumental information about the well inclination for reasons that I hope will become clear. There are no great philosophical quandaries here. A camera sent down to look at the fluid would also be direct information if it were practical.

Category II Inscribed. The inscribed source is exemplified by the Teledrift and TOTCO tools, developed as a significant advance over wax cylinders. These tools communicate their information through physical imprints in the environment or in their own system. The Teledrift tool communicates its results using mud-pulse telemetry, a process that exploits pressure differences to carry binary signals to a surface reader. A pendulum is connected to a plunger and inclination is derived from the number of restrictions through which the plunger passes. The TOTCO tool also uses a pendulum with a stylus attached which penetrates a paper disc marked with concentric rings indicating hole inclinations. As with Dretske's argument that tree rings carry a signal (bits of information) to a receiver – viz. the information that the tree is a certain number of years old – the markings on this disc carry bits of information about the wellbore inclination to the receiver at the surface. (Dretske 1983) The disc is recovered and the marks provide a physical representation of the wellbore inclination. Cameras were used in early wellbore navigation but they do not film or take pictures of the well but of a pendulum and compass card unit.

Category III Inferential. Inferential sources communicate their information not through direct perception of their states or through physical inscriptions but through encoded data. I would not be surprised if this was the most contentious example of a perceptual source. I will argue that this contentiousness derives primarily from a lack of familiarity with the perceptual capacity. That is, if we were more familiar with the capacity it would seem more like perception (see my earlier argument for the immediatacy of experienced perception).

6.2 Senses and sensors: The extension of perceptual capacities with technical artefacts

6.2.1 Seeing with microscopes vs. seeing with gyroscopes

Fortunately, a similar argument to the one I am making has been made by Elisabeth Pacherie. (Pacherie 1995) She argues that there are no stable criteria for distinguishing between the natural perceptual instruments in the skull and artefactual perceptual instruments such as those developed by micro-biologists, petroleum engineers, and so on (other than the arguably tenuous and stipulative distinction between the natural and the artificial.)[84] Here I argue that the images generated by PL tools are perceptual sources of information of the fluid in the well just as the images generated by electron microscopes are perceptual sources of information of electrons.

Pacherie refers to instruments used in generating knowledge as “epistemic artefacts” and compares their contribution to the one made by our naturally-endowed epistemic organs or systems. In particular, she focuses on microscopes and compares these with the visual organs and systems. She notes that all perception is indirect in that it involves acquired skills and learned knowledge beyond the innate abilities we are born with. Since both microscope and ordinary images “carry information about the spatial properties of distal layouts” there are no

---

[84] See Ch. 4 for discussion of this distinction.

stable, non-stipulative grounds for distinguishing between "seeing with eyes" and "seeing with microscopes". However, with practice comes the appearance of instantaneousness. I have argued that there is no difference either between 'seeing with microscopes' and 'seeing with gyroscopes' (so named by Laplace because they 'see' the rotation of the Earth) and we eliminate the one at the expense of the other. I and the engineers I worked with are not the only ones to make the allusion to inferential perception. Here is Oberto Serra, a geologist, in his textbook on well-log analysis:

> More recently, however, there has been an increasing appreciation of the value of log data as a source of more general geological information. Geologists have realized, in fact, that well-logs can be to the subsurface rock what the *eyes and geological instruments* are to the surface outcrop. (Serra 1984, p. 2. Emphasis added.)

And a quote from the BP Amoco handbook:

> Since both well and geological target are invisible from the surface, we can never really know whether the one passes through the other. We can only *infer* it from survey tool measurements and well results. (BP Amoco Upstream Technology Group 1999, p. 3)

In fact, it is not hard to find descriptions of inferential perception in what one would assume to contain descriptions purely of procedural instructions and data-processing.

We have so far discussed four non-biological putative instrumental sources of perceptual knowledge. That is, instruments which communicate signals which carry information that can provide perceptual knowledge: the Enactive Torch; direct instruments such as the wax cylinder and those used in geological tests; inscribed instruments such as the TOTCO and Teledrift tools; and inferential instruments such as the Gyro tool. Of the three classifications of tools in petroleum engineering, the Gyro most readily gives itself to perceptual descriptions of the information it carries. That is, engineers and analysts using the Gyro are more likely to say 'it just looks like…' than someone reading a compass card where the phrase 'looks like' refers to the state of affairs underground and not a measurement of the state of affairs. Why is this? I would suggest that it is because whether something seems like perception is a factor of two things: complexity of information and familiarity. When we are born our eyes are open but we lack the experience to tell things apart, to select from our environment what is relevant and ignore the rest, to become skilled perceivers. The situation is the same for any form of perception. Consequently, any skilled processing of spatial or distal layouts will not *seem* like perception if it is easier than the 'basic' perception that it supervenes on.

6.2.2 The Bayesian approach

What might an inferential theory of perception look like? In this section I will outline an approach to perception from cognitive science which provides an example. Let us say that a *state* (or a situation) is a set of variables in a process that generates sensory data or inputs. States, that is, cause sensory inputs. Typically those variables vary rapidly and continuously over time. Hence, states of the world change rapidly and continuously over time. However, some variables (in processes generating sensory inputs) change discretely and on a slower time scale. Sets of these discrete and slowly changing variables can be taken as *contexts*. The distinction between state and context is important, but for my argument it won't make a significant difference. So, for ease of discussion, if not otherwise specified, I shall use 'state' to refer to both states and contexts. The only access we have to the world is through our senses which can be viewed as sources of information about the states of the world. This information is generally corrupted by random fluctuations, noise, and ambiguity. The same sensory information can be caused by many different states and the same states may cause different types of sensory information. When we act in the world, moreover, our motor signals are also corrupted by noise. Since intelligent and adaptive behaviour is tied to the ability to survive in a changing and uncertain environment, our cognitive system must handle this sensory and motor uncertainty in order to extract information about which state obtains in the world. The Bayesian framework provides one principled way this sensory and motor uncertainty can be handled in order for us to behave adaptively in our world.

Bayesian inference is a type of statistical inference where data (or new information) are used to update the probability that a hypothesis is true. To say that a system performs Bayesian inference is to say that it updates the probability that a hypothesis H is true given some data D by executing Bayes' rule,

$$[1] \quad \mathrm{Prob}\ (H|D) = \mathrm{Prob}\ (D|H)\mathrm{Prob}\ (H)\ /\ \mathrm{Prob}\ (D)$$

In other words, 'the probability of the hypothesis given the data ($P(H|D)$) is the probability of the data given the hypothesis ($P(D|H)$) times the prior probability of the hypothesis ($P(H)$) divided by the probability of the data ($P(D)$).'

In the case of our cognitive system, hypotheses consist of states of the world, and data to sensory inputs. As our cognitive system receives sensory information, the probability distribution over the possible states of the world is updated via [1]. There are in fact accumulating pieces of evidence that *suggest* that our cognitive systems might implement Bayesian inference. (Doya et al 2007; Knill & Richards 1996; Rao, Olshausen, & Lewicki 2002) Even though the main type of evidence comes from psychophysical tasks where people's performance is shown to approximate the Bayesian optimum, broad features of biological sensory systems can be explained in a Bayesian framework. Prior knowledge about the causal structure of the environment would be encoded in the backward connections. Forward connections would provide feedback by transmitting sensory prediction-error up to higher levels. Perception would arise from mutually informed top-down and bottom-up transformations distributed along the hierarchy.

6.2.3 Conclusions on seeing (and knowing) with technical artefacts

> Notoriously, this inference is non-demonstrative: there is typically no conceptual connection between a perceptual category and its sensory indicants; an indefinite number of perceptual analyses will, in principle,

> by compatible with any given specification of sensory input. On this account, then, perceptual integrations are most plausibly viewed as species of inferences-to-the-best-explanation, the computational problem in perceptual integration being that of choosing the best hypothesis about the distal source of proximal stimulations. (Fodor 1975, p. 50)

The Bayesian method is an example of a physical description of knowledge attribution – the agent has perceptual knowledge when her sensory apparatus have performed the appropriate Bayesian inference on the informational state – which can accompany our social description. Hence we have a sociotechnical description, one that combines reference-talk about both the empirical and the social phenomena at play. As with technical artefacts, this is essential to fully characterizing the nature of technical knowledge. A purely social description would not be comprehensive, since it omits the relation between the social consensus and the physical processes in the epistemic system. It is not typically the case that knowledge attributions depend for the truth solely on other social conditions being present. The knowledge must be transmitted through a physical relation between the referent and the epistemic system. Likewise, a purely physical description would not be comprehensive for practical purposes since we need a method whereby we can determine whether this person or this group is right to attribute knowledge. In this thesis I have favoured a collectivist approach where knowledge attributions are correct by virtue of the social interactions within the epistemic system. Other approaches may also be proposed but, in any case, we need some way to characterize the social element of knowledge which gives that epistemology a practical application.

In the next chapter I pursue further the idea that epistemic systems can be analysed as information-processing systems. I argue that an informational approach provides a great deal of heuristic value when analysing epistemic systems and that as long as we give due consideration to the social element of knowledge attribution, we can have a comprehensive, practically useful account.

Chapter 7

Data interpretation and information

7.1 Philosophy of information

> Evans had the idea that there is a much cruder and more fundamental concept than that of knowledge on which philosophers have concentrated so much, namely the concept of information. Information is conveyed by perception, and retained by memory, though also transmitted by means of language. One needs to concentrate on that concept before one approaches that of knowledge in the proper sense. (Dummett 1993, p. 186)

> Social epistemology is the study of the social dimensions of knowledge or information. (Goldman 2006)

7.1.1 Information overload, the need for epistemic virtue in the Information Society

We can think of each instance of knowledge in informational terms.[85] This allows us to draw analogy between ordinary cases of perception

---

[85] Parts of this chapter are adapted from Kerr & Pritchard 2011.

and data analysis. Philosophers of information, according to Floridi, study how information should be "adequately created processed, managed, and used."[86] (Floridi 2010c, p. 32) One question which epistemologists routinely ask is what is the purpose or function of inquiry? (Kvanvig 1998; Owens 2003; Rorty 1995)

We have previously discussed how inquiry is closely linked to answering questions. Consider a simple case: If I visit the doctor to be told that I am suffering from Creutzfeldt-Jakob disease (CJD) I will probably be given rafts of information, in the form of leaflets, documents, verbal explanations and so on. Perhaps this will be a lot for me to take in at once. On a much grander scale, corporate offices are increasingly becoming enormous databases, collecting, disseminating, and processing untold amounts of information. One of the, perhaps unsurprising, discoveries of my fieldwork research was that the majority of analysis of well-log data was not done in-the-field, so to speak, but was conveyed back to large office computer networks and

---

[86] Informative summaries of the philosophy of information are available in Allo 2010; Allo et al. 2012; Floridi 2002, 2003, 2004, 2008a, 2010a, p. 16-44, and 2010b.

teams of analysts who then process and interpret this data. In 2003, researchers estimated that approximately 12 exabytes of data had been amassed throughout human history up until the advent of computers. They also calculated that print, film, magnetic, and optical storage media produced more than 5 exabytes of data in 2002 alone.[87] (Floridi 2010c, p. 11) In August 2008, a company that manages supply chains for retailers processed 100 gigabytes of information per day. By early 2010 this number had increased by a factor of ten. (Stenzel 2011, p. 172) In 1991, research has showed that the name of every US citizen is processed by computers forty times per day. (Shaw 1991, p. 35) This number is likely to be significantly larger today.

This vast growth of information without a concomitant growth in data management tools has been called a 'data deluge' and 'information overload'.[88] It is part of our daily negotiations that we must sift through

---

[87] 1 exabyte = 10^{18} bytes.

[88] For a philosophical description of information overload and its negative consequences see Himma 2007. For more on the data deluge see *The Economist*, 25th February 2010, 394/8671. London: 11.

the overload and select the most relevant information to us. Almost two centuries ago Johann Wolfgang von Goethe wrote,

> The modern world has a false sense of security because of the great mass of data at its disposal. But the valid criterion of distinction is rather the extent to which man knows how to perform and master the material at his command. (quoted in Friman 2003, p. 19)

Here Goethe is expressing the, now common-sensical, view that information (and, indeed knowledge) alone is quite useless if there is not a) the tools to organize and filter it; and b) the necessary time and resources to process it.[89]

---

[89] I say that it is 'common-sensical' that the superabundance of data is a problem if we do not have the ability and tools to use it, but this view has been challenged. It is, for example, quite orthodox within information and communications technology to state that 'the data outperforms the algorithm'. See, for example, Anand Rajaraman's study with his Stanford Data Mining students discussed at http://anand.typepad.com/datawocky/2008/03/more-data-usual.html and also at http://anand.typepad.com/datawocky/2008/04/more-data-beats.html. In other words, to state that for any given problem, the quality of the procedure you use to solve the problem is not as important as the amount of data you have at your disposal. One can think of many examples where sheer weight of data can provide statistically better solutions to a problem than a high-quality algorithm. The amount of data simply overwhelms and disables any imperfections that the algorithm might have. The question is debatable and important but cannot be solved in this dissertation. I will rest on the assumption that, however we decide the algorithm vs. data question, it will be important to have heuristics and tools for sorting and identifying high quality information.

We rarely attend to information for its own sake. The joint purpose of my doctor and I is that I will come to have at least some useful knowledge about the disease. I may be informed that CJD is a rare but fatal brain disorder; it affects about one person in every one million people per year worldwide; symptoms typically occur at about age 60; about 90% of patients die within a year; and so on. I want this information because it is relevant to the knowledge I wish to acquire about CJD. In other words, the information is epistemically valuable to me in the situation I am in.

What do we mean by information? It is a polymorphous word and is used idiomatically with increasing variety in phrases such as the Information Age; information overload; the Information Superhighway; freedom of information; information technology; information science; the Information Society; and so on. 'Information' and 'knowledge' appear together frequently both in popular writing and scientific disciplines either as conflated terms for the same phenomena or related terms in some way involved in practices of inquiry, discovery, knowledge acquisition, and so forth. Whole academic disciplines are devoted to the study and management of information and it is often tasked with responsibility for various cultural ills and developments.

One curious phenomenon surrounding the ubiquity of technology, software and algorithms in our everyday and professional lives is that the more complex this equipment becomes, the more control we surrender to it in return for what we perceive to be increased efficiency and utility. The assumption that computer software and the algorithms are nothing more than time-savers and powerful calculators is coming under increasing scrutiny as they outstretch our ability to read and understand what they do. This is not a Luddite proposal. Instead I propose in this chapter that we understand what it is that the algorithm is capable of by appraising it as an epistemic agent.

We can separate our analysis into four distinct concepts: data, algorithms, information, and knowledge:

Datum The representation of a concept or entity in a form suitable for communication, interpretation, or processing by human beings or by automated systems. (Zins 2007) This may be sensory stimuli that is perceived through human sensory organs. Or, alternatively, it may be coded symbols organized according to an algorithm, readable by an automated computational system.

Algorithm A procedure or finite list of instructions for calculating a function. (Rogers 1987) This may be, e.g., the Bayesian, or other, inferential procedures involved in sensory perception in humans or the heuristics, arithmetical, and computational procedures involved in, e.g., PLA.

Information A meaningful, well-formed instance of one or more data. Well-formed means that the data are clustered together correctly according to the rules or syntax that govern the system. Meaningful means that the data complies with the semantics of the system. (Floridi 2011a)

Knowledge What someone (or an automated system) knows. To know something requires that the epistemic agent a) has access to the

> information required under the chosen system and b) has been appropriately attributed with knowledge by the chosen epistemic community. (Kusch 2002)

Before detailing the implications of this quadripartite system, let us review the relatively small cadre of information-based epistemologists.[90]

7.1.2 Information-based epistemology

The quote at the head of this chapter repeats Dummett's assertion that the philosophy of information is analytically prior to epistemology. Information is, as it were, more basic than knowledge. He continues,

> Information is acquired, for example, without one's necessarily having a grasp of the proposition which embodies it; the flow of information operates at a much more basic level than the acquisition and transmission of knowledge. I think that this conception deserves to be explored. (Dummett 1993, p. 186)

---

[90] By information-based epistemology I mean an epistemology which takes as a central component the concept of information; one which holds that an epistemic agent has knowledge if and only if she possesses the requisite information.

Dummett is not the only prominent epistemologist who has sought to position information as a core concept of epistemology. Fred Dretske and Luciano Floridi have developed sophisticated epistemological accounts which speak of information, data, signals and receivers, transmission, and the like, rather than belief, concepts, ideas, speakers and hearers, communication, and the like. Consequently, their philosophy is appropriate for studying the inquiries of extended epistemic agents such as analyst-plus-PLT or autonomous computational systems. Dretske suggests we connect knowledge with an ordinary dictionary definition of 'information':

> [*By information*] I mean nothing very technical or abstract. In fact, I mean pretty much what (I think) we all mean in talking of some event, signal or structure carrying (or embodying) information about another state of affairs. A message (i.e., some event, stimulus or signal) carries information about *X* to the extent to which one could learn (come to know) something about *X* from the message. (Dretske 1983, p. 10)

Speaking of information in this way, Dretske aims to relate information-based epistemology to the relevant alternatives theory discussed in §3.2. Why does information overload pose a problem for epistemologists? Philosophers such as Fred Dretske and David Lewis have argued that in order to know, one must be in a position to rule out certain relevant alternatives. For example, in order to know that what you see in the garden is a goldfinch you must know that it is not a chaffinch. If you are unable to make this distinction, they argue, then it is not appropriate to say that you know as you may quite easily have been looking at a chaffinch and called that a goldfinch also. Similarly, we can imagine Goldman's 'fake barn' case. Goldman, borrowing from Carl Ginet, asks us to imagine driving through a countryside populated by many barn façades. That is, there are only the front walls of barns but they appear from the roadside to be ordinary barns. It is just that we cannot see what is behind the façade. There is however, one barn that is real. It just so happens that I point to this one and call it a barn.

In this case it seems odd to say that I knew it was a barn as I could just as easily have been pointing to any one of the fake barns surrounding it. I just got lucky. And knowledge, as is now platitudinous in epistemology, cannot be lucky. (Pritchard 2007) The rule of relevant alternatives does not require that we rule out possibilities that no reasonable person would entertain, of course. We do not have to rule out that the bird is a giraffe or that the barn is in fact the square root of three or the concept of pathos. Nevertheless, it seems appropriate that any knower worthy of the title be able to rule out obvious falsehoods.

This poses a problem in a world where there is so much relevant information that we often do not have the necessary expertise to rule much of it out. Even if we did, we do not have the time to be able to do so. Consequently, our hand has been forced. We need to rely on computers which are able to process this information and, crucially, make sense of it, in a satisfactory timeframe.

7.1.3 The reason for creating new epistemic agents: not extended but external

At the beginning of this chapter I described a scenario in which a patient may seek information as a means to gaining knowledge about a medical matter. If information such as this were always susceptible to skeptical challenges then this susceptibility would be uncomfortably passed on to the knowledge claims based upon the evidence it carries. Such worries caused Dretske to abandon a key principle explaining how we reliably expand our knowledge: epistemic closure. I have presented an alternative epistemological picture here which does not have such drastic consequences. However, this informational story cannot be the whole story for collectivists. Return to perceptual knowledge and consider how we discussed that perception can be viewed in information terms.[91]

---

[91] For an extended discussion of the goal of information collection and dissemination see Fallis (2002). Note that even those who deny that the goal of information services is for users to acquire knowledge grant that in a large range of contexts our goal in collection and disseminating information *is* to acquire knowledge. For example, the information management scholar Chun Wei Choo expresses, albeit in different terms, a widely held view when he states that the primary

The problem I deal with in this section is this: Today, algorithms are replacing humans as data miners and knowledge producers. This is a response to an increasingly complex world in which individual humans do not have the processing power to make sense of it or compete with the efficiency of a computer. But are algorithms able to cope with information overload in a reliable and epistemically virtuous way? The answer I submit is that we do not currently know for certain but that there are systems we could put in place that would keep them better in line with the traditional picture of a virtuous epistemic agent. This section outlines the ideas of information overload and the various industries which use algorithmic methods to deal with it. I will pose the problem that a computer lacks the human ability – necessary for epistemic virtue – to understand and reflect on the data it receives and the action it takes. Consequently, as things stand, unsupervised algorithms have the potential to make wild knowledge claims that will

---

goal of information management is to 'harness the information resources and information capabilities of the organization in order to enable the organization to learn and adapt to its changing environment.' (Choo 2002, xv) Later, Choo writes that the 'transfiguration of information into knowledge is the goal of information management.' (Choo 2002, xiv) go under the radar of the rest of us or be subsumed within a larger technical system. I will argue that an information-based epistemology would approach this problem by appraising the algorithm just as any other epistemic agent.

Consider one tool we have developed to combat the problem of information overload. If we take traits like curiosity, open-mindedness, and inventiveness to be epistemic virtues then Google's page rank is not conducive to a flourishing epistemic life. Instead of having my filter quieten unpopular information, I should be trying to receive better quality information as well as searching for lesser known but potentially more useful information. What has now become an instinctive solution to solve many problems that pop up during our everyday lives, the internet search engine seems notably ill-equipped to return high-quality information, at least in this stage of their development.

Let me paint a second picture. Things are not getting any simpler and so greater computer power and computer influence is the direction of travel. Specifically unsupervised autonomous algorithms are left to their own accord because, seemingly, it is more efficient. Take the financial industry, for example. Here, algorithms are increasingly popular because they provide speed and the capacity to dissect an entire day's news instantaneously. Where once stock traders shared a single crowded room and shouted and gestured to try and get their signal out fastest; algorithms now compete with each other over timeframes shorter than that of a finger clicking a computer mouse. In September 2011, plans went ahead for the Hibernian Express, a six thousand km fibre-optic link connecting the New York and London stock exchanges. This link will shave six milliseconds off the transmission time between these cities. Whilst this doesn't seem very worthwhile to many of us, for the algorithms developed by high frequency trading firms this is an eternity and gives the computers a competitive edge.[92]

---

[92] The £300m cable that will save traders milliseconds. *The Telegraph*, 11th September 2011.

Algorithms are replete throughout the financial world and not just for trading stocks. In fact high frequency trading accounts for 70% of US stock market transactions.[93] In 2010 a study showed that algorithms running on twitter can be used to predict the stock markets. They search for patterns of language use and from this attempt to derive users' moods. If moods are good, the markets are predicted to be up and conversely if moods are low, the algorithm predicts a market fall. In addition, some hedge funds are now using this software and, perhaps most significantly, it appears to work (i.e. their fund is up over a period when pretty much everyone else was down). Outside of finance, a Scottish company called Epagogix predicts whether movies will make money using algorithms which compute variables to do with the script. For example, they might say that adding a woman in peril or moving the location to a desert island will add $20 million. Again, this is said to be accurate to an unnerving and, I find, quite dispiriting degree. Amazon marketplace, which is quickly forcing out high-street competitors, uses algorithms to set competitive prices.

---

[93] The Guardian, September 16, 2012, p. 22. Ross, A. K. & Mathiason, N. 'Britain opposes MEPs seeking ban on high-frequency trading.'

In my own fieldwork research I have seen algorithms which can be used to predict oil well blowouts such as the one that took place on Deepwater Horizon and manage risk. There is a tendency in the oil industry to replace software which is relatively simple and input-driven with software which takes much of that control away from the user and performs its own calculations which are not always transparent and not always traceable. In my training I was given instruction in three brands of software that follow this trajectory: PLWin, PLATO, and Emeraude. PLWin was one of the earliest pieces of logging analysis software developed by a German petroleum engineer working in Aberdeen, Scotland. Today, its interface and feature suite immediately strikes one as old-fashioned. However, when compared to the more sophisticated French-based software, Emeraude, it gives the user a lot more control over the variables used in calculations and other manipulating data curves and the like.

The benefits of the modern approach to software development are that it provides a reliable general overview of the state of the well and requires much less in the way of technical proficiency and the training to interpret logs accurately oneself. The benefits of the earlier approach were that analysts who were experienced with the 'look' of well logs could more easily delve into the details of the data and make more accurate predictions and spot any significant outliers that the software may not. The results can sometimes be unpredictable and can go unnoticed either by those who are not interested in great accuracy (because they are only looking for a general overview) or who just do not understand what a reasonable prediction should look like. We will see the potential fall-out from some of these scenarios shortly.

In the following I will compare algorithms like these with a human analyst. They are after all doing a very similar job. Traditionally, all these tasks I have described have been performed by humans and we can compare the kind of knowledge gathering that would ordinarily have been done by humans with the new algoepistemology done by high frequency trading algorithms or Epagogix, the movie valuators. Our algorithmic colleagues are similar to us in several of the respects needed for them to compare as epistemic agents. They behave and act autonomously (i.e. without supervision or control by a human). They search for information and inquire and then act on that information, normally, in a rational manner. They make knowledge claims. Many epistemologists would object to me even considering an algorithm to be worthy of the term ‘epistemic agent’. One central objection they make is that algorithms, like thermostats, do not have belief states. In reply to this I will say two things:

1. It is arguable whether an epistemic agent needs to have belief states. If you adhere to the information-based epistemology of Fred Dretske or Luciano Floridi, for example, it is not required.

2. Even if you do want a belief-based epistemology, algorithms certainly approximate believers in some sense. They do inquire, gather, investigate, and are disposed to act in certain ways depending on their environment and inputs.

This is not so dissimilar to a human inquiring into the shape of a rock, investigating and being disposed to call it a sheep given certain physical stimuli. Most importantly, they act independently of direct coercion. They are, in other words, agents. Of course, they are constrained by their coding and can certainly be stopped and started at will by a left mouse-click but we might as well point out that humans are also constrained by their biological and sociological determinations and can be stopped and started or otherwise coerced by a pointing finger or stern word. A belief-based epistemology that required absolute free agency would be skeptically strict.

The algorithms work. Often they work not because they are very smart but simply by throwing a lot of information in. The old A.I. adage says 'data outperforms algorithms' or we might say 'information outperforms intelligence'. This is statistical fact for many A.I. researchers. Given their success in many industries we can imagine them spreading into other areas: We will use algorithms to assess the likelihood of epidemics and direct hospital resources accordingly; we will use them to predict voting patterns and our political climate will change; no one will make a film, book or work of art without first consulting the algoepistemologist. This all sounds quite optimistic when discussing hospital management or earthquakes but begins to sound quite dystopian when it moves into cultural terrain like art and politics. There are more worries to come. Earlier this year a book appeared on Amazon called 'The Making of a Fly'. It was priced at $23.7 million (plus $3.99 shipping). What had happened was that Amazon's pricing algorithms had fought to outbid each other and the price had gradually escalated until it became absurd. One can imagine what would happen if such algorithms become responsible for automated trading of goods or services. On a more serious note, in 2010 9% of the US stock market just disappeared in 5 minutes.

The confusion is perhaps most apparent in the stock markets. Nobody knows what happened to that 9% because no one authorized it and there's no way to trace it. Kevin Slavin described this at the 2011 TED conference in Edinburgh as 'writing the unreadable'. We have unwittingly coded and set free something which we do not even understand.[94] The algorithm will change the environment and sometimes we will notice that it has done something odd (like wipe 9%) but be impotent to do anything about it and other times we will just not notice.

This is exactly what happens in the petroleum industry with the various software used to analyse well data. As stated earlier, I looked at three pieces of software: PLWin, Emeraude and PLATO. PLWin, the earliest effort at creating software for logging analysis, requires quite a lot of training to use effectively but also allows engineers a reasonable amount of freedom in how they use the data. For example, one engineer might want to set a particular variable which describes the relationship between density and the ratio of oil to water or oil to gas. In another piece of software this variable would be predetermined. Developers generally do this to remove potentially fatal bugs and to protect their rights over the sophisticated algorithms the software makes use of. The result is that newer software potentially masks

---

[94] There is actually a 100-year old novel by E.M. Forster called The Machine Stops that forecasts this very eventuality.

unusual or outlying data that would ordinarily have been removed in a manual process and judgment.

7.1.4 What can computational systems know?

Computational systems do not believe. This simple fact seems to provide a knock-down argument for saying that they cannot *know*. After all, it seems nonsensical for a person to say that they know *p* but they do not believe it.[95] Consequently, an epistemology which wishes to incorporate computational systems as epistemic agents must forego the belief-condition for knowledge. That is not to say that when discussing the knowledge of human agents we can attribute them with knowledge even though they do not believe. It is to say, however, that some agents (e.g. computational systems) can have knowledge without belief.

---

[95] This is known as a Moorean sentence after the first philosopher to formally acknowledge it. (Moore 1993)

The problem for algorithms is a problem of common sense. They cannot recognize an absurd variable of the kind that human operators would instantly spot. Take the Amazon pricing algorithm that kept outbidding other algorithms until the price became absurd. Any human store clerk or data inputter would have stopped such a bidding war very early on. The software behaves very much like a naïve inquirer or some trainee analysts, receiving and processing information without reflecting on whether the information makes sense in the context.

The problem is related to the frame problem in computer science. It appears to be impossible to write a programme that can find whether solutions to quite simple equations are available as the computer will run through infinite iterations without noticing what any competent mathematician could notice just be looking at the formula. One of the most evocative illustrations was provided by Dan Dennett who described three robots with progressively incompetent solutions to a relatively (for us) simple problem. In the experiment, the robots are tasked with removing a precious battery which was in a room strapped to a wagon with a time bomb set to go off in the near future. The first robot, R1, found the room, and the key to the door, and hypothesized that in order to remove the battery it must pull the battery out of the room. It succeeded in pulling the battery out of the room but not long after, of course, the bomb exploded. A second robot, R1D1 was built which would be designed to recognize 'not just the intended implications of its acts, but also the implications about their side-effects, by deducing these implications from the descriptions it uses in formulating its plans.' (Dennett 1984, p. 1) This robot was perhaps even less successful. R1D1 arrived at a similar plan to R1 and as it was working through the various implications, such as that pulling the wagon would not change the colour of the room's walls, the wagon exploded. The third robot, R2D1, was designed with something like a Dretske-Lewisian epistemology in mind. Dennett writes, 'We must teach it the difference between relevant implications and irrelevant implications,' said the designers, 'and teach it to ignore the irrelevant ones.' (Dennett 1984, p. 1) This robot spent some time running through thousands of implications which it then deemed to be irrelevant, keeping a list of irrelevant implications as it went, before the bomb went off. These are robots which suffer from the frame problem, a problem which AI research has to date failed to resolve.

The problem is not, of course, exclusive to algorithms (hence drawing further comparisons between human and artificial agents). In my own fieldwork I have observed the importance for analysts to know the practical meaning of the data they are using—and not just the procedure of entering it. Any experienced engineer familiar with what the data represent would see a value of, let's say, 2.74 and realize that cannot be right and there must be an error or false data earlier on. A novice who does not know what the calculations and information mean will not notice this. In such cases, we would say she has become robotic, merely following procedural rules and not stopping to reflect on the likelihood of that data being correct. Needless to say that I think this is a genuine problem as incoming trainees are trained more on how to use increasingly sophisticated software and less on the back-of-an-envelope calculations and practical work that give that ability to spontaneously error-check.

The solution I propose comes from the philosophy of information, particularly that dealing with knowledge. I will follow a Dretske-Floridi line, as I think that is best suited for this kind of comparison. These two authors contend that knowledge is information-based. That is, that information is a basic component of knowledge. Recall Dretske's definition of information given in §7.1.2. Keep in mind that there is little in this definition that would preclude an algorithm as being the receiver of such a message. In fact talk of signals and receivers seems more suited to a computer than a person. By relating information to knowledge in this way, Dretske's information-based epistemology becomes allied to the relevant alternatives theory that he also endorses. You will recall that this view states that possessing knowledge depends on an agent's capacity to rule out a certain range of alternatives. This notion of relevancy has been notoriously difficult to pin down, as noted by Floridi (2010a), Schope (2002), *inter alia*. Duncan Pritchard defines the position as it applies to perceptual knowledge as follows:

> 'S has perceptual knowledge that p only if S can discriminate the target object at issue in p from the objects at issue in relevant alternative (not-p) propositions, where a relevant alternative is an alternative that obtains in a near-by possible world. (Pritchard 2010a, p. 3)

According to this rendering of the relevant alternatives view, our capacity to possess perceptual knowledge is heavily affected by our environment.

7.1.5 Friendly and unfriendly cognitive environments

Pritchard makes the distinction between epistemically friendly and unfriendly environments. Most of the time, it will be very easy for us to make the necessary discriminations between, for example, hands and stubs, or canaries and crows. But epistemologists are always coming up with weird thought experiments which arrange our environment such that it will not be so easy, for instance in the fake barns example. It is because of the possibility of deceptive environments like this that Dretske denies that information alone could ever answer a skeptical doubt.

The argument for this is as follows: I have many defeasible reasons for thinking that I am writing this in Edinburgh, Scotland (memory, testimony, observation, etc.). This gives me an informational basis for believing that I am writing in Edinburgh. However, I do not have an informational basis for believing that I am not a brain-in-a-vat (BIV) on Alpha Centauri who is being fed the illusion that he is writing in Edinburgh. Even if the standards for knowledge are very low, and even if I know that were I in Edinburgh then I would not be a BIV on Alpha Centauri, this would not give me an informational basis for denying the skeptical hypothesis. The reason for this is my inability to discriminate between the scenario in which I am in Edinburgh and the skeptical BIV scenario in which I am on Alpha Centauri. Accordingly, argues Dretske, it follows that I receive exactly the same information in either scenario, and hence that I can have no informational basis to reject the alternative skeptical scenario.

In general, Dretske argues that no signal can carry the information that a skeptical hypothesis—an hypothesis explicitly designed such that it is indiscriminable from normal circumstances, and yet involves a high degree of error—is false. In his *Knowledge and the Flow of Information*, for example, Dretske writes,

> No signal can rule out all possibilities if possibilities are identified with what is consistently imaginable. No signal, for instance, can eliminate the possibility that it was generated, not by the normal means, but by some freak cosmic accident, by a deceptive demon, or by supernatural intervention. (Dretske 1981a, p. 130)

And, more recently,

> This is true of all indicators, all sources of information. That is why there is nothing in the world…that indicates that there is a material world. (Dretske 2005b, p. 22)

So on Dretske's view I can have an informational basis for believing that I am in Delft but I can have no informational basis for believing that I am not a BIV on Alpha Centauri (a skeptical hypothesis which entails that I am not in Delft), even whilst I know that if I am a BIV on Alpha Centauri then I am not in Delft. It is for this reason that Dretske denies epistemic closure. You will recall that epistemic closure is the principle that if an agent knows one proposition, and knows that it entails a second proposition, then that agent also knows the second proposition. So, for example, if one knows that one is presently in Delft, and one knows that this entails that one is not a BIV on Alpha Centauri, then one knows that one is not a BIV on Alpha Centauri. Although this principle has broad intuitive support, Dretske rejects it.

But why is it that on Dretske's view I can acquire knowledge about a proposition but not about a proposition which I know full well is entailed by it? Dretske is led into this position through two closely related commitments: (i) that perceptual information is never relevant to skeptical hypotheses, and (ii) that information is essentially non-factive evidence.

7.1.6 Unfriendly environments, skepticism and algotrading

Let us transpose this idea of an unfriendly environment into the world of algotrading. Let us say that an unfriendly or skeptical environment is one that would deceive the algorithm just as it would deceive a human trader. The consequence is that the algorithm is misled into either declaring falsehoods or neglecting to sufficiently protect itself against being misled in the future. This would be epistemically vicious. The solution would be to somehow programme in the algorithm the ability to detect an unfriendly environment. In some cases this will be easy. Consider how a human detects an unfriendly environment. Take the barn façade case, for example. How would I check that I am not in that scenario? Fortunately, I have many ways. One way would be to inspect the barn more closely; walk round the back of it; or use a pair of binoculars; and so on. Alternatively, before I took the trip I could ask my friend if he knew of any fake barn countryside which we might be travelling through. I could watch the news for reports of fake barn areas around the country and estimate the likelihood that I will pass through one.

In a previous example I talked about how a novice programmer might not recognize that the value of ‘2.74’ she entered in a programme is absurd because she does not understand what it means. An expert would not make that mistake because she would recognize the data as standing out. Essentially, what we are doing when we do this kind of investigation is induction: comparing our current experience with past ones and checking that nothing is out of the ordinary. The algorithm could be programmed to do the same thing. A trading algorithm could check with previous transactions of a similar type and restrict actions which did anything too unorthodox. Of course, this would be useless in practice as the extra milliseconds it took to do that would lose it the sale and further the very purpose of these algorithms is to break previous behaviours. The problems occur when you or I or the algorithm is in such an unfriendly environment that it cannot discriminate between the good cases in the past and the present bad case.

Since, ex hypothesi, agents cannot discriminate between normal scenarios and skeptical alternatives, so it follows, according to Dretske, that agents lack an informational basis for dismissing skeptical alternatives. Another way out of this would be if information was factive. That is, some guarantee that a clue in the environment implied some relevant truth. Dretske's second commitment (that information is non-factive) becomes clear once we reflect that if information could be factive evidence for what it is evidence for—if, that is, it could entail the truth of what it is evidence for—then it would follow that the information we have to support our beliefs in normal circumstances might well suffice to entail the denial of the target skeptical scenario. Clearly, however, Dretske does not think that we ever have evidence of this sort, and hence a non-factive view of the evidence provided by information is clearly implicit here.

In order to more closely examine these commitments, consider the following local skeptical hypothesis:

> **Zebra$_i$**
>
> Fred is at the zoo. If he perceives what he takes to be a zebra, Fred can have no informational basis for believing that what he perceives is not, in fact, a cleverly-disguised mule. In other words, the signal carrying this information does not allow him to discriminate between 'a zebra in my perceptual field' and 'a cleverly disguised mule in my perceptual field'.

Fred may interpret the signal as evidence that there is a zebra in front of him as a matter of habit, or perhaps relying on other evidence such as the sign on the fence or assumptions about what kinds of animals are in a zoo. However, his information is, it seems, non-factive. Just because I receive a signal such as this does not entail that there is in fact a zebra in the pen. More generally, as Dretske claims, it appears that none of the information that the subject possesses which indicates that he is perceiving a zebra is information which offers him an adequate epistemic basis on which he can dismiss the 'cleverly disguised mule' skeptical scenario.

This way of thinking about our evidential position with regard to skeptical challenges has, however, been challenged. Ram Neta (Neta 2002; 2003), for example, has argued that the scope of your evidence is affected by context. Under this account, there is a range of contexts in which evidence (for which read: information) is factive. Neta argues that the skeptic only appears to succeed by restricting what counts as evidence. In normal contexts my evidence typically is factive, and it only becomes non-factive in skeptical contexts in which very demanding standards for what counts as evidence are in play.

Hence, in the zebra case, my evidence for believing that there is a zebra before me could well be factive in normal contexts. For example, if my evidential state in normal contexts is that of seeing that there is a zebra before me, then, since seeing that p entails p, my evidential state actually entails that there is a zebra before me, and which hence entails that I am not currently being presented with a cleverly disguised mule. Relatedly, if my evidence, in normal contexts, for believing that I have two hands is that I can see them before me, then I have evidence which entails not only that I have two hands, but also that I'm not a handless BIV on Alpha Centauri.

According to Neta, however, the context can change in such a way as to restrict the scope of one's evidence. If I were to gain evidence that cast doubt upon my belief that I have hands—for example, if I were to witness a room of BIVs—then this would make the possibility that I am a BIV a relevant alternative. This is effectively what the skeptic does: to describe such a scenario and cast doubt upon what was previously undoubted. There are two ways in which this may be done. Either the skeptic may simply suggest the possibility of a skeptical hypothesis that had previously been ignored or unexamined by the subject. This may place an onus on the subject to now eliminate that possibility in order to be correctly said to know the proposition. This intuition suggests that we cannot know a proposition until we have ruled out all relevant alternatives and that the range of relevant alternatives is determined by the conversational context. (Pritchard 2010a, p. 19) In other words, being made aware of an alternative, however implausible or absurd, can make that alternative relevant.

The second way in which the skeptic can make the alternative relevant is by actually offering evidence for thinking that a skeptical scenario has obtained. For example, consider an extension to the case of Zebra$_i$:

> **Zebra$_{i2}$**
>
> Fred's friend and skeptic, Frank, mentions to Fred that he once read a science-fiction story in which all the world's zebras are replaced by hologram zebras and the real zebras are taken to a neighbouring planet. A little while later, Frank notices a pot of paint lying beside the animal and brings this to Fred's attention by gesturing towards it. He also tells Fred that the sign on the outside of the pen appears to have been written over an older sign, suggesting that a different message was once written there.

In this example, Frank initially merely presents Fred with a radical skeptical hypothesis. In the view of some epistemologists, such pronouncements can change the conversational context in which evidence requirements and relevant alternatives are set. Frank's story may rob Fred of his knowledge that there is a zebra in the pen before him. In the subsequent details of the story, however, Frank presents Fred with perceptual information and testimonial evidence for calling into doubt Fred's knowledge of what is in the pen.

According to Neta, in these skeptical contexts Fred's evidence is no longer factive. In particular, it is now no longer the case that one's evidence can entail the denials of skeptical hypotheses, given that they are in play and problematising our epistemic position. So although in normal contexts my evidence that I am seeing two hands could be that I see that I have two hands, in skeptical contexts where the skeptical hypothesis is at issue my evidence can at most be that I seem to see that I have two hands, where this evidential standing clearly does not entail the target proposition. It is, on the other hand, possible to gain evidence supporting local skeptical claims. In the case of Zebra, if I were to notice a pot of paint next to the animal or its flaking 'skin', then this may provide an evidential or informational basis for believing that the animal is a cleverly disguised mule. If one subscribes to Dretske's relevant alternatives theory or Neta's contextualism, then the absence of such signals means, respectively, that either we are not required to rule out this possibility or that we are in an ordinary context in which the denials of skeptical hypotheses are known.

Consider O. K. Bouwsma's 'adventures'. (Bouwsma 1965) In the first of these, a demon has constructed a world to confuse a man called Tom. In this world everything is made out of paper, although it looks identical. However, the paper world does not deceive Tom for long. In a horrific telling of the story, when Tom peels away part of his face he receives, as Dretske would say, a signal carrying the information that he is in a world made of paper (i.e., that a skeptical hypothesis—viz., that the world he perceives is not what 'real'—is true). Of course, one could take this a stage further and ask if the perception of a paper world is also the victim of a skeptical trick but there the same test will apply. Whilst Tom is in the paper environment he has the capacity to discriminate and can come to know. Information, in these local skeptical scenarios, is relevant to what Tom knows. In Zebra, it appears, we have perceived signals that carry the information that the animal may be a painted mule.

What is relevant information is constrained by skeptical or non-skeptical environments. Just as the victim of CJD does not need to know about the controversies over the aetiology of CJD (because he is a sufferer not a specialist doctor), he does not need to know the denials of skeptical hypotheses which may cast doubt upon what knowledge he possesses about CJD (because he is an epistemic agent and not an epistemologist). The upshot of this is that information not only has the function of providing a basis for knowledge but also a relevant alternatives or context-defining function. This gives pluralist epistemologies such as relevant alternatives theory and contextualism practical application as epistemic sorting-machines for information managers: in what contexts can we know what we want to know, what information is relevant, what information changes the contexts for knowledge, what are the epistemic limits of information?

To return to our doctor's scenario that began this chapter, the knowledgeable specialist is one who can inform me of relevant information about CJD and also point me in the direction of reliable information sources elsewhere (and steer me away from dodgy websites and unsafe medical treatments). In most cases, these sources will not be denials of skeptical hypotheses but they will be sources of information which will increase the likelihood of my acquiring knowledge about my condition and how to cope with it. It would be an odd special case if local skepticism were the only epistemological problem that can be affected by informational signals in a context. Information services such as libraries, databases and internet search engines can also make use of relevant alternatives in order to organize and structure their resources and content.

Here are two apparent truisms.

1. Our interest as inquirers in information is often motivated by our desire to gain knowledge about something.

2. We are almost always faced with limited information about the target issue. At the very least, one can always think that it would be better if one had more information about this subject matter.

What falls out of these two statements? One might think that, as Aristotle claimed of knowledge (De Anima, 402a1), more information is always better than less and so we should endeavor to collect as much information as possible on the matter in question with the hope of, at some point, turning it into knowledge.[96] Cursory reflection reveals that this is evidently false. (Himma 2007) Internet search engines are a good example. Type in a random search string and it will probably return hundreds of thousands of results. No human could sort through that amount of information and so the search engine is designed to return those results that are likely to be most beneficial to the user first. Thus, in order to deal with problems as they arise one needs to put constraints on what evidence and information is relevant. According to Neta, the skeptic unduly restricts evidence in certain contexts. What information management effectively does is make the same judgments about appropriate restrictions.

---

[96] Another example of this thinking can be found in Bishop Berkeley's *Alciphron or the minute philosopher*, where Berkeley's spokesman for traditional Christianity, Euphranor, states: I love information upon all subjects that come in my way, and especially upon those that are most important. (Berkeley 1803, p. 25)

Dretske's account is primarily an account of perceptual knowledge and information. He therefore feels entitled to conclude that, since the mere appearance of an object cannot communicate its non-skeptical status, any signal which carries information about appearance cannot answer a skeptical doubt. However, we have provided examples (such as Bouwsma's adventures and Zebra) where perceptual information does justify a skeptical hypothesis or a non-skeptical proposition. It would seem that Dretske is wrong to think that information is irrelevant to combat local skeptical scenarios. Agents can receive information (even if we think of information as non-factive) for dismissing such scenarios (once we do not limit their information to the bare visual scene). (Pritchard 2010a) Whether Dretske is right about radical skeptical scenarios depends on whether information is ever factive. If it is always factive then Dretske has no need to deny closure. Even if information is only sometimes factive (i.e., in ordinary contexts, à la Neta) then Dretske is still wrong.

Let us consider an argument that reasons (under which heading we may include perceptual evidence) are factive, which is from John

McDowell. (McDowell 1998) Earlier in the paper, I discussed Neta's comment that external world skepticism is not meant to cast doubt upon certain 'inner' reasons such as 'that I am not having a visual experience of a white expanse before me'. McDowell argues against a tacit assumption throughout epistemology that these inner reflections can encompass factive empirical reasons. (Pritchard 2008a, p. 10, 2009)

However, McDowell does not think that no empirical reasons are factive. In the case of veridical perception, we have a kind of perceptual evidence which is not present in cases of non-veridical perception such as illusion or hallucination. McDowellian epistemological disjunctivism presents an option for Dretske which has so far been left unexplored but which may undermine his case against epistemic closure, with concomitant implications for his theory of information. In brief, if perceptual evidence is (sometimes) factive, then Dretske is wrong to say that there is no perceptual evidence which can serve as evidence against skeptical hypotheses. Dretske's view is that all perceptual evidence is defeasible when it comes to radical skeptical hypotheses. No matter how competently one receives and judges the information one is presented with, these processes never amount to something which entails the denial of the target skeptical hypothesis. The view is intuitive and persuasive but the McDowellian view offers one alternative: that there is a disjunct between cases of factive and non-factive reasons. That is, there is some reason or warrant or a kind of support missing in cases of radical skepticism that is present in so-called 'ordinary' cases.

Dretske takes it for granted that any given knowledge claim can be subject to a skeptical rebuttal. Such rebuttals challenge the upgrading of an information-based belief (that something appears to be the case) to information-based knowledge (knowledge that something is the case). In the case of Zebra_i there is information that carries the signal to Fred that what is in the pen is a painted mule. Dretske might insist that this does not undermine his thesis as these pieces of information may themselves be subject to skeptical hypotheses and are providing only non-factive evidence. However, if one follows McDowell down his disjunctivist path then it is not inevitable that Dretske takes such a position and consequently not inevitable that he is led to reject the principle of epistemic closure.

Neta presents a contextualist account of evidence or reasons in which the evidential requirements for knowledge are affected by context. Dretske closely links information to non-factive evidence but under the contextualist account there are cases of factive evidence which would provide information-based knowledge of the denials of skeptical hypotheses in some cases. Additionally, McDowell provides a non-contextualist account of evidence or reasons in which there is an epistemic component present in some cases, not present in others (such as cases of hallucination or illusion—the hallmark of skeptical hypothesizing), and in which factive evidence warrants the denial of skeptical hypotheses.

As a consequence, these distinctions between skeptical and ordinary contexts or between factive and non-factive evidence present alternatives to Dretske's inference that perceptual information can never give us evidence or reasons to refute skeptical hypotheses. I described a scenario in which a patient may seek information as a means to gaining knowledge about a medical matter. If information such as this were always susceptible to skeptical challenges then this susceptibility would be uncomfortably passed on to the knowledge claims based upon the evidence it carries. Such worries caused Dretske to abandon a key principle explaining how we reliably expand our knowledge: epistemic closure. I have presented an alternative epistemological picture here which does not have such drastic consequences.

How can this help algoepistemology? If many of our epistemic endeavours are to be replaced by algorithms can these be protected against the misleading setting of an unfriendly environment? Maybe. There are informational clues that give away unfriendly epistemic environments. It is just that the algorithms prefer to go for fast computation over this slow monitoring activity. What kind of programme would be able to spot an unfriendly environment? Let's go back to Zebra. In that case what we receive is physical stimuli that cause to pause and reflect. That kind of reflection is inherently unavailable to algorithms. What is needed will likely be some kind of human-like judgment even if this is restricted just to spotting unfriendly epistemic environments which are prone to leading the algorithm astray. It may be that entire fields are highly susceptible to generating unfriendliness.

To sum up, an algorithm fails in a specific epistemic virtue. It does not – cannot – reflect on its conclusions or its inputs with the kind of rounded social context that a human operator can. It cannot think, '$27million is patently absurd for any book and so something must have gone wrong,' or 'The plane has been climbing too high, something must have gone wrong.' Nor does it understand the figures it deals in. They lack the reflective capacity necessary to check for absurd results. Consequently, they will always have problems identifying unfriendly environments. We can programme them to manage some risks as and when they can check previous instances with the current action. They will not be able to identify indiscriminable unfriendly environments in the way that humans can because they lack that reflective capacity needed to solve problems like the frame problem.

In conclusion, algorithms can be epistemic agents but not particularly virtuous ones and so we do, as epistemologists, have reason for concern at their widespread use even whilst they give us more capacity to deal with the complex world we have made. What would an epistemically virtuous algorithm look like? It would be, at the least, one that reflected on unusual cases by comparison with previous cases. But more than this it would be one that searched for unpopular and, on the face of it, unpromising information for the sake of curiosity. It would therefore be a much more open code than is previously used. This would be less efficient but also more able to spot unfriendly environments and be less disposed to the kind of absurd results we have seen to date.

7.2 Information, meaning and knowledge

> "What is it in the territory that gets onto the map?" We know the territory does not get onto the map. That is the central point about which we here are all agreed. Now, if the territory were uniform, nothing would get onto the map except its boundaries, which are the points at which it ceases to be uniform against some larger matrix. What gets onto the map, in fact, is difference, be it a difference in altitude, a difference in vegetation, a difference in population structure, difference in surface, or whatever. Differences are the things that get onto a map…. Of this infinitude, we select a very limited number which become information. In fact, what we mean by information – the elementary unit of information – is a difference which makes a difference. (Bateson 1972, 320–321)

In earlier chapters I argued that in order to understand epistemic concepts and, in turn, to understand knowledge attributions, we must understand their meaning and how they are used.[97] We mean many things by "information". In the quote that heads this chapter, the anthropologist and cyberneticist Gregory Bateson defines information as 'a difference which makes a difference.' This might make information sound like almost anything. To some extent this is true. We are used to hearing the word "information" used in so many

---

[97] Parts of this chapter are adapted from Kerr (2012).

contexts from freedom of information to the information age to the information superhighway to information overload and so on. This chapter looks at the relation between information and meaning: both how information gets its meaning – in other words, what makes information meaningful – and how we can explain meaning using the conceptual tools of information theory. In previous chapters we have asked "What is information?" or "What does "information" mean?" The question that concerns us in this chapter is whether information itself is meaningful.

7.2.1 The theory of meaning and the symbol grounding problem

The theory of meaning is an enormous topic in the philosophy of language. In its most general form, it is the attempt to explain how our language connects to the external world; that is, the relation between what we say and what we are speaking about. The issue is not as straightforward as it sometimes appears at first glance. At the start of the 20th century, Ludwig Wittgenstein wrote: 'The difficulty of my theory of logical portrayal was that of finding a connection between the signs on paper and a situation outside in the world. I always said that truth is a relation between the proposition and the situation, but could never pick out such a relation.' (Wittgenstein 1961, 19e-20e)

Philosophers of language at the end of the 19th century had noticed that the thing that a word refers to is not necessarily the same as the meaning of the word. Frege distinguished two aspects of meaning which he called Sinn (or "sense") and Bedeutung (usually translated as "reference"). (Frege 1960) Consider the example of the planet Venus. Venus was once known by two names: those who saw it at sunrise knew it as "the Morning Star," and those who saw it at sunset called it "the Evening Star." These two names express a difference sense, but they have the same reference, viz. the planet Venus. This reference-based theory of meaning was followed by numerous attempts to find the "hook" that connects our words to the things they are about. Philosophers of information claim that the problem of meaning is actually much wider. Theories that explain meaning have been limited by their focus on language, specifically on speech and writing. Philosophers of information try to explain meaning by referring to a much more general phenomenon than language, which is, of course, information. This gives rise to two central questions: (i) how information acquires its meaning (and consequently, to the question of what meaning is) and (ii) how the theory of meaning can be illuminated by informational concepts.

Let us take each of these in turn. The first question has been called the "meaning grounding problem" and it concerns where meaning comes from or how an entity can acquire meaning. We recognize some marks, symbols, sounds, gestures, signs, and so on as meaningful yet others are meaningless. But just knowing the relationship between symbols or sounds does not ever amount to knowing the meaning of those symbols and sounds. Consider a subspecies of the meaning grounding problem which concerns how symbols or data acquire meaning. This is known as the symbol grounding problem. It can be illustrated with a familiar example. Suppose that you have just arrived in a foreign country with no knowledge of the local language and you are attempting to find your way to your hotel. You notice a sign with some letters printed on it, you hear people speaking and gesticulating, you see traffic directions and symbols. All of this is more or less unintelligible to you. You buy a dictionary but it is entirely written in the local language. Nevertheless, you persevere and look up the word on the sign. You then begin to look up the words in the definition; and then the words in that definition, and so on. Of course, you will probably never be able to derive any meaning from this dictionary as the mere relations between symbols cannot provide you with the "hook" that connects those symbols to the world which they seem to refer to.

Stevan Harnad described the problem as:

> How can the semantic interpretation of a formal symbol system be made intrinsic to the system, rather than just parasitic on the meanings in our heads? How can the meanings of the meaningless symbol tokens, manipulated solely on the basis of their (arbitrary) shapes, be grounded in anything but other meaningless symbols? (Harnad 1990)

The problem he describes is, put somewhat crudely, how can we get something from nothing? If we find meaning in the symbols – for example, the particular colour pattern before our eyes at the moment – then are we not just imposing the meaning that we already have in our minds onto the environment? If so, where did the meaning in our heads come from in the first place? We seem to run into a regress: a circular repetitive sequence which only prompts further questions. A "symbolist" (for example, Jerry Fodor 1980, 1985) would say that the meaning of the symbols is grounded in the relationship between the symbol system and the world. Many symbolists believe that cognition can be reduced to a kind of symbol manipulation. If correct, this is promising as we can theoretically simulate cognition in artificial agents by equipping them with the right rules for manipulating symbols. In humans, this symbol manipulation is subconscious and autonomous and once the brain receives symbols then, assuming it is functioning properly, it will acquire the relevant meanings. Research in cognitive science, however, highlights the interpretative and problematic nature of picking out the objects, events, and states of affairs that symbols refer to and so the symbolist account seems to simplify the function of cognition.

So how ought we to resolve the symbol grounding problem? One option is to rely on perception. We can say that when we learn a language, for example, we learn how to appropriately map symbols to specific referents. Suppose that, in the illustration given above, you abandon your dictionary and try to work out what the people around you are saying. You listen to their words, observe what they point to, and construct a small bilingual dictionary of your own. Each time someone points to an object and appears to refer to it you write down what they say and translate it into your own language. Unfortunately, you sometimes hear someone say a word and seem to refer to a completely different object or situation. Again, no matter how close and detailed your observations it seems that you will never completely understand what the meaning of the word is just by noting the structural relations between words, symbols and gestures.

What is more, whilst symbols are usually discrete, self-contained entities, the external world for the most part is not. If we go along with this approach we are either sorely lacking in symbols if we want to describe the world well or we conflate many things when using the same symbols. In any case, as described by Hilary Putnam, just as a set of mathematical equations can be mapped onto a near-infinite number of structures in the world, a system of relations between symbols can be mapped onto a near-infinite number of structures in the world. Hence, there is no non-arbitrary way in which symbols acquire meaning (Putnam 1999. See also Lakoff 1987).

A second approach is to abandon our assumption that there is a direct relation between symbols and referents. An alternative approach might be found in research in embodied cognition, which takes cognition to be highly dependent on the physical capacities and actions of an agent. Specifically, we can view meaning as a way of coordinating action to achieve certain goals. According to this approach, the meaning of a particular situation for the agent is the combination of actions available to it. For example, the meaning of being close to a car, for a human, might be travel but for a cat it might be shelter. Consequently, the symbol grounding problem is dissolved rather than solved. There is no need to look for the “hooks” that connect references and referents, speech and what we speak about. Rather, what ought to be matched is goals and affordances: what the agent wants to achieve and what opportunities are afforded by a situation to achieve that goal.

7.2.2 Statistical approaches: Shannon and Weaver

Claude Shannon and Warren Weaver were mathematicians often credited with inventing information theory. Their landmark paper, A mathematical theory of communication, published in 1948, proposed an influential model of the communication of data and signals. Instead of focusing on the meaning (in the sense of interpretation or reference) of data, they constrained their model to the level of detail and frequency of the data. In other words, the content of the symbols is irrelevant; it does not matter what the information is about. All that matters is the quantifiable qualities of the data. Their model defines what have become the main components of many subsequent mathematical or probabilistic approaches to information, such as Bar-Hillel and Carnap, and Dretske discussed below. These components are the sender, message, transmission, noise, channel, reception, and receiver. By way of explanation, consider how a telephone system works. We have a person, A – the sender – speaking to another person, B – the receiver. The message is what A, the sender, says. The signal is the sound waves produced by the sender when she speaks, which are then carried through the electrical system (the transmitter) along a channel to B, the receiver. We can represent this as follows,

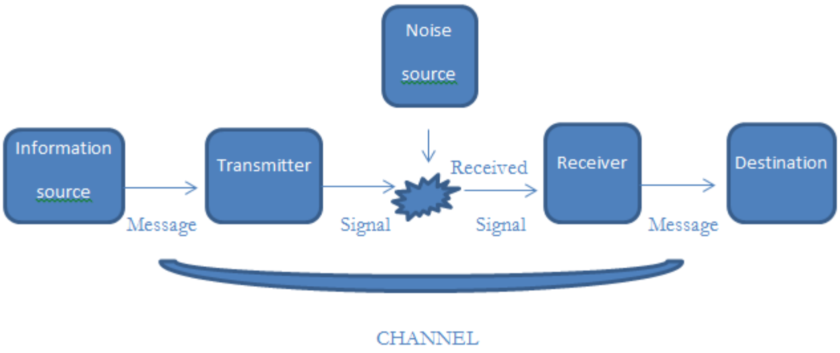

Figure 8. The Shannon-Weaver model of communication

This model has been applied not just to describe the external processes between two cognitive agents but also the internal processes of the mind. George Sperling, for example, described how environmental sounds are processed through cognitive systems and transformed until they reach the conscious level. One failing of the Shannon-Weaver model is its inadequacy for describing the transmission of information between multiagent or distributed systems; for example, when accurately describing how two people communicate or the communication between distributed computational systems where the line between who is the sender and who is the receiver becomes blurred. Today, most researchers recognize the benefits the mathematical theory of communication affords in terms of a useful probabilistic theory of the correlations between states of a sender and a receiver whilst acknowledging its limitations for many forms of communication.

7.2.3 Probabilistic approaches: Bar-Hillel and Carnap

Shannon and Weaver defined information in terms of probability space distribution. Yehoshua Bar-Hillel and Rudolf Carnap developed a related probabilistic approach which sought to do justice to the problem of meaning, which Shannon and Weaver had deliberately set aside. Their approach was based on what is called the inverse relationship principle. According to this principle, the amount of information associated with a proposition is inversely proportional to the probability associated with that proposition. The core idea is that the semantic content of p is measured as the complement of the a priori probability of p,

$$\mathrm{CONT}(p) = 1 - P(p)$$

Where CONT is the semantic content of p (p could be a set of sentences, events, situations or possible worlds). Crudely, CONT(p) is a measure of the probability of p not happening, or not being true. This means that the less probable or possible p is, the more semantic information p is assumed to be carrying. Tautologies, like "all ravens are ravens" have to be true. So they are assumed to carry no information at all. Since the probability that all ravens are ravens is 1, P(p) is 1, so CONT(p) is 1-1, i.e. 0. By extension, we might presume that contradictions – statements which describe impossible states or whose probability is 0, such as “Alice is not Alice” – to contain the highest amount of semantic information. Thus we seem to run into what has been called the Bar-Hillel-Carnap paradox: the less likely a statement is, the greater its informational content, until you reach a certain point at which, presumably, the statement contains no information at all since it is false. As Bar-Hillel and Carnap state:

> It might perhaps, at first, seem strange that a self-contradictory sentence, hence one which no ideal receiver would accept, is regarded as carrying with it the most inclusive information. It should, however, be emphasized that semantic information is here not meant as implying truth. A false sentence which happens to say much is thereby highly informative in our sense. Whether the information it carries is true or false, scientifically valuable or not, and so forth, does not concern us. A self-contradictory sentence asserts too much; it is too informative to be true (1953, p. 229).

7.2.4 Probabilistic approaches: Dretske

Fred Dretske was one of the earliest philosophers to formally connect information to theory of meaning. In a theory he called Indicator

Semantics he wrestled with the problem of how data and information can be "upgraded" to knowledge. Dretske supposed that the relationship between data and meaning was central to resolving this puzzle and key to this relationship was the further connection between signal and receiver. The relationship between signal and receiver is familiar from many scenarios: in a telephone call we receive a signal; a television antenna transmits a signal to an aerial; and so on.

This relationship is analogous to how our minds receive signals (or stimuli) from the world or how we communicate with one another. It is a bit like the relationship that pertains when two mountaineers climb a mountain. Often a length of rope will connect the two climbers so that the follower, Bob, can signal to the leader, Alice, whether he is climbing or not. If Bob wants to climb, he will tug on the rope, thereby sending a signal to Alice to give him more slack in the rope. This relationship only works if there is a tight connection between the signal that is sent and the beliefs formed by the receiver. Dretske suggested that if the probability that a certain state of the world (e.g. that Bob is tugging on the rope) given a certain mental state (e.g. that Alice believes that Bob needs more slack) is 1, then the mental state (the belief) has that world state as its content. Alice's belief, in other words, that Bob wants more slack only contains the information that Bob is tugging on the rope if and only if her belief is a reliable indicator that Bob is tugging on the rope. In this case, her belief is a reliable indicator since it would take a freak occurrence for her to believe that Bob is tugging on the rope if he were not. Philosophers call such beliefs 'sensitive'. Only sensitive beliefs, therefore, are informative.

In *Knowledge and the flow of information*, Dretske uses the notion of information to explain and develop the concepts of knowledge, perception and meaning. (Dretske 1981) He uses the Mathematical Theory of Communication (MTC) described above but considers it to have several limitations. MTC is a theory of data quantification and transmission rather than a theory of information, as we would normally describe it in ordinary language. It is this latter account that Dretske wants to provide. Secondly, MTC is an account of the statistical properties of transmission and the average amount of information generated by a source. Under Dretske's ordinary-language understanding of information, however, information is associated with singular events. As a solution to these limitations Dretske proposes the following definition of a signal's informational content.

> A signal r carries the information that s is F = The conditional probability of s's being F, given r (and k), is 1 (but, given k alone, less than 1) Where k is a variable that represents how what an agent already knows can determine the information carried (for that agent) by a signal.

For example, Fred's utterance that he is a professor does not carry the information that he is a professor of philosophy, but if you already know that Fred's discipline is philosophy then for you, his utterance does carry the information that he is a professor of philosophy. This is similar to the variable I outlined in my description of Bayesian inference that contains relevant background knowledge when receiving and processing states.

Dretske defends the claim that the value of the conditional probability in his definition of information is one and not less than one on three grounds. The first is that we ought to satisfy, Dretske claims, a basic principle of information flow known as the conjunction principle. This principle states that if a signal A carries the information that B, and A carries the information that C, then it also carries the information that B and C. If the conditional probability requirement was less than one, then this principle would not be satisfied. Secondly, Dretske wants his account to be consistent with what he calls the Xerox principle. This is the principle that if A carries the information that B, and B carries the information that C, then A carries the information that C. So information flow is transitive. If the conditional probability requirement was less than one, then again this principle would not be satisfied. Thirdly, if the conditional probability requirement was less than one we would be tasked with setting some kind of threshold on it. There appears to be no non-arbitrary way to do this and, more importantly, were it to be done then the epistemological significance of information would be dramatically reduced.

One problem with setting the conditional probability requirement to one is that since there are very few conditional probabilities of one, the amount of information flow is low. To combat this problem, Dretske proposes the idea of fixed channel conditions and relevant alternatives. The latter I have previously discussed in §3.2. In brief, Dretske states that conditional probability requirements are made relevant to a set of possible relevant alternatives to the communication channel. The problem with this, as discussed earlier, is that there is no obvious method for deciding what counts as a relevant alternative or for excluding what does not count.

7.2.5 Semantic approaches: Floridi, levels of abstraction, and syntax

The fourth approach we will look at is known as the General Definition of Information (GDI). (Floridi 2005) This theory proposes a tripartite definition of information:

> The General Definition of Information (GDI)
>
> p is an instance of information, understood as semantic content, if and only if:
>
> (GDI.1) p consists of one or more data;
>
> (GDI.2) the data in p are well-formed;
>
> (GDI.3) the well-formed data in p are meaningful.

In other words, if what you have is not a well-formed combination of meaningful data, then you do not have information. Here, "well-formed" means that the data are correctly structured according to the relevant syntax. The question that falls out of this statement – viz. when is data correctly structured? – is problematic enough. We have already discussed the importance of a relevant syntax. The question raised by (GDI.3) is the primary concern of this chapter – namely, when is well-formed data meaningful? In other words, how can data acquire their meaning? No doubt, we often attribute meanings to data.

We can point to a datum on a graph and infer that the datum means that, say, the recorded temperature at a certain pressure is 50°c. We can read a weather report and infer that the data in the report mean that it is likely to rain tomorrow. But what are needed are general principles which would be exportable across conceptual platforms. What allows us to make these inferences is the structure of the data and the syntax that we use to interpret it.

Here, syntax means something broad. It is the system, structure, code or language that determines the form, construction or composition of information. It is not necessarily linguistic – think, for example, of an instruction diagram for building flat-pack furniture, or Morse code signals. Bateson, in the quote at the head of this chapter, notes that on a map only differences in altitude, and not consistencies in altitude, would be represented (i.e. as contours). But if we do not understand the syntax of the contour lines – the differences – we may misunderstand what they represent.

Consider the following illustration from Dretske: An airplane pilot uses altimeters to determine the plane's altitude. An altimeter is essentially a pressure gauge – it responds to changes in air pressure and the pilot takes this to represent changes in altitude. Note that we must say "takes this to represent" here as what is crucial is not just what the device measures but what those measurements are taken to represent. Conflating the two can lead to bad inferences being drawn from the altimeter's movements. For example, if the aeroplane strays into an environment that has an unusual air pressure, and the pilot takes its movements to represent altitude, he will acquire false beliefs about the altitude. Similarly, if we take the altimeter out of the plane and place it in a depressurized container, the altimeter will register a much higher altitude than is, in fact, the case. In such examples, it is not that the altimeter is malfunctioning – it is doing its job perfectly well – but that we have incorrectly understood the syntax of the information it provides. We do not understand what this information means. When the altimeter responds to external changes in its environment it is not actually generating any information per se. What is generated might instead be called data.

We need a syntactical structure – what philosophy of information calls a Level of Abstraction – to give that data meaning and determine its informational content. We often hear people say, and have no doubt said ourselves, things like, "It looks cloudy. Does that mean it will rain?" "What does irascible mean?" "That's the third time a black cat has crossed my path. What does it mean?" and so on. What we are saying when we use the word "mean" is given by the kind of answers that we expect to satisfy the question. The answer that clouds are often used in literature to symbolise the coming of rain would not be a helpful answer to the first question but that the presence of certain types of cloud indicates that it is likely to rain would. The answer that it means that an a feline animal with dark coloured fur has travelled in such a manner as to move adjacently to your own trajectory would not be a helpful answer to the last question whereas the answer that black cats crossing your path give you luck might be satisfying (to the superstitious amongst us). We might conclude from this that the meaning of a bit of information depends on our interests or goals.

7.2.6 Pragmatic approaches

We receive a lot of stimuli or signals in our day-to-day lives – either through newspapers, television, and the internet, or just the stimuli and signals we receive from all the objects we pass as we walk down the road. Our minds need a way of managing all that data and selecting only the meaningful, informative, and – hopefully -- true. One of the first steps in this process is deciphering what this putative information means. Charles Sanders Peirce was an American philosopher in the later 19th century who co-founded the Pragmatist movement with William James and John Dewey, but who also tackled some of the problems we face in deciphering the meaning of various signals. Peirce thought that this problem of the meaning of signals, what he called Semiotics, was in fact something of a prima philosophia. He wrote, '[I]t has never been in my power to study anything, – mathematics, ethics, metaphysics, gravitation, thermodynamics, optics, chemistry, comparative anatomy, astronomy, psychology, phonetics, economics, the history of science, whist, men and women, wine, metrology, except as a study of semiotic.' (Peirce 1977, 85–86) Clearly, Peirce thought semiotics could help with quite a range of studies!

Peirce argued that in order to understand the meaning of a signal, we are not concerned with every aspect of the signal but only the signifying element. For example, in order to understand that smoke means fire, we do not need to know everything about the smoke – its shape, the particular way its fumes form, its precise colour, and so on – but only that element of it that signifies the presence of fire. Peirce used a lot of different terms for the signifying element of a signal – "sign", "representamen", "representation", and "ground". We will just call them signs. Consider a beehive in your garden as a sign that there are bees in your garden. It is not every single characteristic of that beehive that signifies that there are bees in your garden. The colour, size, or shape of the hive is not particularly important and plays what Peirce calls a "secondary signifying role". The primary signifying role in this case is the causal connection between the type of object that a beehive is and the presence of bees. This relationship is the sign. The meaning of this sign is "there are bees here!" There may be other signs that there are bees in the garden: the pollen count, a bee sting on a child that unfortunately bothered one of the garden's residents, and the noise that bees make perhaps.

What makes these things signs is their capacity to indicate the presence of bees. The colour of the stripes on a bee, its gender or age, are not essential to indicate the presence of the bee and so are not signs that there are bees. A second element to this connection is, in Dretske's terms, that receiver's interpretation of the sign. Roughly, this is the meaning we take from the relationship between the sign (the beehive) and the object (the bees). Peirce thought that signs determine their interpretation. That is, the beehive draws our attention to the connection between beehives and bees and in so doing determines that we will believe that there are bees.

When we talk about meaning in this communicational form, we often imply two concepts: intention and understanding. When Alice tells Bob that she is currently in Edinburgh, Scotland, she intends to transmit the information contained in what she says. She also intends that Bob is her audience and she assumes or predicts that he will understand what she says. Furthermore, in order for Bob to receive this information he must understand what Alice meant, at least to some degree. If he understands by her sentence something radically different from what Alice intended for him to understand then there has been a "failure to communicate": a failure, in other words, to transmit meaningful information. This is an intuitive picture of human communication.

Of course, the reality is that this picture is largely impressionistic. Humans can communicate effectively – that is, can bring about through communication whatever effects they wish – with surprisingly little common understanding. We can communicate with those who speak other languages or with children and have our intentions understood. Even with our linguistic peers there is often a significant disparity between what one person said and what their interlocutor took them to have said, and nevertheless, whatever that person wished to have happen was not completely blocked by this "indeterminacy" of meaning.

Very little of this is true for computational communication, such as might take place between an internet server and a web browser or between you and “Siri”, Apple’s “natural language” application for its mobile operating system. Thirty years on from the publication of Searle’s “Chinese room” argument and 60 years from Turing’s Imitation Game test for computer intelligence, we have become used to computers talking to us and even conversing with them, to some degree. They ask us questions such as, “Would you like to listen to your messages?” or, “Do you want to save changes?” and when we answer they respond appropriately. Conversely, we sometimes ask them questions such as when we interact with Apple’s “Siri” or when we execute a database query. It is possible for computers to take part in this kind of interaction without any understanding of what they are saying or doing. They are like the subject in the Chinese room passing on lists of symbols she cannot make sense of. This does not mean that it would not one day be possible for an intelligent machine to understand conversations (in the deepest sense of the word) but that it is not necessary for them to take part effectively and appropriately in a wide range of situations.

Computers are, for want of a more precise definition, information processors. That is, data goes in one end, and information comes out the other. Much of the last 60 years of research in Artificial Intelligence has wondered whether humans too are information processors in much the same way. Our brains and varied sensory apparatus pick up data from the environment, process it, and generate information. Some theorists in cognitive science speculate that our visual behaviour can be predicted or even simulated through Bayesian calculations, essentially the claim that the brain is a giant, superfast, probability calculator, trying to choose the most mathematically probable interpretation of any given data set. The metaphor may be tenuous. It seems unlikely, after all, that the human brain could be literally equivalent to an artificial machine. Certainly, the brain has been compared to many artificial machines throughout the history of ideas, from the catapult in Ancient Greece, to a mill in the 17th century, to the telegraph, and hydraulic and electro-magnetic systems in the early 20th century. (Searle 1984) We might conclude that saying the brain is a computer is just the latest in an ancient tradition of saying the brain is whatever technological marvel happened to be nearest to hand. We can refine the question in this chapter to whether computers can mean things in the same way that humans do. If they can mean what they say/display then we might be more tempted to say that they are functioning like human brains (and vice versa if they cannot).

So far we have looked at a number of approaches which try to answer the question of how data acquire meaning. We looked at statistical approaches such as Shannon and Weaver's mathematical theory of communication; probabilistic approaches such as Bar-Hillel and Carnap's or Dretske's theory of meaning; semantic approaches such as Floridi's; and Peirce's pragmatic approach. We also looked at several specific problems in the philosophy of information such as the symbol grounding problem and the Bar-Hillel-Carnap paradox. Finally, we asked what it would mean for non-animal agents to communicate or transmit meaning.

7.3 Agency, algorithms, analysts and the north-seeking gyro

7.3.1 Noise, ambiguity and the North-seeking gyro

Skilled analysts (like skilled infants) are performing analogous inferences on the information they perceive and coming to acquire perceptual beliefs. It is easy to see the similarities beyond mere phenomenology. The only access analysts have to their world is also through the senses (and sensors) which provide information about the states of the well. This information is also generally corrupted by random fluctuations, noise and ambiguity. The same log can be interpreted differently by different analysts. Intelligent and adaptive analysis is tied to the ability to reduce uncertainty by corroborating different sources of information and drawing on prior experience. None of this is to say that noise and ambiguity are altogether negative influences. In fact they are necessary parts of how the system hangs together. A system which could not recognize and deal effectively with noisy signals would be extremely maladaptive. What is more, the history of the technology used in this field has many examples where noise turned out to be of great benefit, once it was recognized for what it was: not messy nonsense, but evidence of another variable.

Recall how I argued that the philosophical and epistemological issues with which this essay is concerned all stem from one practical constraint: we cannot (ordinarily) see what is going on down there in a wellbore, the hole and its contents being located up to 10,000 or 15,000ft under the ground or sea. I add the qualification 'ordinarily' because I hope to demonstrate that the use of these instruments does in fact make possible a kind of seeing, perception enhanced by the use of compasses, cameras, gyroscopes, and other 'perceptual' equipment (note that the word 'gyroscope' derives from the Greek *skopeein*, meaning 'to see' and *gyros* meaning 'turn'; compasses are familiar devices for orienting a person or vehicle and so their connection to location and direction is plain; cameras are probably, of these three instruments, the most credible tool for perception-enhancement and will be discussed in some detail). The second point to make is that while some of the knowledge derived from the reports of these instruments can be accessed by other means, much cannot and the reliability and accuracy of most is greatly enhanced by their use.

Early surveying instruments would produce complicated data reports or 'logs' of various measurements which engineers would have to read and interpret into some meaningful propositional form. Given that no instrument is ever perfect many of these logs would be quite messy or 'noisy' and not obviously indicate any one particular interpretation. Mud-pulse telemetry is a case in point: The circulating mud system has a noisy pressure wave from the pumps and rig systems of filtering hydraulics do not always reveal a clear pulse. Any air in the 'mud' (typically a composite of clay, chalk and other minerals) causes it to compress and contract and so one cannot easily see the pressure spike on the graph (if you can imagine, the graph produced is a wave with peaks and plateaus and the imprecision of the results makes it difficult to read where the wave is peaking). Similarly gyroscopes used in surveying must be corrected for 'drifting'—the inevitable loss of a gyro's initial orientation over time—as well as other imprecisions and differences depending on the well's geographical location.

Analysts and engineers have a suite of techniques for correcting for these noisy interference patterns. However, just as we correct for what we perceive to be 'noise' or 'junk' we may also be ignoring what, if only we knew more, could be useful, valuable information.[98] In the case at hand, it was only later that engineers realized that they were junking the gyro's measurement of the natural rotation of the Earth! Recall that the Earth is in fact an enormous gyro, a spinning mass trying to hold a particular orientation in space established around 4.5 billion years ago. This natural process was later exploited in the 1980s by Don van Steenwyk, the engineer who developed the North-seeking gyro. This gyroscope is capable of finding True North without having to be primed at a known orientation.

---

[98] The history of science and technology is replete with examples of technologists purportedly 'correcting for' or treating as 'junk' information which may, at a later date, be sought after. Think, for example, of contemporary scientific attitudes towards so-called 'junk DNA'.

The 'junk' seen by early surveyors was a function of the rotation of the Earth. Suppose we accept my earlier argument that what we are dealing with is perceptually-sourced knowledge. In that case, what did the engineers at this time know of the data they took to refer to junk, noise, messy signal, imperfection in the transmitter, and so on? It is not that the data actually referred to the rotation of the Earth, for it was good and useful procedure for an analyst to correct for this corruption. What the analysts and engineers knew was indexed to a particular framework of interpretation. This framework changed when some ingenious engineers noticed that Earth rotation would have an effect on the tool. Thus, the conditions for knowing what the data referred to and correctly interpreting the log had changed.

Nor is it that the standards or thresholds for knowing changed. It was an interpretive framework that changed and this is best accounted for, I would argue, by the Kuschian, communitarian, collectivist analysis I have presented throughout this thesis. In many fields of inquiry—archaeology, forensic science, here in petroleum engineering—entities leave traces of information that are waiting to be read by interested observers who notice that what they are seeing is not meaningless. Since their status as relevant evidence is truly up in the air we cannot be asked to rule them out (and hence prove our knowledge) before this status is confirmed. Just as our own perceptual cognitive systems must adapt to handle noisy or ambiguous signals, any engineering science which is perception-like or whose instruments function as extensions to perceptual capacities must also have algorithms for handling noisy or ambiguous signals.

7.3.2 Return to Deepwater Horizon: Information and prediction

Over the course of a week in April 2010, two events caused widespread destruction to our environment and daily affairs, including tragically in the latter case, 11 deaths. The first involved hundreds of thousands of carbon dioxide emissions per day, the second involved tens of thousands of barrels of crude oil. Both were (in practice) not predicted. The first was the eruption of the Icelandic volcano, Eyjafjallajökull. The second was the blowout and subsequent leak at the Deepwater Horizon mobile offshore drilling unit (MODU). Both events attracted great public and political interest and prompted questions to be asked of authorities, scientists and engineers tasked with predicting and, where possible, preventing or mitigating their impact. The 2010 eruption cancelled thousands of flights throughout Europe—the most significant restriction in living memory—and cost hundreds of millions of pounds to business and the tourism industry in the UK.

Whilst that event was significant, the effects of the Deepwater Horizon blowout (DHB) were much more devastating. 11 men were killed in the explosion, from Texas, Louisiana and Mississippi. 17 others were injured. In total, approximately 4.9 million barrels of petroleum escaped in the three months before the well was capped. The consequences to marine and wildlife habitats, local fishing and tourism are by now known internationally. A reasonable question to ask is, why did we not know these events were going to happen? Volcanic eruption forecasting is far from being a perfect science. Scientists and engineers can, at best, predict some volcanic eruptions by spotting early indicators just before an imminent eruption. As regards DHB, allegations were made that BP, Halliburton and others could have known—indeed did know—about the blowout well before it actually occurred.

How reliable is the science and engineering for predicting blowouts in the oil and gas industry? In February 2009, BP stated that it was "unlikely that an accidental surface or subsurface oil spill would occur from the proposed activities." Without a detailed and protracted analysis of the internal machinations of the companies involved we may never conclusively know the answer to this but we can engage in an analysis of the issues of collective responsibility and group knowledge which arise out of this case as well as with the kinds of perceptually-sourced beliefs that come into play when using technology to observe subterranean or submarine events. Social epistemologists have a role to play in diagnosing how it is we can hold companies or other groups responsible for lack of knowledge and what kind of knowledge can be had of properties and events in an offshore oil rig.

What kind of knowledge were newspaper headlines like those listed above inquiring about? (§1.1.1) They ask about knowledge pertaining to the integrity of the well cement casing; whether it was known that the well was likely to suffer from a blowout (i.e. that the well's pressure control systems would fail); they ask who was responsible for knowing the risks and who was responsible for ensuring adequate safety procedures were in place. Now, it may be the case that these questions expect too much from the drilling engineers and safety officers. It would seem that there are many positions a person may hold in contemporary society where they are expected to be able to predict events which are beyond the predictive capacities of current scientific and engineering methods. For now let us consider what an agent might be responsible for. One cannot, presumably, hold the individual significantly responsible for what is largely creditable to the efforts of a group. However, the group cannot entirely be responsible for the kind of knowledge we have been concerned with. In part, this knowledge is generated through a process of informational interpretation and prediction. Let us consider this aspect now.

7.3.3 Capping the well and recapping the hypothesis

First, let us ask what can be known of the properties and events in an oil rig and then we will ask who can know it. I have already presented how the relevant knowledge is perceptually sourced and then put to communal arbitration. To recap, the fieldwork research involved working for a petroleum engineering company, specifically in the areas of well logging analysis and wellbore navigation. This is the analysis of data gathered from oil and gas wells in order to work out if the well is sufficiently productive, to diagnose problems (such as the risk of a blowout), and to direct drilling or interception of a well, and so on. One of the striking findings of working with other analysts was that it became apparent that the better qualified someone was to read and analyse these logs—the more experience and expertise they had—the more likely they were to describe their interpretations in perceptual terms. In truth, those who were at the earlier stages of their career tended to follow quite formulaic processes in deciding, for instance, whether to remove an outlying data point or emphasize a particular value. At other times their modus operandi appeared, to a novice and a non-engineer, to border on the arbitrary. Skilled and experienced analysts on the other hand were likely to describe what they were doing in casual perceptual terms:

> "It looks like zone 6 is thieving water" – Analyst, SP
>
> "[pointing] Can you see that the temperature is rising here? Because of oil." – Analyst, MT
>
> "Looking at that... It doesn't look right." – Engineer, BK

Or to simply pass over the fact that we were interpreting a data log, as in the following dialogue:

> Eric Kerr: How do I know I should cut it [the on-screen representation of the PLT data in Emeraude] here?
> Analyst, SP: You see where the spinner goes up here? That's where there's a perforation.'

By way of explanation, I should say that ‘spinner’ refers to a particular tool inserted into the PLT and the velocity it is rotating at. Therefore, it sounds odd to say that one ‘sees’ that the spinner goes up, when one is looking at a graphical representation of data received from the PLT. I do not wish to place too much significance on these conversations as they can be interpreted as metaphorical usage. Nevertheless, we have all had, I believe, experiences where we are reading a novel, inputting data on a spread-sheet or, if we are computer programmers, writing a line of code for a piece of software, where we transcend the experience of reading information and see what the information refers to. This is indicative, I suggest, of the somewhat unromantic fact that what we are doing when we are seeing is processing information (albeit often so quickly as to give the impression of seeing entire discreet objects).

7.3.4 Conclusions on data interpretation and information

In Ch. 6 I remarked that whilst we readily accept, on the basis of biological evidence, that perceptual knowledge (for example, seeing that an object is blue) is constructed, we are reluctant to accept the same about communal knowledge (knowledge that gains its status, in part, by virtue of shared consensus). And yet, in some sense, there is little principled distinction other than the contingent one that eyes are located in the head whilst other people are located elsewhere in the world. In this chapter I have explored one way in which cognitive science can expand on this approach. An informational approach to epistemology allows us to explore epistemic systems and evaluate their knowledge attributions.

I also reviewed the history of meaning in information to diagnose the importance of the meaning of concepts and their use to understanding the practices that make use of them.

Chapter 8

CAT, knowledge, and groups

8.1 Groups and perception

8.1.1 Introduction

Consider the following headlines from print and online newspapers from the months following the blowout.

> BP, Halliburton Knew Oil Spill Cement Unstable[99]
> BP and Halliburton knew of Gulf oil well cement flaws[100]
> Did BP know more than they led on?[101]
> Studies suggest MMS knew blowout preventers had ‘critical’ flaws’[102]

---

[99] BP, Halliburton knew oil spill cement unstable. *Discovery News*, 29th October 2010. http://news.discovery.com/earth/bp-halliburton-oil-spill.html.

[100] Goldenberg, S. & Kollewe, J. 2010. BP and Halliburton knew of Gulf oil well cement flaws. *The Guardian*, 29th October 2010. http://www.guardian.co.uk/environment/2010/oct/29/bp-oil-spill-bp.

[101] Wilson, J. 2010. Did BP know more than they led on? *KETK NBC*, 30th October 2010. http://www.ketknbc.com/news/did-bp-know-more-than-they-led-on.

[102] Clayton, M. 2010. Studies suggest MMS knew blowout preventers had ‘critical’ flaws. *The Christian Science Monitor,* 17th June 2010. http://www.csmonitor.com/USA/2010/0617/Studies-suggest-MMS-knew-blowout-preventers-had-critical-flaws.

> What did BP know?[103]
> BP wasted a lot of time in responding to the Gulf of Mexico oil disaster last year with remedies it knew wouldn't work, the chief executive of Exxon Mobil said.[104]
> BP Oil Spill: 7 Secrets BP Doesn't Want You to Know[105]

Cases such as the Deepwater Horizon incident are particularly interesting in this respect as it seems to be common to ask questions about the epistemic responsibility of a group, rather than the more typical case of an individual. These ordinary ways of speaking may be intended to indicate the social character of the knowledge involved. The idea that BP should or could know *anything* about what happened is philosophically puzzling. When we first think of what kind of thing can have knowledge many of us will tend to think of biological entities – usually a human being. That is, we think of something which has a mind, beliefs, will, agency, a brain, perceptual faculties, perhaps, and so

---

[103] Lugarten, A. & Knutson, R. 2010. What did BP know? *The Columbus Dispatch*, 9th June 2010. http://www.dispatch.com/content/stories/national_world/2010/06/09/what-did-bp-know.html.

[104] Exxon: BP wasted time in Gulf of Mexico. *United Press International*, 20th April 2011. http://www.upi.com/Business_News/Energy-Resources/2011/04/20/Exxon-BP-wasted-time-in-Gulf-of-Mexico/UPI-31641303300666/.

[105] Emami, G. 2010. BP oil spill: 7 secrets BP doesn't want you to know. *Huffington Post*, 7th May 2010. http://www.huffingtonpost.com/2010/05/05/7-secrets-bp-doesnt-want_n_563102.html.

on. Legal entities such as corporations do not possess any of these things. Put this way, it sounds very odd to ask whether BP knew anything at all. And yet it is a common feature of ordinary language use. (Goldman 2004, p. 12; Hakli 2007, p. 249; Quinton 1975/1976, p. 17; Schmitt 1994a, pp. 257-8)

Here, knowledge is attributed not to individuals per se, as is traditional in epistemology, but to groups, collectives, corporations, and so on. 'What did they know and when did they know it' became something of a mantra for various groups asking questions of BP, Halliburton, MMS, Goldman Sachs, and others. Why do we use this language, and what do we mean by it? Coincidentally, Lackey also uses an example from the Deepwater Horizon disaster. She points out that when discussing matters of such import in a serious setting such as a courtroom or when discussing serious matters such as the moral and legal responsibility of certain parties, no one could read the group knowledge attribution as,

> …merely metaphorical or as an exaggeration, as in the case when one says, The injured mouse just sat there as I approached him—it was as if he knew that I was going to help him,' or 'The wind blew the stick precisely to mark our location—it seemed to know that we needed to be rescued.' (Lackey, forthcoming a, p. 4)

In these cases there is no suggestion that the mouse or the wind might have really known the propositions in question. When compared to these common ordinary language cases the proposition in favour of group knowledge attributions seems more plausible.

An obvious objection to the strong view that groups can possess knowledge is that knowledge requires belief and groups do not have belief. Across a range of philosophical disciplines, the possibility that groups may indeed have beliefs or even minds has been considered by a number of researchers. (Hutchins 1995a, 1995b; Giere 2002; Gilbert 2004; Meijers 2003, 2003; Pettit 2003; Tollefsen 2002, 2007, 2009; Tumollini & Castelfranchi) In this chapter I will be taking a different tack. Namely, I will forego the condition that knowledge requires belief. This is not to deny the obvious point that if I know that snow is white I must also believe that snow is white—of course, it sounds ridiculous to claim to know something whilst also claiming that one does not believe it—but it is to assert the methodological principle that the study of knowledge does not require a concomitant study of belief. Since the beliefs of agents are often quite opaque to the outside observer (and occasionally to the introspector) an epistemological methodology that does not require any investigation into the beliefs held by particular agents may prove more fruitful. The communitarian and informational approaches outlined in this thesis provide such an approach. So, if belief is not required for our study of knowledge—and we do not have the tricky problem of ascribing mental states to entities that do not seem to have minds—what *is* there to say about our ordinary language use of group knowledge attributions?

8.1.2 Group knowledge attributions: The spokesperson view

Let us suppose that these ordinary language uses are indicative of a systematic and widespread phenomenon that deserves our attention. A number of routes have been taken when explaining this feature of language use. One option is that it is shorthand for saying that a particular individual knows it. This has been termed the 'spokesperson' view. (Lackey, forthcoming a, and forthcoming b) According to this view, when we say that a group—e.g. BP—know *p* we mean that there is an individual within the group—e.g. former chief executive of BP, Tony Hayward—knows *p*. The spokesperson view is useful for describing certain kinds of groups but wholly inadequate for others. For example, when describing the knowledge ascriptions of groups which have a formally-designated representative, the spokesperson view seems to capture much of what we are aiming at when we use the group knowledge attribution.

When the knowledge attribution regards a Court of Session decision, for instance, there is a definite decision (containing knowledge claims, beliefs, opinions, and so on) to which all members of the group (all sitting judges) adhere. All members are formally taken to subscribe to the decision by virtue of their membership of a group whose formal structure requires it. Even though one or more of the judges may not concur with the decision, it is appropriate to say that the decision represents the decision of the Court of Session.[106] Similarly, after a public inquiry into ethical standards of the British press, a particular spokesperson from a newspaper may be nominated to express the settled views of the newspaper's editorial board. It is appropriate in such cases to attribute knowledge to the group—to 'The Daily X' newspaper, for example—as we know that the members of the group have accepted (perhaps formally signed) a statement which the spokesperson then reads out. These 'formal' groups fit neatly with a 'spokesperson' approach to group knowledge attributions.

An odd quirk of the spokesperson view becomes apparent when we consider negative attributions: That is, attributions of the form: no one in the group knows *p*. Consider the following statement regarding the Deepwater Horizon oil spill,

---

[106] The Court of Session is the highest civil court in Scotland. The same point could be made with regard to any court where typically a panel of two or more judges sits and so there is a designated quorum.

> 'Nobody believed there was going to be a safety issue,' Mr. Hafle told a six-member panel of Coast Guard and Minerals Management Service officials.[107]

How would the spokesperson view deal with such cases? We must assume from what Hafle says that if not one member of the group believed there was going to be a safety issue, the group does not believe it either. Alternatively, he may be saying that every member of the group was agnostic with respect to the proposition, 'There is going to be a safety issue.' The latter, that every member of the group had no opinion with respect to whether there was going to be a safety issue, seems unlikely and negligent so we must assume that he meant the former. However, the consequence of the former statement is that every member of the group believed that there was not going to be a safety issue. The spokesperson view states that Hafle represents the group when he implies that every member of the group believed that there was not going to be a safety issue but this seems to be 'beyond his pay grade' as spokesperson of the group and beyond what the

---

[107] From Urbina, I. 2010. Documents show early worries about safety of rig. *The New York Times*, 29th May 2010. http://www.nytimes.com/2010/05/30/us/30rig.html?pagewanted=all.

spokesperson would be justified in asserting. He may be able to express the formally accepted opinion of the group but it seems to misrepresent what is believed by the members. The spokesperson view entitles the spokesperson to express the formally accepted opinion of the group. However, when it is couched in such terms as to say, 'Every member of the group believes *p*,' when it may well be the case that some of the members were agnostic towards *p*, this seems to stretch the limits of what the spokesperson is entitled to state.

Nonetheless, it may be argued that these are odd quirks and the spokesperson view typically does capture and represent the opinion of the group. However, if we consider groups which have no formally appointed representative or formal (or informal/implicit) system of agreeing to the group belief, then it is wholly inappropriate. This seems to apply to most of the cases we have discussed in this dissertation. In scientific laboratories, oil rigs, engineering offices, and so on, the knowledge that is produced by the group is rarely possessed by any single member, nor is any single member given the authority to represent a group view. In most scientific research groups or engineering offices, it is typically not the case that all members (or even most members) tacitly or formally accept a set of beliefs held by any one member. It is likely in these cases that the knowledge held by the group as a whole greatly exceeds the knowledge held by single one member of that group. Therefore, there will be cases where knowledge is attributed to the group (knowledge of the final product of the research, for example) which no single member would be justified in asserting.

8.1.3 Group knowledge attributions: Summative and aggregation procedures

Another option is that when we say that the group knows *p*, we mean that the individuals that comprise the group know *p* either in the sense that they all know it or the majority know it or some other quorum knows it. (Quinton 1975/1976; Gilbert 1994, p. 235) This is called a summative approach or judgment aggregation approach: the process of aggregating group members' individual beliefs or judgments into corresponding collective beliefs or judgments endorsed by the group as a whole. (See Cariani 2001; List and Pettit 2002, 2004; and List 2005) The problems of collective knowledge, collective belief, group minds, etc. are difficult and well-covered. The approach does, however, have problems. It is easy to see how some groups fit into the aggregation approach better than others. Court decisions, for example, are typically formally canonized. It is clear that, in the case of the U.S. Supreme Court, for instance, a 'majority opinion' is so-called because it represents the opinion of a majority of justices. (Rupert 2005)

The same applies for other formal groups such as medical panels or boards of directors where there is a formally recognized system for stabilizing and aggregating the beliefs of individual members. These kinds of groups, however, are rare when compared to the many groups that have no such formal procedures. These other kinds of groups—informal groups, if you like—are much more difficult to aggregate and much more interesting for epistemologists working on issues of testimony, communitarianism or social epistemology. Teams of logging analysts are informally grouped. So can logging analysts' beliefs be aggregated in some way just as court justices can? Consider the following dilemma:

Logging Dilemma

Three analysts are debating whether or not a particular zone in the well contains more than 90% water. Analyst A believes that it does and she believes that if it does then we should recommend that BP halt production. Analyst B also believes that there is more than 90% water in that zone but does not think that the team should make the recommendation to halt production. Analyst C believes that there is less than 90% water in the zone but does agree with A that if there was more than 90% water, then they should make the recommendation.

Following List, We can represent this dilemma as follows.

	p	p → q	q
Analyst A	True	True	True
Analyst B	True	False	False
Analyst C	False	True	False
Majority	True	True	False

Figure 9. A discursive dilemma

The table shows an odd inconsistency for aggregation theories. Whilst a majority of the group believe *p*, and believe that if *p* then *q*, a majority does not believe the simple logical conjunct of those two premises, namely *q*. List calls this the rationality challenge. As he puts it, the problem illustrates that, 'under the initially plausible aggregation procedure of majority voting, a group may not achieve consistent collective judgments even when all group members hold individually consistent judgments. (List; see also Pettit 2001; List and Pettit 2002, 2004; List 2005a)

List and Pettit (2002) have shown that no aggregation procedure can generate consistent collective judgments i) for any logically possible combination of complete and consistent individual judgments on the propositions, ii) where each individual judgment has equal weight in determining collective judgments and iii) the collective judgment on each proposition depends only on the individual judgments on that proposition, and the same pattern of dependence holds for all propositions. There are some alternatives available to those who wish to preserve aggregation procedures. The first is to give up on i). In some cases, this may be desirable. Where, for example, the disagreements are sufficiently limited or there are other mechanisms in place for reducing disagreement. For example, a group that would not experience the odd combinations of individual judgments as in the dilemma, need not necessarily worry about i). However, could the group be sure that such irrational occurrences could not happen? As the procedure does not work in general, it is philosophically unsatisfying. Alternatively we could reject ii). Pauly and van Hees (2005) have shown that if we give up only on ii), the only possible aggregation procedure is a 'dictatorial procedure'. (List 2005)

That is, the collective judgments are always those of some antecedently fixed group member. In some groups this will be satisfactory (perhaps one person will be nominated committee chair, for instance). In teams of logging analysts, it will rarely be appropriate. Finally, we could give up on iii). In other words, we could treat the 'premises' in the table as different in priority from the 'conclusions'. However, as List points out, aggregation procedures which violate iii) 'may be vulnerable to manipulation by prioritizing propositions strategically'. We can easily see how this might happen in logging analysis. Suppose all three propositions in our example are designated as premises. That is, i) there is more than 90% water in the zone; ii) if there is more than 90% water, we should halt production; and iii) we should halt production. If all these are designated as premises, then all three are collectively judged to be true. If, however, just p and q are designated premises, then p is judged to be true but q and p->q is judged to be false.

Remember that it may not always be the case that it is obvious what the premises are and what the conclusions are.

There are a lot of problems with aggregation procedures and here is my solution for teams of logging analysts. At issue, ultimately, is collective responsibility. We are here asking what BP knew because we wish to hold whoever is responsible for the accident accountable. What is required is not a general solution but attention to the make-up of the sociotechnical system that forms the group. This includes the reliability of the technology providing information as well as the reliability of the receivers of that information. Epistemic responsibility requires that epistemic agents take responsibility not only for the beliefs that they form but the procedures they use to acquire those beliefs. Collective epistemic responsibility means that groups cannot shift all of the blame for 'bad' beliefs or 'bad' belief-forming processes onto one individual scapegoat but that responsibility lies with the whole system. In the wake of the Deepwater blowout, media attention focused on one individual, Tony Hayward, then CEO of BP. Partly this was because of his perceived unwillingness to take responsibility for the crisis.

Whilst we have become accustomed, as members of the general public, to group 'dictators' (as List calls them) taking the fall for the epistemic failings of the group, it is surely irresponsible (as well as being unproductive in the sense of being conducive to better systems) in most cases to apportion all of the blame to one agent for the failings of a system.

In conclusion, the lone analyst cannot take epistemic responsibility for forming false beliefs about the well on a perceptual basis. This is because those perceptual judgments must pass the arbitration of the group for it to become group knowledge. Secondly, a dictator or scapegoat should not be pinned with all the blame since that form of aggregation does not generalize and we need a more sophisticated way of describing particular sociotechnical systems. Since none of these can currently be properly blamed for the system's epistemic failings the only conclusion we have is to, as the anarchists say, 'blame the system'.

8.1.4 Group knowledge attributions: Corporate epistemic responsibility

Production logging tools (PLTs) provide us with logs which are open to interpretation in a way that perceiving the colour of an apple or even examining rock cuttings are not. A range of sometimes conflicting readings must be corroborated in such a way as to construct a convincing narrative of what is going on in the well. Sophisticated computer applications can sort through some of the mess by running statistical algorithms but can that alone provide justification? Note that two analysts reading the same log will often derive different conclusions, particularly so when it comes to specifics such as the proportion of hydrocarbons in the fluid. Here we have controversy and conflict. Nor can feedback on the truth-value of any assertions be provided from the world itself. In most cases, the analysts will present their case to production companies on the basis of internal discussions, not by testing their conclusions against the behaviour of the well. By contrast, if there was a dispute about whether a church spire was cubic or spherical, we could usually test this against the world. We could, for instance, walk to the church and take a closer look. We need not rely solely on testimony. In most cases of logging analysis, once the data is retrieved, analysts must sort out any disputes between themselves.

8.2 Epistemic communities and epistemic agents

> [Discussing the corroboration of well-log data using the temperature curve to work out total flow] EK: Are these results supposed to correspond to other analysts' judgments or results on the ground 'out there'?
> Analyst, SK: Mostly compared to whether we believe that is right, whether it fits our overall interpretation of the well. Like a police suspect. You have a perpetrator in mind, have to piece together the evidence to see if it fits. Evidence might overturn your initial suspicions but it's harder to overturn.

8.2.1 What is a community?

As we have seen, individual analysts will have their own predilections when analysing data. Although this phenomenon is not often acknowledged by epistemologists and philosophers of perception, current research suggests that we do, each of us, see the world in non-uniform ways.[108] This individuality of perception is comparable to the logging analyst.

How does this individuality relate to the group knowledge discussed in §8.1? Karin Knorr Cetina makes a persuasive case in *Epistemic cultures: How the sciences make knowledge*, that the traditional taxonomies of disciplines are not particularly helpful in contemporary science and we ought instead to divide different groupings as 'epistemic cultures'. (Knorr Cetina 1999) I have used this term elsewhere in the thesis as it captures a specific binding feature of certain sociotechnical groups: collective knowledge production. Cultures (or communities, as I use in this thesis) are to be divided by their different 'knowledge machinery'. She uses metaphors such as 'knowledge machinery', 'tools and instuments' and so on to engineer a notably pragmatic and praxis-centered approach to her two laboratory case studies: high energy

---

[108] Recent research on 'super-recognizers' is amongst a growing number of studies which indicate that there is much less uniformity to how each of us sees our environment than previously thought. See, e.g., Russell, Duchaine, & Nakayama 2009.

physicists at CERN during the UA2 and ATLAS program and molecular biologists at the Max Planck Institute (conducted with Klaus Amann). The research groups varied in size from 300 to 2000 scientists at CERN and from 8 to 30 scientists at MPI. How does studying the ‘construction of the machinery of knowledge construction’ help us understand what is particular to these epistemic communities? Her answer is because it allows researchers to track the different ways by which an original thought transforms from a belief into knowledge or fact.

> Different cultural systems of behaviour, as we know, construe the world differently. If they involve sign processes, as they invariably do, the question is nonetheless on what, figuratively speaking, do they place their bets and stake their money - on signs or not on signs. They may choose to process signs in ways that upgrade their impact and make the most of their potential, or in ways that downgrade their role and minimize interaction with them. In short, they may construct their world in terms of sign processes, or continuously construct it away from such processes. Molecular biology practice, in the laboratories studied, shows the latter tendency. It shows a preference for experiential knowledge acquired through mechanisms that reduce representations and turn away from sign-mediated experience. (Knorr Cetina 1999 p. 80)

These disciplines are made up by their instrumentation and the way that these instruments communicate their information and the way they are used to construct new knowledge is of their own making. For example, recall the North-seeking gyro and how its capability to derive its own orientation was only discovered from a reinterpretation of noisy data logs. Knorr Cetina observes similar phenomena underlying knowledge in high energy physics. Like Dretske, she considers knowledge to be the culmination of an transmitter-signal-receiver relationship. However, in high energy physics, these signals are invisible, shadowy footprints whose producers are never themselves observed.

> [The machinery of knowledge construction] moves in the shadowland of mechanically, electrically, and electronically produced negative images of the world – in a world of signs and often fictional reflections, of echoes, footprints, and the shimmering appearances of bygone events. (Knorr Cetina 1999, p. 217)

This is in marked opposition to the culture of molecular biologists where,

> ...the person remains the epistemic subject…. The laboratory, experimentation, procedures, and objects obtain their identity through individuals. The individual scientist is their intermediary—their

> organizing principle in the flesh, to whom all things revert. (Knorr Cetina 1999, p. 217)

Essentially, the difference Knorr Cetina is noting here is between a system where instruments and technical artefacts do the perceiving and where it is primarily done by humans *seeing with* technical artefacts. In the former, the particular configurations of the instrument and its capacities are the primary factor in producing knowledge; in the latter it is the intentions and skill of the biologist who retains control and maintains direction over the experiment. This is the distinction between an extended epistemic system and an enhanced epistemic agent. We can describe the biologist using autoradiographs or X-ray films as, in some sense, an epistemic agent extended or distributed over the artefacts but it seems a rather attenuated description. The biologist is at all times in total control of what to see, what to produce and what and how to interpret. The epistemic agency remains with her. However, when the knowledge machinery is so complicated, with so many parts and configurations, the scientist herself becomes lost as but one cog in a larger machine. It becomes in these cases more appropriate to describe the agency in a distributed or extended manner.

> Those amalgams of arrangements and mechanisms -- bonded through affinity, necessity, and historical coincidence -- which, in a given field, make up how we know what we know. (Knorr Cetina 1999, p.1)

Knorr Cetina describes HEP as operating 'in terms of a negative and reflexive epistemics.'[109] (Knorr Cetina 1999, p. 46) The individual is thus 'erased' as an epistemic subject (Knorr Cetine 1999, p. 166) as she becomes part of this larger unit. (Knorr Cetina 1999, p.171)

Bird makes a distinction between two kinds of social group: 'established' and 'organic'. (Bird 2010) Whilst Bird is skeptical of the various accounts of what 'establishes' a group—in other words, what provides the necessary cohesion for us to say that the group believes or knows as a group—he concedes that less general claims about the nature of such groups can be made. He cites Gilbert, Tuomela and Schmitt as holding that established groups are held together by 'a mechanism for forming joint commitment'. (Bird 2010, p. 37) Juries,

---

[109] By 'epistemics' she means the 'strategies and practices assumed to promote the "truth"-like character of results.' (p46, original emphasis)

medical ethics panels, residents committees, and the like are 'established' in the sense that there is some procedure in place for binding the members together in commitment to a belief or assertion. Organic groups, on the other hand, receive their cohesion from an inter-dependence based on a division of labour. Collaborative scientific research projects are typical examples of organic groups. In such groups each member relies on the work of others but need not be actively involved. But the engineering groups we have looked at are a mix of both establishment and organic. We need to look more closely at the company itself.

8.2.2 Professional knowledge, cross-cultural knowledge, and attitudes

Natural language is notoriously ambiguous and the term 'knows' and its cognates are used in a variety of different senses. Epistemologists are concerned primarily with so-called 'informational' or 'propositional knowledge'. Interestingly, this is the least common usage of 'knows' in ordinary language. Epistemologists are usually concerned with what may be called 'knowledge that' as opposed to 'knowledge how' but this distinction reveals the limits and ambiguities of ordinary language. Other European languages appear to accommodate this distinction more clearly. In the French, it is the distinction between '*savoir*' and '*connaitre*'; in the German, between '*wissen*' and '*kennen*'. However, these languages also seem inadequate to define this distinction although, paradoxically, instances of each type may be generally accepted once the distinction is understood. It seems that there is an intelligible distinction to be made, though an acceptable definition has not yet been provided.

I have said before that different societies have different attitudes towards what kind of thing knowledge is and what other doxastic attitudes or metaphysical concepts it is cognate with. (see, e.g., Goldman 2004; Hallen 2004; Harding 1998; Hongladarom 2002a; Maffie 1995; Mulder 1996; Weinberg, Nichols, & Stich 2001) It should be noted that my own fieldwork research was conducted in another culture and that this dissertation may run the risk of neglecting this fact, particularly considering a central thesis that proper knowledge attribution is partly-conditioned by communities (and, by extension, cultures). The company in the case study is American and multinational but the vast majority of employees are from different parts of Southeast Asia and it would be unwise, I think, to translate Western intuitions and preconceptions onto, for example, Thai cultural practices.

There is some, although not nearly enough, literature about the differences in epistemic practices in the West and the East. Alvin Goldman discusses Japanese culture in *Knowledge in a Social World.* (Goldman 1999, p. 125-126) Soraj Hongladarom argues for the possibility of alternative epistemic cultures to the Western or veritistic tradition. He argues that Thailand is a possible example of a culture which does not (intuitively or academically) consider knowledge as justified true belief and values social harmony over truth-as-correspondence. (Hongladarom 1996 and 2002a) Goldman would likely respond that if that is the case then it is a culture with poor epistemic practices or one that does not value knowledge after all. Some epistemologists wish knowledge to be an acultural phenomenon—that is, to have no variance over cultures with regard to proper epistemic practices. It seems untenable to me but that is not what I wish to argue here. If engineers in Asia are less likely to connect knowledge and truth it is because engineers do not typically think of their work as aimed towards truth (as a scientist might). If I start to talk about truth—e.g. is this statement true—I am typically greeted with an exasperated expression or thrown up hands: 'Who knows, right? It's just a guess.'

In this dissertation, we have been discussing perceptual knowledge and it is worth reflecting on the prominence given to perception in different communities. It is common in modern analytic epistemology for the author to use perceptual cases to make their case.[110] There have been (and still exist) societies in which perceptual knowledge is privileged above all other forms of knowing. In 'Yoruba moral epistemology', Barry Hallen describes how the Yoruba language culture

---

[110] See, for example, the fake barn case cited in Goldman 1976: 'Henry's belief that the object is a barn is caused by the presence of the barn; indeed, the causal process is a perceptual one.' G. E. Moore, in his refutation of skepticism which is not particularly focused on perceptual knowledge, frequently used perceptual examples. (Moore 1903, p. 41; and Moore 1959, p. 128) See also Cohen (1998) and the 'sheep-shaped rock' thought experiment; Palle Yourgrau's 'zebra and painted mule' thought experiment (1983, p. 183; and DeRose 2000b); and countless others.

make a distinction between *imo* (putative knowledge) and *igbagbo* (putative belief). One can only have *imo* if they have witnessed something first-hand or if it is a personal matter. Hallen writes, 'The example most frequently cited by discussants, virtually as a paradigm, is visual perception of a scene or an event as it is taking place.' (Hallen 2004, p. 298)

The take-home message is that if one wishes to make broad principled statements about the nature of knowledge – its relation to truth, its relation to certain virtues – one had better be prepared to back that up either with evidence that the intuition is more than cultural prejudice or that it does not matter what other communities prefer as there are solid reasons for thinking that this way is correct. I would prefer in this thesis not to pronounce on the best route here but do wish to say that the case is still to be made one way or the other. In the absence of such a case it is always worthwhile pointing out that other communities do not share the intuitions of other communities when it comes to knowledge, truth, and cognate concepts.

If neither of those routes is preferable I would argue that the most plausible explanation is that knowledge is a community-dependent term. And further that one of the best approaches for explaining that community-dependence is that developed in Kusch's communitarianism. If 'knowledge' does indeed mark a social status then it is little surprise that the application of this status varies across communities since we know already how varied other social statuses are even amongst quite closely-related communities.

8.2.3 Epistemic responsibility revisited

In recent years we have seen scientists and epistemic agents held responsible for prediction and interpretations which, arguably, extend beyond our current epistemic capacities. In September 2011, six Italian scientists and a former government official appeared on trial for manslaughter on the grounds that they failed to predict and communicate the risk of an earthquake which occurred in 2009.[111] At the beginning of 2011, a proposed amendment to the Weather Service Bill meant that South Africa's independent weather forecasters could be imprisoned for up to ten years – or fined up to £800,000 – if they issue incorrect severe weather warnings without official permission.[112] Increasingly, questions will be asked of those who hold positions of authority and expertise and who are trusted to advise parties on the probability of geological events.

To put it another way, there appears to be a legitimate question to ask concerning those whose social status confers a particular epistemic responsibility upon them to know the probability of certain events occurring. One of Congressman Markey's tasks as chairman of the

---

[111] The defendants were scientists who held various positions to do with advising government and the public on seismic activity in the area—head of the Serious Risks Commission, director of the Civil Protection Agency's earthquake risk office, etc. For more details see Hall, S. S. 2011; and *The New York Times*, Fountain, H. 3rd October 2011. Trial over earthquake in Italy puts focus on probability and panic. Distressingly, the scientists were found guilty of manslaughter by the Italian courts. http://www.bbc.co.uk/news/world-europe-20025626. Retrieved 22nd October 2012.

[112] Reported in The Telegraph, 12th January 2012. Newling, D. South Africa weather forecasters threatened with jail if predictions wrong. http://www.telegraph.co.uk/topics/weather/9010030/South-Africa-weather-forecasters-threatened-with-jail-if-predictions-wrong.html.

select committee was to investigate what BP knew so that they may, if necessary, be held accountable for failing to act appropriately on that knowledge. Epistemologists ought to also be of help when it comes to holding agents responsible for their epistemic behaviour. A large research project in epistemology has attempted to discover the conditions under which an agent is epistemically responsible. (Amaya 2008; Bishop 2000; Code 1984, 1994; Corlett 2008; Engel 2009; Greco 1990, Hetherington 2002, 2003; Kornblith 1983; Steup 2001; Vaesen 2011a) More work is needed to deliver an accepted general theory of when an agent is epistemically responsible and to inform legal and political policy.

If we cannot do this then we have little chance of answering questions asked at the outset such as 'Who is responsible for knowing this?' and 'What can we reasonably expect an epistemic agent to know in this context?' Epistemic responsibility is the idea that an epistemic agent acquires certain responsibilities accruing to her role as an epistemic agent. What those responsibilities are is a matter of great interest and debate amongst epistemological academia.[113] My tentative account at this point is that the responsibilities derive from the agent's social status. The agent's responsibilities are what is reasonable to expect of an agent given her social status and context. This is a very common legal test in most jurisdictions and may be roughly applicable in epistemic tests.

An October 2008 special issue of *Episteme* marked a revival in interest in legal theory amongst analytic epistemologists. Alvin Goldman has claimed that legal adjudication systems primarily aim '(at least in part) to secure accurate judgment about material facts arising in a legal dispute' (Goldman 2008, p. 129). In fact, disputes about legal facts are relatively rare and are most commonly disputed in criminal cases. Yes, this is a part of legal process but it is not at all obvious that epistemic rules of this sort ought to be applied to court systems in general when 'what happened' is often not at issue and the case centres around disputes of law, legal interpretation. Courts are interested in epistemic

---

[113] See, e.g., Axtell 2008; Code 1984 and 1987; and Greco 2001.

terms—truth, knowledge, evidence—but the 'truths' they are interested in are primarily those stipulated by the Will of Parliament. A legal adjudication system is a quasi-epistemic system: that is, it should be evaluated significantly but not primarily on the basis of its epistemic performance.

Does epistemic responsibility pertain to groups? Return to the courtroom for a moment. It is common to talk of rights and responsibilities pertaining to natural persons, that is individuals. But do juristic or legal persons *in general* have rights? That is, do corporations and other legal entities have rights?

8.2.4 Group knowledge attributions and distributed-extended cognition theories

Here is one possible answer to the problem of group knowledge attributions based on the arguments I have set out in this thesis. Let us suppose that the reader accepts all the main arguments contained in this thesis: That knowledge is a social status and depends on attributions from others; that perceptual knowledge is an informational and inferential kind of knowledge; and that knowledge typically depends on a network of people and technical artefacts. One conclusion we might draw from these three principles is that attributing knowledge to a group is no more mysterious or in need of explanation than attributing knowledge to an individual. It does seem to be socially appropriate to attribute knowledge to groups. (see 7.1.1) Groups are capable of possessing perceptual informational states as described in Chapter 5. They are capable of processing information and updating their internal states on the basis of external stimuli just as we are. Finally, groups are typically constituted of networks of people and technical artefacts. In the literature, this is described as a sociotechnical system. Most groups do their epistemic work (inquiry, investigation, belief-formation, assertion) within a technological framework (a court room, a science laboratory, an oil rig, an office).

The conclusion I draw from this is that it is possible to argue that groups often know just in the same way that individuals know. Just as we might think of the eye as an evolutionary augmentation of the existing sensory capacities our ancestors had, and as we might think of the microscope as another augmentation (albeit an artificial one), we can think of other technical artefacts and other people performing specific professional roles as further augmentations. Granted these conclusions rest on the axioms presented in this dissertation and I have made my best efforts in presenting them fairly and deliberatively. If the reader finds the communitarian account of knowledge, the informational account of perceptual knowledge, or CAT implausible on whatever grounds, the conclusion that groups have knowledge becomes increasingly tenuous.

8.2.5 Final conclusions

At the start of this thesis I presented a suite of research problems.

1. What is social epistemology and how can it help contextualize the particular epistemic responsibility that a professional person holds. In particular, how can a communitarian epistemology – which treats knowledge as a kind of social status—define different kinds of knowledge?

2. What kind of knowledge do petroleum engineers have? What is technical knowledge?

3. What are technical artefacts and how does this affect technical knowledge?

4. Some technical knowledge appears to be perceptual. How do perceptual sources generate technical knowledge?

5. What epistemic agents can have this knowledge (individuals, groups, animals, artefacts, computers, etc.)?

Let us now recap the answers.

1. Social epistemology was shown to be a broad church with many non-complementary factions.

I selected communitarianism, contextualism, and the Strong Programme as being of particular utility as the main purpose of the thesis was to give a practical application to the sophisticated theories of social epistemology. These three, I have argued, are most apt to perform this task. I argued that, with communitarians, we should see knowledge attributions as conferring a particular social status upon an individual or group. In this thesis I have focussed on a particular group of society, petroleum engineers in South-East Asia, and argued that they have a distinctive kind of technical knowledge. I have also explored the extent to which this knowledge is perceptual.

2. I argued that technical knowledge is knowledge of how to design, produce or operate technical artefacts.

I presented a history of the design and production of technical artefacts in the field of study and then analysed the texts used in their operation.

3. I explored the ontological nature of technical artefacts and its connection to the ontological nature of knowledge.

Since I claim that technical artefacts can affect our human concepts (and thus what counts as knowledge) it was important to establish what technology, and technical artefacts, is. I argued that technical artefacts are a branch of artificial kind. I proposed a five-fold necessary criteria for a technical artefact:

a) Must use both structural and social terms in order to be comprehensively described;

b) Is designed (or chosen) to meet a need;

c) Has a proper function;

d) Has proper conditions of use;

e) Is not appropriated as part of a single unorthodox instance of use but is part of a systematic structure of tool use of that kind which determines 1-4.

4. I explored in detail how our perceptual faculties may be said to acquire knowledge.

I proposed a theory of an extended epistemic system which incorporates technical artefacts into its knowledge-gathering mechanics. I looked at the phenomenological experience of workers in engineering and assessed whether this experience was indicative of a more substantial point about the nature of the knowledge they aim to generate.

I proposed the following conditions on technical knowledge:

a) Technical knowledge is a distinctive kind of knowledge.

b) It includes use know-how as well as knowledge of how to design, produce, and operate artefacts.

c) This knowledge is a central component of engineering and provides a distinction between the knowledge possessed by engineers and that possessed by natural scientists (although there is undoubtedly a great deal of cross-over between the disciplines).

d) Technical artefacts have a proper function. Comprehensive knowledge requires knowledge both of this proper function and its empirical characteristics.

e) Proper function is constrained by normative considerations which are collectively and performatively sustained.

f) Technical artefacts are an artificial kind.

g) Artificial kinds are social institutions (due to self-reference)

h) Technical knowledge (knowledge of technical artefacts) is knowledge of a social institution.

i) Proper function is known, in the main, through testimony.

j) This substantiates Kusch's claim at least in this case that the community is prior to the individual in the order of explanation.

k) Technical knowledge is social.

5. I have throughout this thesis attempted to relate any theoretical questions back to lived experience and, in particular, to the findings of my fieldwork study.

Whether or not one accepts the conclusion that knowledge attributions extend across non-human or human-plus-x systems, I have strongly defended the claim that any comprehensive analysis of technical knowledge attribution which wishes to have utility must provide both social and physical descriptions of the knowledge attributions in question just as any comprehensive analysis of the ontological nature of technical artefacts requires a description of both its social and physical components. I would submit that this claim justifies the relevance of applied social epistemology to epistemology in general.

Doubtless, to substantiate many of the claims in this thesis, more empirical research needs to be done. We need to know more about the epistemic practices of engineers, the differences between epistemic communities across professions as well as nations. Finally, we need to know more about the design, production, and operation of technical artefacts and how this conditions our epistemic practices. I put this thesis forward as an initial inquiry into what needs to be done and an answer to critics skeptical of the need for social epistemology.

Appendix

Source	**Mea**	**Tra**	**Int**	**Inf**	**Rel**	**Fun**	**Soc**	**Oth**	**Tot**
BP Amoco (1999)	64	23	41	128	149	102	21	2	402
%				32	37	25	5	1	100
Hawkinson (2000)	60	28	20	108	98	99	10	3	318
%				34	31	31	3	1	100
Darling (2005)	43	7	51	101	83	61	30	5	280
%				36	30	22	11	2	100
Kaulback (2009)	33	36	12	81	117	364	58	8	628
%				13	19	58	10	1	100
SDI (2010)	22	2	6	30	19	20	9	2	80

%				38	24	25	11	3	100
Total	222	96	130	448	466	646	128	20	1708
Total %				26	27	38	7	1	100

Table #.1 'Tool' reference in five petroleum engineering texts

When Kaulback (2009) is excluded, the total percentages become respectively 34, 32, 26, 1, 6.

Cross-validation results:

Source	**Mea**	**Tra**	**Int**	**Inf**	**Rel**	**Fun**	**Soc**	**Oth**	**Tot**
BP Amoco (1999)	x	x	x	98	110	92	20	10	330
%				30%	33%	28%	6%	3%	

Cross-validated results were not requested for the subset of informational references (measurement, transmission and interpretation). The results are only to be used to cross-validate the quantity and spread of informational, reliability, functional, social and other references.

Key to table #.1:

Source, means the text used.

Mea, means number of tool-references to measurement.

Tra, means number of tool-references to transmission.

Int, means number of tool-references to interpretation.

Inf, means number of tool-reference to informational (the combined total of tool-references to measurement, transmission and interpretation).

Rel, means number of tool-references to reliability.

Fun, means number of tool-references to proper function.

Soc, means number of tool-references to societal.

Oth, means number of other tool-references.

Tot, means the total number of references and percentages thereof.

References

Adams, F. & Aizawa, K. 2001. The bounds of cognition. *Philosophical psychology* 14: 43-64.

Alcoff, L. 2001. On judging epistemic credibility: is social identity relevant? In N. Tuana & S. Morgan (eds.) *Engendering Rationalities*. Albany: SUNY Press.

Alhanati, F. J. S. 2011. Production operations. *Journal of petroleum technology* 76, March 2011: 76-85.

Allinson, M. 2009. Well-being futures: Images of their futures for the young engineers of 2009. *International conference on science, technology and innovation for sustainable well-being*, Mahasarakham University, Thailand.

Allo, P. (ed.) 2010. Luciano Floridi and the philosophy of information. *Metaphilosophy* 41 (3).

_____ et al. 2012. *The philosophy of information – a simple introduction*. Retrieved from http://www.socphilinfo.org/teaching/book-pi-intro.

Alston, W. 2005. *Beyond "justification": Dimensions of epistemic evaluation*. Ithaca, NY: Cornell University Press.

Amaya, A. 2008. Justification, coherence, and epistemic responsibility in legal fact-finding. *Episteme* 5(3): 306-319.

Anderson, E. 1995. Feminist epistemology: an interpretation and defence. *Hypatia* 10: 50-84.

Anscombe, G. E. M. 1976. The question of linguistic idealism. *Acta philosophica fennica* 27: 188-215.

Arkiner, S. & Aldis, A. (eds.) 2004. *The Caspian: Politics, energy and security*. New York: Routledge.

Austin, J. L. 1975. *How to do things with words*. Cambridge MA: Harvard University Press.

Auyang, S. Y. 2009a. Knowledge in science and engineering. *Synthese* 168: 319-331.

_____. 2009b. What is wrong with 'technology as applied science'? Retrieved from http://www.creatingtechnology.org/eng/apply.htm. February 21st 2011.

Axtell, G. 2008. Expanding epistemology: A responsibilist approach, *Philosophical papers* 37(1): 51-87.

Ayer, A. J. 1956. *The problem of knowledge*. London: Macmillan.

Bach, K. 2005. The emperor's new 'knows'. In G. Preyer & G. Peter (eds.) *Contextualism in philosophy: knowledge, meaning, and truth*. Oxford: Oxford UP.

Baird, D. 1997. Scientific instrument making, epistemology, and the conflict between gift and commodity economies. *Philosophy and technology* 2(3-4): 25-43.

_____. 2004. *Thing knowledge: a philosophy of scientific instruments*. Berkeley: University of California Press.

Baker, L. R. 2004. The ontology of artefacts. *Philosophical explorations* 7: 99-111.

_____. 2006. On the twofold nature of artefacts. *Studies in the history and philosophy of science* 37: 132-136.

_____. 2008. The shrinking difference between artifacts and natural objects. *APA newsletter on philosophy and computers* 7(2): 2-5.

Balbi, J. 2008. Epistemological and theoretical foundations of constructivist cognitive therapies: post-rationalist developments. *Dialogues in philosophy, mental and neuro sciences* 1(1): 15-27.

Barnes, B. 1977. *Interests and the growth of knowledge*. Routledge.

_____. 1982a. *T.S. Kuhn and social science*. London: Macmillan.

_____. 1982b. The science-technology relationship: A model and a query. *Social studies of science* 12/1: 166-72.

_____. 1983. Social life as bootstrapped induction. *Sociology* 17: 524–545.

_____. 2000. *Understanding agency: Social theory and responsible action.* London: Sage.

_____. 2001. On '*the construction of social reality*' Revue Internationale de Philosophie, 55 (216): 263–267.

Barnes, B. & Bloor, D. 1982. Relativism, rationalism and the sociology of knowledge. In M. Hollis and S. Lukes (eds.), *Rationality and Relativism*. Oxford: Blackwell: 21-47.

Barnes, B., Bloor, B., & Henry, J. 1996. *Scientific knowledge: a sociological analysis*. London: Athlone.

Baumann, P. 2011. WAMs: why worry? *Philosophical papers* 40(2): 155-177.

Beck, B. B. 1980. *Animal tool behaviour: the use and manufacture of tools by animals*. New York: Garland.

Bell, D. 2012. Communitarianism. *The Stanford encyclopedia of philosophy (Spring 2012 edition)*. E. Zalta (ed.) http://plato.stanford.edu/archives/spr2012/entries/communitarianism/

Berkeley, G. 1803. *Alciphron, or the minute philosopher*. New-Haven: Increase Cooke & Co.

Biagioli, M. (ed.) 1999. *The Science Studies Reader*. Routledge.

_____. 2000. Replication of monopoly? The economies of invention and discovery in Galileo's observations. *Science in context 13*, 3(4): 547-590.

Bijker, W. E., Hughes, T. P., & Pinch, T. (eds). 1987. *The social construction of technological systems: New directions in the sociology and history of technology*. Cambridge, MA: MIT Press.

Bijker, W. E. & Law, J. (eds.) 1992. *Shaping technology/building society: Studies in sociotechnical change*. Cambridge, MA: MIT Press.

Bird, A. 2010. Social knowing: The social sense of 'scientific knowledge.' *Philosophical perspectives* 24: 23-56.

_____. & Tobin, E. 2008. Natural kinds. *The Stanford encyclopedia of philosophy (Autumn 2008 edition)* E. Zalta (ed.) http://plato.stanford.edu/entries/natural-kinds/.

Bishop, M. A. 2000. In praise of epistemic irresponsibility: how lazy and ignorant can you be? *Synthese* 122(1-2): 179-208.

Blaauw, M. 2003. WAMming away at contextualism. *Nordic Journal of Philosophy* 4: 88-97.

Blackburn, S. 1984. The individual strikes back. *Synthese* 58: 281-301.

Bloor, D. 1983. *Wittgenstein: A Social Theory of Knowledge*, New York: Columbia UP.

_____. 1986. Wittgenstein and Mannheim on the sociology of mathematics. In S. Shankar (ed.) *Ludwig Wittgenstein: critical assessments*. New Hampshire: Crook Helm.

_____. 1991 [1976]. *Knowledge and Social Imagery* (2nd ed.) Chicago: University of Chicago Press.

_____. 1997. *Wittgenstein, rules and institutions*. London: Routledge.

_____. 1999a. Anti-Latour. *Studies in history and philosophy of science* 30/1: 81-112.

_____. 1999b. Discussion. Reply to Bruno Latour. *Studies in history and philosophy of science* 30/1: 131-136.

_____. 2007. Epistemic grace: Antirelativism as theology in disguise. *Common knowledge* 13: 250-280.

_____. 2011. The enigma of the aerofoil: Rival theories in aerodynamics, 1909-1930. Chicago: University of Chicago Press.

Bluff, L. A., Troscianko, J., Weir, A. A. S., Kacelnik, A., & Rutz, C. 2010. Tool use by wild New Caledonian crows corvus moneduloides at natural foraging sites. *Proceedings of the royal society b*: 1377-1385.

Bonen, Z. 1979. *Evolution of a sociotechnical system: a model and some implications*. Cambridge, MA: Alfred P. Sloan School of Management, MIT.

Bouwsma, O. K. 1965. Descartes' evil genius. In A. Sesonske & N. Fleming (eds.) *Meta-meditations: Studies in Descartes*. Belmont, CA: Wadsworth.

Boyd, R. 1991. Realism, anti-foundationalism, and the enthusiasm for natural kinds. *Philosophical studies* 61: 127-148.

BP Amoco Upstream Technology Group. 1999. *Directional survey handbook*. Aberdeen: ODL Wood Group.

BP. 2010. Deepwater horizon: Accident investigation report. Retrieved from http://www.bp.com/liveassets/bp_internet/globalbp/globalbp_uk_english/gom_response/STAGING/local_assets/downloads_pdfs/Deepwater_Horizon_Accident_Investigation_Report_Executive_summary.pdf

Brad Wray, K. 2009. The epistemic cultures of science and Wikipedia: a comparison. *Episteme* 6(1): 38-51.

Brown, H. 2012. *Production logging for base managers*. Unpublished.

Brown, J. 2006. Contextualism and warranted assertibility manoeuvres. *Philosophical studies* 130: 407-35.

Bucciarelli, L. L. 1996. *Designing engineers*. Cambridge: MIT Press.

____. 2003. *Engineering philosophy*. Delft: Delft UP.

____. 2009. The epistemic implications of engineering rhetoric. *Synthese* 168: 333-356.

Bunge, M. 1966. Technology as applied science. *Technology and culture* 7: 329-347. Amsterdam: Spinger.

____. 1977. *Treatise on basic philosophy: ontology I: the furniture of the world,* vol. 3. Amsterdam: Springer.

Callon, M. 1986. Some elements of a sociology of translation: Domestication of the scallops and the fishermen of St Brieuc Bay. In J. Law (ed.) *Power, action and belief: A new sociology of knowledge?* London: Routledge: 196-223.

Cariani, F. 2011. Judgment aggregation. *Philosophy compass* 6: 22-32.

Carroll, L. 2012. *Through the looking glass and what Alice found there.* Retrieved from http://etext.lib.virginia.edu/toc/modeng/public/CarGlas.html.

Carter, J. A. & Kerr, E. T. In preparation. Richard Rorty and epistemic value. *Draft paper.*

Cavanagh, T. 2008. Diverse designing: sorting out function and intention in artifacts. In P. Vermaas, P. Kroes, A. Light, & S. A. Moore (eds.) *Philosophy and design: from engineering to architecture.* Springer: 301-316.

Choo, C. W. 2002. *Information management for the intelligent organization: The art of scanning the environment* (3rd edition) Medford, NJ: Information Today, Inc.

Churchland, P. M. 1986. The continuity of philosophy and the sciences. Mind and language 1 (1): 5-14.

Clark, A. 2001. *Mindware.* Oxford: OU Press.

_____. 2002. Minds, brains and tools. In Hugh Clapin (ed.) *Philosophy of mental representation.* Oxford: Clarendon Press.

_____. & Chalmers, D. 1998. The extended mind. *Analysis* 58: 7–19.

Code, L. 1984. Toward a 'responsibilist' epistemology. *Philosophy and phenomenological research* 45/1: 29-50.

_____. 1987. *Epistemic responsibility.* Hanover: University Press of New England and Brown University Press

_____. 2010. Testimony, advocacy, ignorance: Thinking ecologically about social knowledge. In A. Millar, A. Haddock & D. Pritchard (eds.) *Social Epistemology.* Oxford University Press: 29-50.

Cohen, S. 1987. Knowledge, context, and social standards. *Synthese* 73(1): 3-26.

_____. 1988. How to be a falliblist. *Philsosophical perspectives* 2:91-123.

_____. 1994. Some steps towards a general theory of relevance. *Synthese* 101: 171-85.

_____. 1998. Contextualist solutions to epistemological problems: scepticism, Gettier, and the lottery. *Australasian journal of philosophy* 76: 289-306.

Collins, H. 2001. Tacit knowledge, trust and the Q of sapphire. *Social studies of science* 31 (1): 71-85.

_____. & Evans, R. 2008. You cannot be serious! Public understanding of technology with special reference to "Hawk-eye". *Public understanding of science* 17 (3): 283-308.

Constant, E. 1980. *The origins of the turbojet revolution.* John Hopkins UP.

Corlett, J. A. 1996. *Analyzing social knowledge*. Lanham, MD: Rowman and Littlefield.

____. 2008. Epistemic responsibility. *International journal of philosophical studies* 16 (2): 179-200.

Craig, E. 1990. *Knowledge and the state of nature: An essay in conceptual synthesis*. Oxford: Clarendon Press.

____. 2007. Genealogies and the state of nature. In A. Thomas (ed.) *Bernard Williams*. Cambridge: Cambridge UP: 181-200.

Darling, T. 2005. *Well logging and formation evaluation*. Amsterdam: Elsevier.

Day, T. & Reynolds, C. J. 2000. Cosmologies, truth regimes, and the state in Southeast Asia. *Modern Asian studies* 34: 1-55.

De Roure, D. & Frey, J. 2007. Three perspectives on collaborative knowledge acquisition in e-science. *Workshop in semantic web for collaborative knowledge acquisition*. Twentieth international joint conference on artificial intelligence, Hyderabad, India.

Dennett, D. C. 1984. Cognitive wheels: the frame problem in artificial intelligence. In C. Hookway (ed.) *Minds, machines and evolution*. Cambridge UP.

____. 1990. The interpretation of texts, people and other artifacts. *Philosophy and phenomenological research* 50/S: 177-194. Reprinted in M. Losonsky (ed.) *Language and mind: Contemporary readings in philosopohy and cognitive science*. Blackwells [1995].

DeRose, K. 1992. Contextualism and knowledge attributions. *Philosophy and phenomenological research* 52: 913-929.

____. 1995. The ordinary language basis for contextualism, and the new invariantism. *Philosophical quarterly* 55 (219): 172-198.

____. 1999a. Solving the sceptical problem. In DeRose, K. & Warfield, T. A. *Skepticism: A contemporary reader*. Oxford: Oxford UP: 183-219.

____. 1999b. Contextualism: An explanation and defense. In J. Greco & E. Sosa (eds.) *The Blackwell guide to epistemology*. Oxford: Blackwell: 187-205.

____. 2002. Assertion, knowledge and context. *Philosophical review* 111.

____. 2009. The case for contextualism: knowledge, scepticism, and context, vol. 1. Oxford: Oxford UP.

____. 2012. Bamboozled by our own words: semantic blindness and some arguments against contextualism. *Philosophy and phenomenological research* 73 (2): 316-338.

____. & Warfield, T. A. (eds.) 1999. *Skepticism: A contemporary reader.* Oxford: Oxford UP.

Downey, G. L. & Lucena, J. C. 1995. Engineering cultures. In S. Jasanoff, G. E. Markle, J. C. Petersen & T. Pinch (eds.) *Handbook of science and technology studies*. Thousand Oaks, CA: Sage: 167-88.

Dretske, F. I. 1969. *Seeing and knowing*. Chicago, IL: University of Chicago Press.

____. 1970. Epistemic operators. *Journal of philosophy* 67: 1007-23.

____. 1971. Conclusive reasons. *Australasian journal of philosophy* 49: 1-22.

____. 1981a. *Knowledge and the flow of information*. Cambridge, MA: MIT Press.

____. 1981b. The pragmatic dimension of knowledge. *Philosophical studies* 40 (3): 363-78.

____. 1983. The epistemology of belief. *Synthese* 55 (1): 3-19.

____. 2005a. Is knowledge closed under known entailment? In E. Sosa & M. Steup (eds.) *Contemporary debates in epistemology*. Malden, MA: Blackwell.

____. 2005b. The case against closure. In E. Sosa & M. Steup (eds.) *Contemporary debates in epistemology*. Malden, MA: Blackwell.

____. 2005c. Reply to Hawthorne. In E. Sosa & M. Steup (eds.) *Contemporary debates in epistemology*. Malden, MA: Blackwell.

____. 2006. Information and closure. *Erkenntnis* 64: 409-13.

Dummett, M. 1993. *Origins of analytical philosophy*. London: Duckworth.

Dupré, J. 1993. *The disorder of things: Metaphysical foundations of the disunity of science*. Cambridge MA: Harvard University Press.

____. 1996. Promiscuous realism: Reply to Wilson. *British Journal for Philosophy of Science* 47: 441-444.

Durkheim, E. 1956. *Education and sociology*. London: Collier-Macmillan.

Dyson, F. 1979. *Disturbing the universe*. London: Harper & Row.

Engel, P. 2009. Epistemic responsibility without epistemic agency. *Philosophical explorations* 12/2: 205-219.

Ernest, P. 1998. *Social constructivism as a philosophy of mathematics*. Albany, NY: State University of New York Press.

Evans, G. 1973. The causal theory of names. *Proceedings of the Aristotelian Society, Supplementary Volume* 47: 187–208.

Fallis, D. 2002. Introduction. Social epistemology and information science. *Social epistemology* 16/1: 1-4.

Finn, J. K., Tregenza, T., Norman, M. D. 2009. Defensive tool use in a coconut-carrying octopus. *Current biology* 19/23: R1069-R1070.

Floridi, L. 2002.What is the philosophy of information. *Metaphilosophy* 33/1-2: 123-145.

____. 2003. Two approaches to the philosophy of information. *Minds and machines* 13/4: 459-469.

____. 2004. Open problems in the philosophy of information. *Metaphilosophy* 35/4: 554-582.

____. 2005. Is semantic information meaningful data?' *Philosophy and phenomenological research* 70/2: 351-370.

____. 2008a. Trends in the philosophy of information. In P. Adriaans & Johan van Benthem (eds.) *Handbook of philosophy of information*. Amsterdam, Oxford: Elsevier: 113-32.

____. 2008b. Understanding epistemic relevance. *Erkenntnis* 69/1: 69-92.

____. 2010a. *The philosophy of information*. Oxford: Oxford University Press.

____. 2010b. The philosophy of information as a conceptual framework. *Knowledge, technology and policy* 23:1-2.

____. 2010c. *Information: a very short introduction*. Oxford: Oxford UP.

____. 2011a. Semantic conceptions of information. *The Stanford encyclopedia of philosophy (Spring 2011 edition)*. E. N. Zalta (ed.) http://plato.stanford.edu/archives/spr2011/entries/information-semantic/.

____. 2011b. Enveloping the world for AI. *The philosopher's magazine* 54: 20-21.

____. 2012. Replies. In H. Demir (ed.) *Luciano Floridi's philosophy of technology: critical reflections*. Springer.

Franssen, M. 2006. The normativity of artefacts. *Studies in history and philosophy of science* 37: 42-57.

____. & Bucciarelli, L. L. 2004. On rationality in engineering design. *Journal of mechanical design* 126: 945-949.

Frege, G. 1960 [1892]. On sense and reference. In P. Geach & M. Black (eds.) *Translations from the philosophical writings of Gottlob Frege*. Oxford: Basil Blackwell: 56-78.

Fricker, M. 2007. *Epistemic injustice: Power and the ethics of knowing*. Oxford: Oxford UP.

____. 2008. Scepticism and the genealogy of knowledge: situating epistemology in time. *Philosophical papers* 37(1): 27-50.

Friedman, M. 1998. On the sociology of scientific knowledge and its philosophical agenda. *Studies in history and philosophy of science* 29(2): 239-71.

Friman, O. 2003. Adaptive analysis of functional MRI data. *Linköping Studies in Science and Technology, Dissertation No. 836*. Linköping, Sweden: LTAB Linköpings Tryckeri.

Froese, T., McGann, M., Bigge, W., Spiers, A., & Seth, A. K. 2012. The enactive torch: a new tool for the sciences of perception. IEEE transactions on haptics 5(4): 365-375.

Fuller, S. 1987. On regulating what is known: a way to social epistemology. *Synthese* 73(1): 145-83.

____. 1992. Epistemology radically naturalized: Recovering the normative, the experimental, and the social. In R. N. Giere (ed.) *Cognitive models of science*. Minneapolis: University of Minnesota Press: 427-59.

____. 1996. Recent work in social epistemology. *American philosophical quarterly* 33: 149-66.

____. 2002. *Social epistemology* (2nd ed.) Bloomington and Indianapolis: Indiana UP.

Gagnier, R. & Dupré, J. 1998. Reply to Amariglio and Ruccio's 'Literary/cultural "economies," economic discourse and the question of Marxism. In M. Woodmansee & M. Osteen (eds.) *The new economic criticism*. Routledge: 401-407.

Garbacz, P. 2012. What can an armchair philosopher do for a "dirty-hands" engineer? *Axiomathes* 22 (3): 385-401.

Garfinkel, H. 1967 [1984]. *Studies in ethnomethodology*. Malden, MA: Polity Press, Blackwell.

Gettier, E. 1963. Is justified true belief knowledge? *Analysis* 23: 121-23.

Giere, R. N. 2002. Discussion Note: distributed cognition in epistemic cultures. *Philosophy of science* 69(4): 637-44.

Gilbert, M. 1987. Modelling collective belief. *Synthese* 73(1): 185-204.

____. 1989. *On social facts*. London: Routledge.

____. 1994. Remarks on collective belief. In Schmitt, F. (ed.) 1994c. *Socializing epistemology: The social dimensions of knowledge*. Rowman and Littlefield: 235–56.

____. 2004. Collective epistemology. *Episteme* 1: 95–107.

Gille, B. 1986. Essay on technical knowledge. In B. Gille (ed.) *The history of techniques* vol. 2: 1136-85.

Goldberg, S. 2007. *Anti-individualism: mind and language, knowledge and justification*. Cambridge: Cambridge UP.

_____. 2012. Epistemic extendedness, testimony, and the epistemology of instrument-based belief. *Philosophical Explorations* 15(2): 181-197.

Goldman, A. I. 1976. Discrimination and perceptual knowledge. *Journal of philosophy* 73: 771-791.

_____. 1999. *Knowledge in a social world*, Oxford University Press.

_____. 2004. Group knowledge versus group rationality: Two approaches to social epistemology. *Episteme* 1(1): 11-22.

_____. 2006. Social epistemology. *The Stanford Encyclopedia of Philosophy* (Summer 2010 Edition), E. N. Zalta (ed.), <http://plato.stanford.edu/archives/sum2010/entries/epistemology-social/.

_____. 2008. The social epistemology of blogging. In J. van den Hoven & J. Weckert (eds.) *Information technology and moral philosophy*. Cambridge: Cambridge UP.

_____. 2009. Social epistemology: theory and applications. In *Royal insitute of philosophy*, supplement 64: 1-18.

_____. 2010a. Why social epistemology is real epistemology. In A. Haddock, A. Millar, & D. Pritchard. *Social epistemology*. Oxford: Oxford UP: 1-28.

_____. 2010b. Systems-oriented social epistemology. In T. Gendler & J. Hawthorne (eds.) *Oxford studies in epistemology* vol. 3.

Grasswick, H. 2008. Feminist social epistemology. *The Stanford encyclopedia of philosophy* (Fall 2008 Edition), E. N. Zalta (ed.) http://plato.stanford.edu/archives/fall2008/entries/feminist-social-epistemology/.

Greco, J. 1990. Internalism and epistemically responsible belief. *Synthese* 85/2: 245-277.

_____. 2001. Virtues and rules in epistemology. In L. Zagzebski & A. Fairweather (eds.) *Virtue epistemology: Essays on epistemic virtue and responsibility*. Oxford: Oxford UP.

Gregory, R. L. 1980. Perceptions and hypotheses. *Philosophical transactions of the Royal Society of London* Series B, Biological Sciences, 290/1038: 181-197.

Hackett, E. J., Amsterdamska, O., Lynch, M., & Wajcman, J. (eds.). 2008. *The handbook of science and technology studies*. Cambridge, MA: MIT Press.

Hacking, I. 1990. Natural kinds. In R. B. Barrett & F. R. Gibson (eds.) *Perspectives on Quine*. Cambridge, MA: Blackwell.

____. 1991a. A tradition of natural kinds. *Philosophical studies* 61: 109-126.

____. 1991b. On Boyd. *Philosophical studies* 61: 149-154.

____. 1995. The looping effects of human kinds. In D. Sperber, D. Premack & A. J. Premack (eds.) *Causal cognition: A multi-disciplinary debate*. Cambridge, MA: Harvard UP: 351-383.

Haddock, A. 2004. Rethinking the 'strong programme' in the sociology of knowledge. *Studies in history and philosophy of science* 35: 19-40.

Hakli, R. 2007. On the possibility of group knowledge without group belief. *Social epistemology* 21: 249-66.

Hall, K. R. L. 1963. Tool-using performances as indicators of behavioral adaptability. *Current anthropology* 4/5: 479-494.

Hallen, B. 2004. Yoruba moral epistemology. In Wiredu, K. (ed.) *A companion to African philosophy*. Oxford: Blackwell: 296-303.

Hankinson Nelson, L. 1993. A question of evidence. *Hypatia* 8 (2): 172-189.

Harding, S. 1991. *Whose science? Whose knowledge?* Ithaca, NY: Cornell UP.

____. 1998. *Is science multicultural?* Bloomington, IN: Indiana UP.

Harsanyi, J. C. 1955. Cardinal welfare, individualistic ethics, and interpersonal comparison of utility. *Political economy* 63. 309-21.

Hattois, G. 1984. *Le signe et la technique: la philosophie à l'épreuve de la technique*. Paris: Aubier.

Hawkinson, B. 2000. *Keeper operation manual.* Re-printed as *AdK Casing multishot, Drillpipe multishot, Memory slickline, single shot orientation* [2008]. Paso Robles: Scientific Drilling International, Inc.

Hetherington, S. 2002. Epistemic responsibility: a dilemma. *The monist* 85(3).

____. 2003. Alternate possibilities and avoidable moral responsibility. *American philosophical quarterly* 40(3): 229-239.

____. 2012. The extended knower. *Philosophical Explorations* 15(2): 207-218.

Heidelberger, M. 2006. Applying models in fluid dynamics. *International studies in the philosophy of science* 20 (1): 49-67.

Himma, K. E. 2007. The concept of information overload: A preliminary step in understanding the nature of a harmful information-related condition. *Ethics and information technology* 9: 259-72.

Hoffman, J. & Rosenkrantz, G. 1997. *Substance: its nature and existence*. Routledge.

Hongladarom, S. 1996. How is Thai philosophy possible? *The sixth international conference on Thai studies*. Retrieved from http://pioneer.chula.ac.th/~hsoraj/web/Thai.html.

_____. 2002a. Cross-cultural epistemic practices. *Social epistemology* 16/1: 82-92.

_____. 2002b. *Science in Thai society and culture*. Bangkok: Institute for Academic Development.

_____. 2008. Floridi and Spinoza on global information ethics. *Ethics and information technology* 10 (2): 175-187.

Hookway, C. 1999. Questions of context. *Proceedings of the Aristotelian society* 96: 1-16.

Houkes, W. 2006. Knowledge of artefact functions. *Studies in history and philosophy of science* 37: 102-113.

_____. 2009. The nature of technological knowledge. In A. W. M. Meijers (ed.) *Philosophy of technology and engineering sciences*, p. 187-189. Amsterdam: Elsevier.

_____. & Meijers, A. J. 2006. The ontology of artefacts: The hard problem. *Studies in the history and philosophy of science* 37: 118-131.

_____. & Vermaas, P. E. 2004. Actions versus functions: a plea for an alternative metaphysics of artifacts. *Monist* 87: 52-71.

_____. & _____. 2009. Contemporary engineering and the metaphysics of artefacts: Beyond the artisan model. *The monist* 92: 403-419.

_____., Vermaas, P. E., Dorst, K. & de Vries, M. J. 2002. Design and use as plans: An action-theoretical account. *Design studies* 23: 303-20.

Howarth, S. 1997. *A century in oil*. London: Weidenfeld & Nicolson.

Hutchins, E. 1995a. How a cockpit remembers its speeds. *Cognitive science: A multidisciplinary journal* 19/3: 265-288.

_____. 1995b. *Cognition in the wild*. Cambridge, MA: MIT Press.

Inglis, T. A. 1987. *Directional drilling*. London: Graham & Trotman.

James, W. 1890 [1983]. *The principles of psychology*. Cambridge, MA: Harvard UP.

Johnson-Frey, S. H. 2004. The neural bases of complex tool use in humans. *Trends in cognitive sciences* 8: 71-8.

Jones, T. B. & Kamil, A. C. 1973. Tool-making and tool-using in the northern blue jay. *Science* 180: 1076-1078.

Kaplan, D. B. 1989. 'Demonstratives' and 'Afterthoughts.' In Almog, et al. (eds.) *Themes from Kaplan*. Oxford: Oxford UP.

Kaplan, D. M (ed.). 2004. *Readings in the philosophy of technology*. Maryland: Rowman & Littlefield.

Karplus, I., Fielder, G. C., & Ramcharan, P. 1998. The intraspecific fighting behavior of the Hawaiian boxer crab, Lybia edmondsoni: fighting with dangerous weapons? *Symbiosis* 24: 287-302.

Kaulback, M. 2009. *Memory production logging: Operations manual V1.6*. Paso Robles: Scientific Drilling International, Inc.

Kerr, E. T. 2012. Information and meaning. In Allo, P. et al. *Philosophy of information – a simple introduction*. Retrieved from http://socphilinfo.org/teaching/book-pi-intro.

_____. 2013. Rock, paper, scissors: engineering differences between natural, social and artificial kinds. In M. Franssen, P. Kroes, P. E. Vermaas (eds.) *The metaphysics of technical artefacts*. Synthese.

_____. & Pritchard, D. H. 2012. Skepticism and information. In H. Demir (ed.) *Luciano Floridi's philosophy of technology: critical reflections*. Springer.

Kersten, D., Mamassian, P. & Yuille, A. 2004. Object perception as Bayesian inference. *Annual review of psychology* 55: 271-304.

Khun, T. 1962. *The structure of scientific revolutions*. Chicago: University of Chicago Press.

Kitcher, P. 1994. Contrasting conceptions of social epistemology. In Schmitt, F. F. (ed.) *Socializing epistemology: The social dimensions of knowledge*. Rowman and Littlefield: 111-34.

_____. 2002. Veritistic value and the project of social epistemology. *Philosophy and phenomenological research* 64/1: 191-198.

Kline, R. 1995. Constructing 'technology' as 'applied science'. *Isis* 86: 194-221.

Knill, D. C., Kersten, D. & Yuille, A. 1996. Introduction: A Bayesian formulation of visual perception. In Knill, D. C. & Richards, W. (eds.) *Perception as Bayesian inference*, 1-21. Cambridge: Cambridge UP.

_____. & Pouget, A. 2004. The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in neurosciences* 27 (12): 712-719.

Knoblich, G., Butterfill, S. & Sebanz, N. 2011. Psychological research on joint action: Theory and data. In B. Ross (ed.) *The psychology of learning and motivation*, vol. 54. Burlington: Academic Press: 59-101.

Knorr Cetina, K. D. 1982a. Relativism: What now? *Social studies of science* 12/1: 133-36.

____. 1982b. Scientific communities or transepistemic arenas of research? A critique of quasi-economic models of science. *Social studies of science* 12/1: 101-30.

____. 1999. *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge: Harvard University Press.

Koen, B. V. 2003. *Discussion of the method: Conducting the engineer's approach to problem solving*. New York: Oxford UP.

Kornblith, H. 1983. Justified belief and epistemically responsible action. *Philosophical review* 92/1: 33-48.

____. 2002. *Knowledge and its place in nature*. New York: Oxford University Press.

Kortlandt, A. & Kooij, M. 1963. Protohominid behaviour in primates. *Symposia of the zoological society of London* 10: 61-88.

Kripke, S. 1980. *Naming and necessity*. Cambridge, MA: Harvard UP.

____. 1982. *Wittgenstein: on rules and private language*. Oxford: Blackwell.

Kroes, P. A. 2002. Design methodology and the nature of technical artefacts. *Design studies* 23: 287-302.

____. 2003. Screwdriver philosophy: Searle's analysis of technical functions. *Techné* 6: 22-35.

____. 2010. Engineering and the dual nature of technical artefacts. *Cambridge Journal of Economics* 34: 51-62.

____., Franssen, M., van de Poel, I. & Ottens, M. 2006. Treating socio-technical systems as engineering systems: Some conceptual problems. *Systems research and behavioural science* 23: 803-814.

____ & Meijers, A. W. M. (eds.) 2000. The empirical turn in the philosophy of technology vol. 20. *Research in philosophy and technology*. Amsterdam: JAI/Elsevier Science.

____. & ____. 2002. The dual nature of technical artifacts: Presentation of a new research program. *Techné* 6/2: 4-8.

____. & ____. 2006. The dual nature of technical artefacts. *Studies in history and philosophy of science* 37: 1-4.

Kuhn, T. 1970 [1962]. *The structure of scientific revolutions* (2nd ed.) Chicago: University of Chicago Press.

Kusch, M. 1997. The sociophilosophy of folk psychology. *Studies in history and philosophy of science* 28: 1-25.

____. 1998. Social epistemology. *University of Cambridge research guide*. http://www.hps.cam.ac.uk/research/se.html.

____. 1999. *Psychological knowledge*. London: Routledge.

_____. 2002. *Knowledge by agreement: The programme of communitarian epistemology*. Oxford: Clarendon Press.

_____. 2004. Reply to my critics. *Social studies of science* 34/4: 615-620.

_____. 2011. Reflexivity, relativism, microhistory: Three desiderata for historical epistemologies. *Erkenntnis* 75 (3): 483-494.

Kvanvig, J. 1998. Why should inquiring minds want to know? Meno problems and epistemological axiology. *Monist* 81 (3): 426-451.

Lackey, J. *Forthcoming a*. Group knowledge attributions. In Brown, J. & Gerken, M. (eds.) *New essays on knowledge ascriptions*. Oxford: Oxford UP.

_____. *Forthcoming b*. A deflationary account of group testimony. In Lackey, J. (ed.) *Essays in collective epistemology*.

Lakoff, G. 1987. *Woman, Fire, and Dangerous Things*. Chicago: University of Chicago Press.

Lamberth, D. C. 1999. William James and the metaphysics of experience. Cambridge: Cambridge UP.

Lange, M. 2007. *Philosophy of science: an anthology*. Oxford: Blackwell.

Latour, B. 1987. Science in action: How to follow scientists and engineers through society. Cambridge, MA.: Harvard UP.

_____. 1998. On actor network theory: a few clarifications 1/2. *Pit Schultz*. Retrieved from http://amsterdam.nettime.org/Lists-Archives/nettime-l-9801/msg00019.html.

_____. 1999. For David Bloor… and beyond: a reply to David Bloor's 'anti-Latour'. *Studies in history and philosophy of science* 30/1: 113-129.

_____. 2003. The promises of constructivism. In Ihde, D. (ed.) *Chasing technology: Matrix of materiality*. Indiana: Indiana UP: 27-46.

_____. 2005. *Reassembling the social: an introduction to actor-network theory*. Oxford: Clarendon.

_____. & Woolgar, S. 1979. *Laboratory life: The social construction of scientific facts*. Sage: Beverly Hills.

Laudan, R. (ed.) 1984. *The nature of technological knowledge*. Dordrecht: Reidel.

Layton, E. T., Jr. 1974. Technology as knowledge. *Technology and culture* 15/1: 31-41.

_____. 1976. American ideologies of science and engineering. *Technology and culture* 17: 688-701.

_____. 1991. A historical definition of engineering. In P. T. Durbin (ed.) *Critical perspectives on non-academic science and engineering*. Bethlehem: Lehigh UP: 60-79.

Leite, A. 2005. Some worries for would-be *wam*mers. *Grazer philosophische studien* 69: 101-125.

Lewis, D. 1979. Scorekeeping in a language game. *Journal of philosophical logic* 8: 339-359.

List, C. 2005. Group knowledge and group rationality: a judgment aggregation perspective. *Episteme* 2/1: 25-38.

_____. 2012. The theory of judgment aggregation: an introductory review. *Synthese* 187 (1): 179-207.

_____. & Pettit, P. 2002. Aggregating sets of judgments: an impossibility result. *Economics and philosophy* 18: 89-110.

_____. & _____. 2004. Aggregating sets of judgments: two impossibility results compared. *Synthese* 140/1-2: 207-35.

_____. & _____. 2006. Group agency and supervenience.*Southern journal of philosophy* XLIV (spindle supplement): 85-105. Reprinted in J. Hohwy & J. Kallestrup (eds.) *Being reduced.* [2008] Oxford: Oxford UP.

Longino, H. 2002. *The fate of knowledge*. Princeton: Princeton UP.

Lynch, M. & Peyrot, M. 1992. Introduction: a reader's guide to ethnomethodology. *Qualitative sociology* 15/2: 113-122.

Macaulay, G. C. 1890. *The history of Herodotus*. Public domain. http://www.sacred-texts.com/cla/hh/index.htm

MacIntyre, A. 1967. The idea of a social science. *Proceedings of the Aristotelian society supplementary volume* 41: 95-114.

_____. 1984. *After virtue: A study in moral theory* (2nd ed.) Indiana: University of Notre Dame Press.

MacKenzie, D. 1993. From Kwajalein to Armageddon. In D. Gooding, T. Pinch & S. Schaffer (eds.) *The uses of experiment: studies in the natural sciences*. Cambridge: Cambridge UP: 409-436.

_____. 2006. *An engine, not a camera: how financial models shape markets*. Mass.: MIT Press.

_____. 2007. Producing accounts: Finitism, technology and rule-following. *IPEG papers in global political economy* 33.

Maffie, J. 1995. Towards an anthropology of epistemology. *The philosophical forum* 26: 218-241.

Magnus, P. D. 2009. On trusting Wikipedia. *Episteme* 6/1: 74-90.

Mahoney, M. 1999. *Reading a machine*. Princeton: Princeton University.

Malcolm, N. 1952. Knowledge and belief. *Mind* 61: 178-189.

Marsh, L. and Onof, C. 2008. Introduction to the special issue 'Perspectives on social cognition'. *Cognitive systems research* 9: 1-2.

McDermid, D. J. 2001. What is direct perceptual knowledge? A five-fold confusion. *Grazer Philosophische Studien* 62 (1): 1-16.

McDowell, J. 1998. *The content of perceptual experience.* In McDowell, J. *Mind, value and reality*, 341-358. Cambridge, MA: Harvard UP.

McFarland, D., & Bösser, T. 1993. *Intelligent behavior in animals and robots.* Cambridge, MA: MIT Press.

Merton, R. K. 1979 [1942]. The normative structure of science. In Merton, R. K. *The sociology of science: Theoretical and empirical investigations.* Chicago, IL: University of Chicago Press.

____. 1965. *On the shoulders of giants: a Shandean postscript.* Free Press.

Meijers, A. W. M. 2002. Collective agents and cognitive attitudes. *Protosociology* 16: 70-85.

____. 2003. Can collective intentionality be individualized? *American journal of economics and sociology* 62: 167-183.

Millar, A. 2000. The scope of perceptual knowledge. *Philosophy* 75: 73-88.

____. 2008a. Disjunctivism and skepticism. In J. Greco (ed.) The Oxford handbook of skepticism. Oxford: OU Press: 581-604.

____. 2008b. Perceptual-recognitional abilities and perceptual knowledge. In A. Haddock & F. Macpherson (eds.) *Disjunctivism: perception, action, knowledge.* Oxford: OU Press.

Millikan, R. 1987. Language, thought, and other biological categories. Cambridge, MA: MIT Press.

____. 1989. In defense of proper functions. *Philosophy of science* 56: 288-302.

Misak, C. J. 1995. *Verificationism: Its history and prospects.* London: Routledge.

Mitcham, C. 1999. *Thinking through technology: The path between engineering and philosophy.* Chicago: University of Chicago Press.

Mitroff, I. 1974. Norms and counter-norms in a select group of the Apollo moon scientists: a case study of the ambivalence of scientists. *American sociological review* 39/4: 579-95.

Moore, G. E. 1903. The refutation of idealism. *Mind 12*: 433-53. Reprinted in *Philosophical studies.* London: Allen and Unwin.

____. 1959 [1925]. A defence of common sense. In J. H. Muirhead (ed.) *Contemporary British philosophy* (2nd ed.) London: Allen and Unwin. Reprinted in *Philosophical papers.* London: George Allen & Unwin.

____. 1993. Moore's paradox. In T. Baldwin (ed.) *G. E. Moore: selected writings*. London: Routledge: 207-212.

Mulcahy, N.J., & Call, J. 2006. Apes save tools for future use. *Science* 312: 1038-1040.

Mulder, N. 1996. *Inside Thai society: Interpretations of everyday life.* Amsterdam: Pepin Press.

Mulkay, M. 1991 [1976]. Norms and ideology. Reprinted in Mulkay, M. *Sociology of science: a sociological pilgrimage*, p. 62-78. Milton Keynes: Oxford UP.

Muller, F. A. & van Fraasen, B. C. 2008. How to talk about unobservables. Analysis 68: 197-205.

Nagel, J. 2008. Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian journal of philosophy* 86/2: 279-294.

Neta, R. 2002. S knows that P. *Noûs* 36: 663-81.

____. 2003. Skepticism, contextualism, and semantic self-knowledge. *Philosophy and phenomenological research* 67/2: 396-411.

Nishida, T., Hasegawa, T., Hayaki, H., Takahata, Y., & Uehara, S. 1992. Meat-sharing as a coalition strategy by an alpha male chimpanzee. In T. Nishida, W. C. McGrew, P. Marler, & M. Pickford (eds.) *Topics in primatology vol. 1: human origins*. Tokyo: University of Tokyo Press.

Okasha, S. 2000. The underdetermination of theory by data and the "strong programme" in the sociology of knowledge. *International studies in the philosophy of science* 14/3: 283-97.

____. 2002. Darwinian metaphysics: species and the question of essentialism. *Synthese* 131 (2): 191-213.

Owens, D. J. 2003. Does belief have an aim? *Philosophical studies* 115/3: 283-305.

Pacherie, E. 1995. Do we see with microscopes? *The monist* 78/2: 171-88.

Paulson, J. 2005. Surveying in Ancient Egypt. *Proceedings of FIG working week*, April.

Pettit, P. 2003. Groups with minds of their own. In F. Schmitt (ed.) *Socializing metaphysics.* Oxford: Rowman and Littlefield.

Pickering, A. 1984. *Constructing quarks: A sociological history of particle physics*. Edinburgh: Edinburgh UP.

____. 1995. *The mangle of practice: Time, agency, and science*. Chicago: Chicago University Press.

Pierce, J. D. 1986. A review of tool use in insects. *Florida entymologist* 69: 95-104.

Pinch, T. & Bijker, W. E. 1984. The social construction of facts and artefacts: Or how the sociology of science and the sociology of technology might benefit each other. *Social studies of science* 14: 399-441.

Pitt, J. 2001. What engineers know. *Techné* 5/3: 17-29.

Plant, B. 2005. *Wittgenstein and Levinas: ethical and religious thought.* Oxford: Routledge.

Polanyi, M. 1962. *Personal knowledge: Towards a post-critical philosophy.* Chicago: University of Chicago Press.

____. 1969. *Knowing and being.* Chicago: University of Chicago Press.

Preston, B. 1998. Cognition and tool use. *Mind and language* 13/4: 513-547.

____. 2008. The shrinkage factor: Comment on Lynne Rudder Baker's 'The shrinking difference between artifacts and natural objects.' *APA newsletter on philosophy and computers* 8/1: 26-28.

Pritchard, D. H. 2002. Recent work on radical scepticism. *American philosophical quarterly* 39/3: 215-257.

____. 2005. Contextualism, scepticism, and warranted assertability manoeuvres. Retrieved from http://www.philosophy.ed.ac.uk/people/full-academic/documents/ContextAndWAMsFINAL_000.pdf.

____. 2007. *Epistemic luck.* Oxford UP.

____. 2008a. McDowellian Neo-Mooreanism. In A. Haddock & F. Macpherson (eds.) *New Essays on Disjunctivism.* Oxford: Oxford UP: 283-310.

____. 2008b. Knowing the answer, understanding and epistemic value. *Grazer philosophische studien* 77: 325-39.

____. 2009. Wright *contra* McDowell on perceptual knowledge and scepticism. *Synthese* 171: 467-79.

____. 2010a. Relevant alternatives, perceptual knowledge and discrimination. *Noûs* 44: 245-68.

____. 2010b. Cognitive ability and the extended cognition thesis. *Synthese* 175: 133-51.

Putnam, H. 1975. *Mind, language and reality: Philosophical papers vol. 2.* Cambridge: Cambridge University Press.

Putnam, H. 1999 [1981]. Brains in a vat. In K. DeRose & T. A. Warfield (eds.), p. 27-42. Cambridge: Cambridge UP.

Quine, W. V. O. 1969. Epistemology naturalized. In his *Ontological relativity and other essays*. New York: Columbia UP.

____. 1976. The scope and language of science: In the ways of paradox and other essays, enlarged edition. New York: Random House.

____. 1986. Reply to Morton White. In L. E. Hahn & P. A. Schilpp (eds.) *The philosophy of W. V. O. Quine*. La Salle: Open Court: 663-665.

____. 1990. *Pursuit of truth*. Cambridge, MA: Harvard UP.

Quinton, A. 1975/1976. Social objects. *Proceedings of the Aristotelian society* 75: 1-27.

Rabinow, P. 1999. *French DNA: trouble in purgatory*. Chicago: University of Chicago Press.

Raskin, J. D. 2002. Constructivism in psychology: personal construct psychology, radical constructivism, and social constructionism. In J. D. Raskin & S. K. Bridges (eds.) *Studies in meaning: exploring constructivist psychology*. New York: Pace UP: 1-25.

Rawls, J. 1971. *A theory of justice*. Cambridge, MA: Belknap Press.

Restivo, S. 1992. *Mathematics in society and history*. Dordrecht: Kluwer.

Reynolds, C. J. 1976. Buddhist cosmography in Thai history with special reference to nineteenth-century culture change. *Journal of Asian studies* 35: 203-220.

Ritsert, J. 1990. *Models and concepts of ideology*. Atlanta, GA: Rodopi.

Rizzolatti, G. et al. 2002. Motor and cognitive functions of the ventral premotor cortex. *Current opinions in neurobiology* 12: 149-154.

Roberts, T. 2012. You do the maths: rules, extensions, and cognitive responsibil-ity. *Philosophical Explorations* 15(2): 133-145.

Rogers, Jr., H. 1987. *Theory of recursive functions and effective computability*. Cambridge, MA: MIT Press.

Rorty, R. 1989. *Contingency, irony and solidarity*. Cambridge: Cambridge UP.

____. 1995. Is truth a goal of enquiry? Davidson vs. Wright. *The philosophical quarterly* 45/180: 281-300.

____. 1997. What do you do when they call you a 'relativist'? *Philosophy and phenomenological research* 57/1: 173-77.

____. 1998. *Truth and progress: Philosophical papers, vol. 3*. Cambridge: Cambridge UP.

Rouse, J. 1987. *Knowledge and power: Toward a political philosophy of science*. Ithaca and London: Cornell UP.

Rupert, R. 2005. Minding one's cognitive systems: When does a group of minds constitute a single cognitive unit? *Episteme: a journal of social epistemology* 1: 177-88.

Russell, R., Duchaine, B., & Nakayama, K. 2009. Super-recognizers: people with extraordinary face recognition ability. *Pschonomic bulletin & review* 16/2: 252-57.

Rust, N. C. & Stocker, A. A. 2010. Ambiguity and invariance: two fundamental challenges for visual processing. *Current opinion in neurobiology* 20: 382-388.

Ryle, G. 1962. Abstractions. *Dialogue (Canadian philosophical review)* Reprinted in *Collected papers* 2: 435-445.

Schaffer, J. 2001. Knowledge, relevant alternatives and missed clues. *Analysis* 61/3: 202-208.

____. 2004. From contextualism to contrastivism. *Philosophical studies* 119: 73-103.

____. 2005. What shifts?: thresholds, standards, or alternatives? In G. Preyer & G. Peter (eds.) *Contextualism in philosophy: knowledge, meaning, and truth*. Oxford: Oxford UP.

____. 2007. Knowing the answer. *Philosophy and phenomenological research* 75/2: 383-403.

Scharff, R. C. & Dusek, V. 2003. *Philosophy of technology: the technological condition: an anthology*. Wiley-Blackwell.

Schlumberger, Ltd. 1982. *Reservoir and production fundamentals*. New York: Schlumberger, Limited.

Schmitt, F. F. 1994a. The justification of group beliefs. In Schmitt, F. F. 1994c: 257–87.

____. 1994b. Socializing epistemology: An introduction. In Schmitt, F. F. 1994c: 1-27.

____. (ed.) 1994c. *Socializing epistemology: The social dimensions of knowledge*. Rowman and Littlefield.

Schütz, A. 1962. *Collected papers vol. 1: the problem of social reality*. The Hague: Martinus Nijhoff.

Schütz, A. & Luckmann, T. 1974. The structures of the life-world: Volume 1 (Zaner, R. M., Engelhardt, H. T., Trans.). London: Heinemann.

Schyfter, P. 2009. The bootstrapped artefact: a collectivist account of technological ontology, functions, and normativity. *Studies in history and philosophy of science* 40: 102-111.

Scientific Drilling International, Inc. 2006. *Borehole survey information guide*. Paso Robles, CA: Applied Technologies Associates and Scientific Drilling International.

____. 2009. *Basic survey manual*. Paso Robles, CA: Scientific Dilling International, Inc.

____. 2010. *Production logging*. Paso Robles, CA: Scientific Drilling International, Inc.

Searle, J. R. 1984. *Minds, brains, and science*. Cambridge, MA: Harvard UP.

____. 2002. Why I am not a property dualist. *Journal of consciousness studies* 9/12: 57-64.

Serra, O. 1984. *Fundamentals of well-log interpretation*. Amsterdam: Elsevier.

Shapin, S. & Schaffer, S. 1985. *Leviathan and the air pump: Hobbes, Boyle and the experimental life*. Princeton: Princeton UP.

Shaw, R. 1991. *Computer-aided marketing and selling: information asset management*. Oxford: Butterworth-Heinemann Ltd.

Shumaker, R. W., Walkup, K. R., & Beck, B. B. 2011. *Animal tool behavior: the use and manufacture of tools by animals*. Baltimore, MD: John Hopkins UP.

Smith, B. & Welty, C. 2001. Ontology: Towards a new synthesis. In Welty, C. & Smith, B. (eds.) *Formal ontology in information systems*. Ongunquit, Maine: ACM Press: iii-x.

Smolen, J. 1996. *Cased hole and production log evaluation*. Tulsa, OK: Pennwell.

Sokal, A. & Bricmont, J. 1998. *Intellectual impostures: postmodern philosophers' abuse of science*. London: Profile Books.

Sosa, E. 2011. *Knowing full well*. Princeton UP.

Stanley, J. 2005. *Knowledge and practical interests*. Oxford: Oxford UP.

____. & Williamson, T. 2001. Knowing how. *Journal of philosophy* 98: 411-44.

Stenzel, J. 2011. *CIO best practices: Enabling strategic value with information technology*. Hoboken, NJ: John Wiley & Sons, Inc.

Steup, M. (ed.). 2001. *Knowledge, truth, and duty: essays on epistemic justification, responsibility, and virtue*. Oxford: Oxford UP.

Stone, J. 2007. Contextualism and warranted assertion. *Pacific philosophical quarterly* 88 (1): 92-113.

Stopka, P. & Macdonald, D. W. 2003. Way-marking behaviour: an aid to spatial navigation in the wood mouse (Apodemus sylvaticus). *BMC* *ecology* 3. Retrieved from http://www.biomedcentral.com/content/pdf/1472-6785-3-3.pdf.

Tebbich, S. & Bshary, R. 2004. Cognitive abilities related to tool use in the woodpecker finch, Cactospiza pallida. *Animal behaviour* 67/4: 689-697.

Thomasson, A. L. 2007. Artifacts and human concepts. In S. Laurence & E. Margolis (eds.) *Creations of the mind: Essays on Artifacts and their representations.* Oxford: Oxford UP: 52-73.

_____. 2008. Artifacts and mind-independence: Comments on Lynne Rudder Baker's 'Shrinking difference between artifacts and natural objects'. *APA newsletter on philosophy and computers* 8/1: 25-26.

Tollefsen, D. 2002. Organizations as true believers. *Journal of social philosophy* 33/3: 395-410.

_____. 2007. Group testimony. *Social epistemology* 21: 299-311.

_____. 2009. Wikipedia and the epistemology of testimony. *Episteme* 6/1: 8-24.

Triplett, T. 1986. Relativism and the sociology of mathematics: remarks on Bloor, Flew and Frege. *Inquiry* 29/1-4: 439-50.

Trist, E. 1981. *The evolution of socio-technical systems: a conceptual framework and an action research program*. Ontario: Ontario Ministry of Labour, Ontario Quality of Working Life Centre.

Tummolini, L. & Castelfranchi, C. 2006. The cognitive and behavioural mediation of institutions: Towards an account of institutional actions. *Cognitive systems research* 7 (2-3): 307-323.

Tuomela, R. 1992. Group beliefs. *Synthese* 91: 285–318.

_____. 2004. Group knowledge analyzed. *Episteme* 1: 109–127.

Unger, P. 1975. *Ignorance: A case for scepticism*. Oxford: Clarendon Press.

_____. 1984. *Philosophical relativity*. Oxford: Basil Blackwell.

Vaesen, K. 2011a. Knowledge without credit, exhibit 4: extended cognition. *Synthese* 181: 515-529.

_____. 2011b. The functional bias of the dual nature of technical artefacts program. *Studies in history and philosophy of science* 42 (1): 190-197.

_____. 2012. The cognitive bases of human tool use. *Behav. Brain Sci.* 35 (4): 203-218.

Vaesen, K. & van Amerongen, M. 2008. Optimality vs. intent: Limitations of Dennett's artefact hermeneutics. *Philosophical psychology* 21/6: 779-797.

Van Fraasen, B. C. 1980. The scientific image. Oxford: Clarendon Press.

Van Lawick-Goodall, J. 1970. Tool-using in primates and other vertebrates. In D. S. Lehrman, R. A. Hinde, & E. Shaw (eds.) *Advances in the study of behaviour vol. 3*. New York: Academic Press: 195-249.

Van Schaik, C. P., Deaner, R. O., & Merrill, M. Y. 1999. The conditions for tool use in primates: implications for the evolution of material culture. *Journal of human evolution* 36: 719-741.

Vermaas, P. E. & Garbacz, P. 2009. Functional decompositions and mereology in engineering. In A. W. M. Meijers (ed.) *Philosophy of technology and engineering sciences*. Amsterdam: Elsevier: 235-271.

____. & Houkes, W. 2003. Ascribing functions to technical artefacts: A challenge to etiological accounts of functions. *British journal for the philosophy of science* 54: 262-89.

____. & ____. 2006. Technical functions: A draw-bridge between the intentional and structural natures of technical artefacts. *Studies in the history and philosophy of science* 37: 5-18.

Vevers, G. M. & Weiner, J. S. 1963. Use of a tool by a captive capuchin monkey (Cebus appela). *Symposia of the zoological society of London* 10: 115-117.

Vincenti, W. G. 1990. *What engineers know and how they know it: Analytical studies from aeronautical history*. Baltimore, MD: John Hopkins UP.

____. 1992. Engineering knowledge, type of design, and level of hierarchy: Further thoughts about *What engineers know…* In P. A. Kroes & M. Bakker (eds.) *Technological development and science in the industrial age*. Dordrecht, Kluwer Academic.

____. & Bloor, D. 2003. Boundaries, contingencies and rigor: Thoughts on mathematics prompted by a case study in transonic aerodynamics. *Social studies of science* 33/4: 469-507.

Von Kármán, T. 1967. Creativity is a task, not a trait. *Machine design*, May 25.

____. 1970. Quoted in *Engineering education*, vol. 60. American society for engineering education: 467.

Weinberg, J., Nichols, S., & Stich, S. 2001. Normativity and epistemic intuitions. *Philosophical topics* 29: 429-460.

Williams, B. 2002. *Truth and truthfulness: an essay in genealogy*. Princeton: Princeton UP.

Williamson, T. 2001. Comments on Michael Williams' 'Contextualism, externalism and epistemic standards.' *Philosophical studies* 103: 25-33.

Winch, P. 1958. *The idea of a social science*. London: Routledge & Kegan Paul.

____. 1964. Understanding a primitive society. *American philosophical quarterly* 1: 307-324.

Wittgenstein, L. 1953. *Philosophical investigations*. Oxford: Blackwell.

____. 1961. In *Notebooks: 1914-16 by Ludwig Wittgenstein*. G. Anscombe & G.von Wright (eds.) Chicago: University of Chicago Press.

____. 1979. *On certainty* (2nd edition). G. E. M. Anscombe & G. H. von (eds.), D. Paul & G. E. M. Anscombe (trans.). Oxford: Basil Blackwell.

Wolpert, D. M. & Ghahramani, Z. 2006. Bayes rule in perception, action and cognition.In Gregory, R. L. (ed.) *Oxford Companion to the Mind*. Oxford: Oxford UP.

Wright, C. 1986. Does Philosophical Investigations I. 258-60 suggest a cogent argument against private language? In P. Pettit & J. McDowell (eds.) *Subject, Thought and Context*. Oxford: Oxford UP: 209-266.

Yourgrau, P. 1983. Knowledge and relevant alternatives. *Synthese* 55: 175-190.

Zelić, B. & Stahl, B. C. 2005. Does ontology influence technological projects? The case of Irish electronic voting. In *Lecture notes in computer science* 3782: 657-667.

Zins, C. 2007. Conceptual approaches for defining data, information, and knowledge. *Journal of the American society for information science and technology* 58/4: 479-93.

Acknowledgements

First and foremost I would like to thank my supervisors Duncan Pritchard and Jesper Kallestrup for reading countless iterations of this dissertation and always commenting with precision and clarity. I would also like to thank Scientific Drilling International, Halliburton, Chevron, and the wonderful people that work there for allowing me access to their buildings and operations and for always humouring me by answering what must have seemed at times like the questions of a very strange academic. For feedback on particular chapters of my thesis I'd like to thank Barry Barnes, Martijn Blaauw, David Bloor, Adam Carter, Matteo Colombo, Luciano Floridi, Soraj Hongladarom, Spyros Orestis Palermos. I would also like to thank anonymous editors at Minds & Machines, Social Epistemology, and Synthese for pointing out some necessary corrections.