RePoSS: Research Publications on Science Studies

RePoSS #28:

Philosophical Perspectives on Interdisciplinary Science Education: An Introduction

Mads Goddiksen

November 2014

Centre for Science Studies, University of Aarhus, Denmark
Research group: History and philosophy of science

Please cite this work as:

Mads Goddiksen (2014). *Philosophical Perspectives on Interdisciplinary Science Education: An Introduction*. RePoSS: Research Publications on Science Studies 28. Aarhus: Centre for Science Studies, University of Aarhus. URL: http://www.css.au.dk/reposs.

Philosophical Perspectives on Interdisciplinary Science Education: An Introduction

Mads Goddiksen, Centre for Science Studies, Aarhus University.
E-mail: madsgoddiksen@yahoo.dk

Preface

The following served as the introductory chapter to my Ph.D.-dissertation entitled *Philosophical Perspectives on Interdisciplinary Science Education: Characterizing Important Expertises Though a Practice Oriented Analysis of Integration and Explanation*. The dissertation was presented to the faculty of science and technology at Aarhus University and successfully defended in October 2014. I have only made minor linguistic changes in the present version of the introductory chapter compared to the original text.

The dissertation as a whole consisted of an introductory chapters followed by four research papers and two appendices. This introductory chapter contains numerous references to the later chapters and appendices. Public versions of these texts can be found in the following places:

- "Chapter 2" is: Goddiksen, M. (2014): "Clarifying interactional and contributory expertise", *Studies in History and Philosophy of Science,* vol. 47, pp. 111-117.
- "Chapter 3" is: Goddiksen, M. & Andersen, H. (in preparation): "Expertise in Interdisciplinary Science and Education". The chapter is publically available here: http://philsci-archive.pitt.edu/11151/
- "Chapter 4" is: Goddiksen, M. (Forthcoming): "An empirical Method for the Study of Exemplar Explanations" in Andersen, H., Wagenknecht, S. & Nercessian, N. (eds.): *Empirical Philosophy of Science: Introducing Qualitative Methods into Philosophy of Science*, Springer.
- "Chapter 5" is: Goddiksen, M. (in preparation): "Explanatory Standards in Physics and Biology Textbooks: The Case of Polymers". The chapter is publically available here: http://philsci-archive.pitt.edu/11150/
- "Appendix A" is: Goddiksen, M. (2013): "Interdisciplinarity and explanation: How interdisciplinary education calls for a new approach to research on scientific explanations" in *Proceedings of the 12th biennial IHPST conference*. Available here: http://archive.ihpst.net/2013-pittsburgh/conference-proceedings/
- "Appendix B" is: 'Philosophy-of-Science in Practice vs. Philosophy of Science-in-Practice'. Co-authored with Hanne Andersen, Sara Green, Brian Hepburn & Susann Wagenknecht. Printed in *SPSP Newsletter* Autumn 2012.

Contents

1 Aims and outline

Interdisciplinary problem solving is an elusive concept. Its specific meanings are hard to pin down, and one definition covering all current use is probably unattainable. And yet, interdisciplinary problem solving as a concept is central in many descriptions of the vision for the future of research and innovation. Influential voices among politicians, funding agencies, private businesses and practicing researchers emphasize that in the future, research and innovation should focus increasingly on solving problems through interdisciplinary collaborations (see sec 2). The vision for the future of research and innovation affects the practice of present day education. If interdisciplinary problem solving is the future of research and innovation then the educational system better start preparing future practitioners for that reality. The process has already begun. High schools and universities are changing their educations in an attempt to prepare their students as much as possible for a professional life that is expected to be filled with demands for interdisciplinary problem solving. A striking example of this can be found in the science faculties of the world's universities and other higher education institutions. Alongside the traditional science degree programs – in physics, chemistry, geology etc. – new interdisciplinary programs are emerging. Many institutions now offer degree programs in for instance nanotechnology, environmental science or biotechnology. The specific aims and contents of these interdisciplinary programs vary from institution to institution. Some interdisciplinary programs are interdisciplinary only by name, whereas their aims and content can be very hard to distinguish from other kinds of science degree programs. Others aim specifically to train their students in interdisciplinary problem solving drawing on specific sciences, and the contents of the program have been developed accordingly. This dissertation deals with this latter kind of interdisciplinary science programs. The specific aims and argument structure of the dissertation is outlined in the remainder of this section 1, and elaborated as the introduction progresses.

1.1 The general structure

One question facing educators involved in interdisciplinary science programs is whether the contents of their program can be adapted to further help their students reach the predefined teaching goal: to enable the students to do interdisciplinary problem solving. This important question may be very difficult to answer, partly because it can be very hard to measure whether a given adaptation of the program's content works.

The standard way of testing whether a change in the content of an educational program – be it a change in the materials used or in the instructional method – is to try out the new content on a group of students selected through reasonable criteria and compare their progress (relative to the teaching goal) to that of a control group (Fraenkel, Wallen & Hyun 2012). In the case where the teaching goal is the ability to perform interdisciplinary problem solving the challenge of constructing such a test is dual. The first challenge is the illusiveness of the teaching goal: What is interdisciplinary problem solving (besides important)? It is very hard to measure, let alone teach, something which educators have only a vague characterization of. The second challenge arises once educators have an idea about what they are trying to measure: What are the relevant indicators for the ability to engage in interdisciplinary problem solving (Boix Mansilla, Duraisingh 2007)?

This dissertation relates to both of these questions although it does not provide a complete answer. The first main aim of this dissertation is to present a characterization of interdisciplinary problem solving, and based on practitioners experiences *highlight specific skills that, while not essential for engaging in traditional disciplinary problem solving, are nevertheless highly valuable when engaging in interdisciplinary problem solving.* The problem of how to test for these skills is not addressed in any detail.

Larson and collaborators took on a similar task when they set out to characterize the specific skills needed in interdisciplinary healthcare. Their paper starts from the following argument referring among others to the review in (NAS 2004):

> ***The assumption is sometimes made that anyone can engage in interdisciplinary research should they choose to do so, but it is clear from a considerable body of literature that interdisciplinary efforts require mastery of specific competencies […]. It follows, then, that if such competencies are explicated, it might be possible to enhance researcher's abilities to participate in and conduct interdisciplinary scholarship (Larson, Landers & Begg 2011, p. 38)***

The account of Larson and collaborators focusing on healthcare is consistent with my account (presented in chapter 3) focusing on the sciences. Their account is also to some extent complementary to mine. The special challenges facing professionals engaging in interdisciplinary problem solving can be divided into two general categories: social or *institutional* challenges and *epistemological* challenges. Larson and collaborators also discuss the institutional challenges related to initiate, fund, publish and sustain interdisciplinary problem solving. These institutional challenges are real and significant but will not be in focus is this dissertation[1]. My account focuses on giving a more detailed characterization of the special skills needed to overcome the epistemological challenges.

The epistemological challenges facing practitioners can be further subdivided into two categories: Challenges related to communication across disciplines and challenges related to the integration of cognitive resources – knowledge, models, concepts, methods etc. – from different disciplines (see chapter 3).

Integration of cognitive resources is commonly taken to be the central characteristic of interdisciplinary problem solving (Klein 2010). It is perhaps also the most elusive part of the interdisciplinary problem solving process. Moving towards a characterization of interdisciplinary problem solving – and thus towards the first aim of the dissertation - will thus have to go via a characterization of integration (see section 3.3). There are further benefits of having such a characterization. Closing in on what integration is also points to a set of skills that are valuable when aiming to perform integration; what we (my co-author Hanne Andersen and I) in chapter 3 shall characterize as *integrational expertise*[2]. As argued in chapter 3 the process of integration involves negotiating differences in the standards used to judge the relevance of problems and the quality of potential problem solutions within individual disciplines. I use the term *epistemological standards* throughout to designate these standards. Negotiating differences in epistemological standards means grappling with questions that are largely philosophical in nature (Eigenbrode et al. 2007). Thus, looking closer at

[1] For discussions see (Felt et al. 2013; NAS 2004; Campbell 2005; Rhoten, Parker 2004).
[2] In this context 'expertise' simply designates a set of skills that enable a person to perform a certain task (see chapter 2)

interdisciplinary problem solving - a central teaching goal of interdisciplinary programs – leads to the conclusion that to achieve this goal, students will benefit not only from training in grappling with scientific questions but also with questions of a more philosophical nature. Reaching this conclusion suggests an adaptation of the contents of interdisciplinary science programs: Just as scientific research is drawn on to enable the students to grapple with scientific questions, so too should research in the history, philosophy and sociology of science (henceforth HPSS) be drawn upon in an effort to prepare students for grappling with the more philosophical questions. How this is best done is ultimately an empirical question that requires suitable testing methods. I shall not address this question here. Instead I, as a trained philosopher of science, turn the gaze inwards. The reality of interdisciplinary programs has created a demand for relevant philosophical research that can somehow be drawn upon in the teaching of interdisciplinary problem solving. Can philosophers fulfil this demand? In many important respects I think the answer is yes. However, there are also shortcomings, specifically when it comes to our knowledge about the detailed differences in epistemological standards across disciplines. The practice of interdisciplinary teaching has thus revealed a hole in the philosophical literature. How can this hole be filled? Providing a partial answer to this question constitutes the second aim of this dissertation. More specifically I aim to address this question by *exemplifying how a practice oriented approach to philosophy of science can be used to characterize differences in explanatory standards across scientific disciplines.* Explanatory standards - i.e. the standards used in a given discipline to judge the quality of explanations - are a subset of the discipline's epistemological standards that can readily be studied using empirical methods (chapter 4). They are emphasized in the literature as an important subset of the epistemological standards of a discipline in relation to interdisciplinary problem solving ((Green, Fagan & Jaeger 2014; O'Rourke, Crowley 2013; Love 2008; Hepburn, Thorén In preparation) see also appendix A).

I move towards the second aim of the dissertation in three steps. I first characterize a general approach to doing philosophy of science: The *practice oriented* approach. This, I argue in section 7, is a suitable approach to studying differences in epistemological standards if a partial aim is that the studies should be relevant to educators. Within the practice oriented approach several specific methodologies suggest themselves. I have chosen to develop a general methodology based on textbook analysis for comparing explanatory standards across disciplines (discussed in detail in chapter 4 and 5). In the third and final step I apply this methodology to a concrete case as I compare the explanatory standards of molecular biology and polymer physics (chapter 5). These two fields are especially interesting to compare as interdisciplinary problem solving efforts aiming to integrate branches of physics and biology are common. However, the process is not without controversy (e.g. (Morrison 2009; Green 2014)). I explore some of the source of this controversy as I compare explanatory standards of molecular biology and polymer physics, two fields that meet for instance in the field of nanoscience.

1.2 Outline

The preceding paragraphs sketched the general aims of this dissertation, my argument for why the two aims are intimately connected and my overall strategy for reaching these aims. Knowing precisely what you aim for is an important step towards reaching your aims, and the sketch provided above will not suffice. My first task will therefore be to further clarify the central concepts introduced above and the relations between them. I do this in the remaining sections of this introduction chapter and continue in chapter 2. Section 2 returns to the vision that interdisciplinary problem solving is the future of research and innovation. I show that this vision has profound influence on how research is funded and organized. I also point out that this is not uncontroversial. The controversy and the difficulties with teaching interdisciplinary problem solving share a common source: The missing consensus on what interdisciplinary problem solving is. Section 3 surveys the discourse on interdisciplinary problem solving: how is interdisciplinary problem solving commonly characterized in the philosophical and educational literature? Furthermore, the section clarifies what I take interdisciplinary problem solving to mean. Section 4 returns to the notion of an interdisciplinary program in higher education. Which educational programs can be considered interdisciplinary given the characterization of interdisciplinary problem solving provided in section 1.3? Answering this question will help determine which educational programs my arguments are relevant to. Section 5 adds some more detail to the claim that interdisciplinary problem solving means grappling with questions of a philosophical nature. Section 6 reflects on the relevance of drawing on HPSS in teaching in interdisciplinary programs, and how the aim to address philosophical questions relevant for interdisciplinary science education affects the way philosophy is done. Sections 7 and 8 provide additional background and depth to these reflections. Section .7 characterizes the ongoing practice turn in philosophy of science, and argues that a practice oriented approach is well suited for addressing questions relevant to science education. Section 8 focuses specifically of explanations. I review the philosophical literature on scientific explanations and clarify what I take explanations to be.

The remainder of the introduction thus delineates the project, clarifies important concepts and places the project in the broader research context. Chapter 2 continues along this line. At the core of the overall argument of this dissertation is the claim that interdisciplinary programs should aim to develop a certain kind of *expertise* in the students – namely integrational expertise. This brings the project in contact with a growing literature on experts and expertise characterized by Harry Collins and Robert Evans as the *third wave of science studies* (Collins, Evans 2002). Chapter 2 positions the project in this literature, and points to important consequences that this project has for the research program on expertise started by Collins and Evans.

By the end of chapter 2 the central concepts should be sufficiently clarified, and it is time to start attacking the central aims of the dissertation more directly. I begin in chapter 3 by considering the first main aim. Based on practitioners experience, common challenges faced in interdisciplinary problem solving are characterized along with the kinds of expertise relevant for overcoming them. In particular integrational expertise is characterized in detail. HPSS is highlighted as a valuable resource when teaching integrational expertise. In chapter Chapter 4 I turn to the second main aim, and present a general methodology for comparing explanatory standards across disciplines. Chapter 5 adapts and applies the general methodology in a comparison of the explanatory standards of statistical physics and molecular biology. I show that the two domains have differing ontologies of change and promote different modes of explanatory reasoning. Furthermore, I point to differences in the kinds of explanation-seeking questions deemed relevant in the two domains.

In the "Future perspectives" chapter following the four core chapters I consider the new research questions that are raised by the work presented in this dissertation, and the directions that my own research will take in the future.

2 The vision of interdisciplinarity

"There are three reasons for doing interdisciplinary research: Funding, funding and funding." This conclusion, drawn by young research physicist during a private discussion, illustrates a tension in present day academia. For many practicing researchers there appears to be very little *scientific* incentive to go into interdisciplinary research[3]. On the other hand interdisciplinary research is so heavily promoted by funding agencies that scientists must do it anyway (or at least pretend they do). In this section I elaborate on this tension and show that it arises partly because it is not obvious to everyone how interdisciplinary research differs from ordinary research and why this difference is for the better[4]. This situation is not only problematic in relation to current research, but also in relation to education. It is often emphasized that the contents of educational programs will have to change in order to prepare students for the interdisciplinary future of research and innovation (for early arguments see (Jantsch 1972; Rohrer 1993), for more recent examples see (Roco, Bainbridge 2003; NAS 2004)). But how specifically educations are to be changed can be difficult to see if it is not clear how interdisciplinary science differs from disciplinary science.

Innovative research that is not particularly concerned with disciplinary borders is strongly promoted by politicians and some researchers. This type of research is often called interdisciplinary[5]. The political emphasis has led to major changes in the way science is funded and, partly as a consequence, changes in the way research is structured. To mention just two examples of the former I briefly consider the European *Horizon 2020* program and the *National Nanotechnology Initiative* in the USA.

The EU's *Horizon 2020* program will distribute nearly €80 billion to research and innovation between 2014 and 2020. Rather than being structured around the traditional disciplines, *Horizon 2020* is structured around a set of predefined "societal challenges" that applicants must somehow address. Many of these challenges are understood to require collaboration and integration across disciplines in order to be solved (European Commision 2014). So, by focusing on societal challenges the *Horizon 2020* program is, indirectly, putting 80 billion euro behind a vision, if only a vague one, that Europe needs more interdisciplinary research.

The American analogy to the European focus on societal challenges is the structuring of research funding around the "converging" or *NBIC* technologies – that is, nano-, bio-, and information technology and cognitive science (Roco, Bainbridge 2003). Again, it is understood that developing these new technologies will *require* interdisciplinary research, at least within the science and engineering disciplines.

The most prominent example of the promotion of the NBIC technologies is the National Nanotechnology Initiative (NNI) launched by President Bill Clinton in 2000. The NNI has had a huge and steadily growing budget ever since, and in 2014 alone it will distribute nearly $1.5 billion to research and education in nanotechnology[6].

---

[3] Researchers and funding agencies mainly talk about interdisciplinary *research and innovation*. I have largely remained true to these terms in this section. However, I find it more enlightening to use the umbrella term 'interdisciplinary problem solving' and will do so consistently in the following sections.
[4] My analysis of interdisciplinarity is conceptual. For a more historical account see (Klein 1990).
[5] In fact, Weingart (2000) notes that in many contexts 'interdisciplinary' is simply used as a synonym for 'innovative'.
[6] http://www.nano.gov/about-nni/what/funding [accessed 24/6/14]

With such vast sums of funding available for interdisciplinary research it is no wonder that university leaders are reconsidering the aims and structure of their institutions in order to promote interdisciplinary research. A striking example is Aarhus University which underwent a major restructuring process from 2010 to 2012. Under the heading "deeper connections" the entire financial and organizational structure of the university was restructured in order to encourage collaborations across disciplinary boarders[7]. Among other things, this meant that a number of previously independent institutions were made part of the university, it also meant that the number of faculties was reduced from 9 to 4 and that the number of departments was reduced from 55 to 26. (The most debated merge involved the faculty of humanities with 9 departments, the faculty of theology and the Danish Pedagogical University each with a number of separate departments, which were all merged into the new faculty *ARTS* with just 3 departments.) This reduction in the number of institutions, faculties and departments was thought to break down major barriers to interdisciplinary research. New "interdisciplinary" centers were also established, for instance a center for neuroscience, where researchers from different disciplines could come together (literally) under the same roof and collaborate.

For various reasons this restructuring process was heavily criticized by both students and employees at Aarhus university. In this context it is worth noting that a major point of criticism has the idea that more interdisciplinary collaborations will automatically arise when larger and/or different units are created. This critique indicates that researchers perceive other, more important factors than the administrative organization of the universities as discouraging interdisciplinary research. What could such inhibiting factors be? Why might the reorganizing of research institutions be insufficient for promoting interdisciplinary research? Environmental scientists Lélé and Norgaard provide a partial answer:

> ***The main barrier to interdisciplinary work – for example collaboration between a botanist and a soil scientist – lies in the relative absence of motivation. […] Most scientists do not see the low level of cross-disciplinary collaboration as a problem. Most are happily addressing the questions that have already been identified within disciplinary boundaries, in the belief that pushing the frontiers of each discipline will eventually lead to the convergence of all knowledge. Crossing boundaries […] distracts from pure research, where academic prestige is still highest. Some funding agencies […] increasingly support applied research in their attempt to address pressing problems. Nevertheless, we believe that the motivation for crossing disciplinary boundaries even within the natural sciences remains generally low. (Lélé, Norgaard 2005, p. 969)***

This quote has much the same flavor as the one presented in the beginning of this section: Many scientists find that the only incentive for going into interdisciplinary research is economical. Importantly, the reward structures distributing academic prestige still favor traditional disciplinary research. This means that especially young researchers are discouraged from going into interdisciplinary research. This is an important part of the institutional challenges facing the interdisciplinary researcher mentioned in the previous section. Furthermore,

[7]For details (in Danish) see (AU 2011).

the quotation indicates that promoters of interdisciplinary research have not been sufficiently good at conveying an *epistemological* argument for why interdisciplinary research is valuable. That is, it is not clear what is special about interdisciplinary research and why this difference means that interdisciplinary research is needed to achieve the aims of research.

Developing such an epistemological argument for the value of interdisciplinary research in full detail would be a suitable topic for a separate dissertation, so I will not pursue it in detail. In relation to education this means that I will not go into the discussion of whether the educational system *should* aim to prepare students for interdisciplinary problem solving. Instead, I take as a starting point the aim to promote interdisciplinary problem solving and ask how it affects the contents of interdisciplinary programs. The preceding paragraphs show not only that this is an important question to answer, but also illustrates why it can be difficult to answer: Those who promote interdisciplinary research and innovation have not been very clear about what exactly it is that they are promoting. In the next section I therefore take a closer look at the many ways interdisciplinary problem solving has been described, and clarify what I take it to be.

3 Interdisciplinary problem solving

Interdisciplinary problem solving is commonly characterized either by the kind of problems it addresses, the kind of solutions it aims for or both. I argue that we need to look at both in order to reach an acceptable characterization. In section 3.1 I consider interdisciplinary problems. Interdisciplinary solutions – or *integrated* solutions - are considered in section 3.3, after a necessary intermezzo in section 3.2 clarifying the central concept 'discipline'.

3.1 Interdisciplinary problems

It is often useful to start the analysis of a given concept by contrasting it to other related concepts. In the case of interdisciplinary problems a relevant place to start is to discuss how they are different from disciplinary problems. As illustrated by the Horizon 2020 and NNI documents there is a pervading idea in the discourse on interdisciplinary research that some problems are inherently interdisciplinary, because they somehow require an interdisciplinary approach to be solved (see also (Brigandt 2010)). Other problems are understood to be inherently disciplinary because they can be solved entirely within one discipline. But why is it that some problems seem to *require* an interdisciplinary approach in order to be solved? Answering this question will give some relevant insights into the nature of integration, and ultimately into the skills needed to perform interdisciplinary problem solving[8].

[8]Besides this practical benefit, answering this question is also important in order to make the characterization of interdisciplinary problem solving non-circular. If interdisciplinary problem solving is the process of solving interdisciplinary problems, and interdisciplinary problems are simply problems that require interdisciplinary problem solving to be solved then the characterization is not really enlightening.

The differences between disciplinary and interdisciplinary problems can be illustrated by considering the following questions[9]:

1. Do magnetic monopoles exist?
2. How can we counter the effects of the changing climate?

The first question is an example of what Kuhn called a *puzzle* (Kuhn 1996, pp. 36-39). When contrasting interdisciplinary problems with disciplinary problems, it is often something like Kuhnian puzzles that is referred to when talking about disciplinary problems. The main characteristics of Kuhnian puzzles are that they arise and can be solved entirely with one and the same discipline – or *community* as Kuhn preferred to call them. Question number one illustrates this, at least if we consider the period up until about 1975.[10]

When particle physicists looked deep into their theories about the fundamental particles in the universe they found that these theories allow for the existence of magnetic monopoles (just as they allow for the existence of electrons, positrons and many other particles). Furthermore, when they looked to the theories of electro-magnetism they found that these would be a lot 'prettier' – i.e. more symmetric – if the existence of magnetic monopoles could be established. And yet, to the best of our empirical knowledge, magnetic monopoles do not exist. This raised a puzzle within particle physics: is the reason why magnetic monopoles have never been detected – in spite of many attempts – that they simply do not exist or is the reason that experimentalists have not yet been clever enough to catch one? Puzzles like these arise within a given community, because they have chosen to accept a specific theory. Should they choose to adopt another theory the puzzle might disappear entirely (Kuhn 1996, p. 37)[11]. Furthermore, there were, within the paradigm of particle physics, well established methods for detecting new particles and rigorous standards for determining when a detection has occurred. In this sense the puzzle not only arose within particle physics, it was also solvable entirely within the paradigm of particle physics. Most will therefore agree that question number one was not an interdisciplinary problem during this period, but a disciplinary one[12]. Note, however, that it requires assistance from non-physicist engineers to build and carry out the actual experiment that could detect magnetic monopoles. The fact that solving a problem requires specialized expertise from a number of disciplines does not, therefore, in itself distinguish interdisciplinary problems from disciplinary problems. Problem solving, interdisciplinary or not, often requires the coordination of contributions from multiple specialists.

[9] As most problems can be formulated as questions I use 'problem' and 'question' interchangeably in the following.
[10] After 1975 the status of the monopole problem is more ambiguous. For an illuminating discussion see (Kragh 1981).
[11] Similarly, earlier theories about the fundamental building blocks of the universe claimed the existence of an aether, which gave rise to a number of puzzles about the aether. As General Relativity, which does not postulate the existence of an aether, became accepted in the physics community, these puzzles about the aether ceased to be puzzles.
[12] Those who do not, and still want to maintain the distinction between disciplinary and interdisciplinary problems, need to counter the *reductio* argument that if the monopole problem was an interdisciplinary problem in this period then very many, if not most, of the problems that physicists have solved over the years have been interdisciplinary, even when they were working under conditions very close to what Kuhn called normal science. Admitting this makes it very hard to provide a useful definition of disciplinary problems and thus to maintain the distinction between disciplinary and interdisciplinary problems.

Coordination of contributions requires at least some degree of communication. The closer coordination one is aiming for the more, and the better, the communication needs to be. Integration, a characteristic of interdisciplinary problem solving (explored in sec. 3.3), requires a fairly high degree of coordination among the contributing experts, and so the communication among them needs to be rather fluent. As we shall see in chapter 3 establishing efficient communication among experts from different disciplines is often reported by practitioners as a major challenge when attempting interdisciplinary problem solving. The skills needed to overcome this kind of challenge – described in chapter 3 as *interactional expertise* (not to be confused with *integrational* expertise mentioned above) – are thus particularly useful in interdisciplinary problem solving, but can also be highly valuable in disciplinary problem solving. In this respect, students in interdisciplinary programs need to learn skills very similar to those of their colleagues in disciplinary programs, although the priority among the skills may be different (interactional expertise may be more important for students in interdisciplinary programs than for students in disciplinary programs). Looking to the features that set interdisciplinary problems apart from disciplinary problems will allow us to see which special skills students in interdisciplinary programs need to gain in addition to the disciplinary skills.

The question about how to deal with the changing climate (question 2. above) exemplifies an interdisciplinary problem. It is similar to question 1. about magnetic monopoles in the sense that it requires expertise from a number of different disciplines to be solved. It is different from question number one in (at least) three important respects: First, there is a difference in the kind of system that is being asked about. Contrary to magnetic monopoles the earth's climate is a *complex system* (see below) which implies that it is hard to study as independent parts. Second, whereas question one sat squarely within the research interest of one discipline (particle physics) question number two has already been addressed in part by a number of disciplines. Thirdly, there seems to be a more immediate societal interest in finding an answer to question 2. compared to question 1. In fact, question 1. has to some extent been imposed on science from the outside. As we look closer into these differences we find that especially the combination of the first two characteristics can be used to distinguish interdisciplinary problems from disciplinary problems. Societal relevance is often cited as a defining feature of interdisciplinary problems, but I will argue that this conflates matters that should be kept separate (sec 3.1.1).

Although neither magnetic monopoles nor the earth's climate can be decomposed into parts that can be studied independently this is so for very different reasons. Magnetic monopoles are elementary particles. They have no parts. The earth's climate has plenty of parts – oceans, forests, cities and the atmosphere to name a few – but these parts are intimately connected and interact in so intricate ways that the behavior of the system cannot be studied simply by looking at the individual parts separately. In this sense the earth's climate is a *complex system*. Complexity is almost unanimously cited as a defining feature of interdisciplinary problems (See reviews in (Klein 2010; NAS 2004; Repko 2008), see also (O'Rourke, Crowley 2013; Brigandt 2010; Gerson 2009; Love 2008; Hansson 1999)), although few spell out in detail what it means for a problem to be complex.

A first approximation would be to say that complex *problems* deal with complex *systems* and then turn to complex systems theory for a definition (Repko 2008, pp. 152)[13].

It is often implied that it is the complexity of a problem that alone "forces" those who try to solve it to adopt an interdisciplinary approach. I find it important to point out that this is not the case, otherwise it becomes difficult to understand what it means to solve an interdisciplinary problem, which, ultimately, is what we want, so that we might better teach others how to do it. Just as problems that require expertise from a number of disciplines to be solved are not necessarily interdisciplinary so too is dealing with a complex system not a sufficient condition for a problem to be interdisciplinary. The chaotic pendulum is a case in point. This is a complex system in the sense described above, but the study of chaotic pendulums still remains the task of physics. Why is this so? In this context, an important difference between the chaotic pendulum and the earth's climate seems to be that at the current point in history only one established discipline has taken an interest in studying any of the parts of the chaotic pendulum[14]. Quite the opposite is the case with the earth's climate. Here a number of established disciplines have taken an interest in at least one of the parts of the complex system. Furthermore, some of the questions that these disciplines ask about the complex system are *relevant sub-questions* to the overall problem of how to deal with the changing climate in one or both of the following ways[15]: The sub-questions may be broadly recognized as relevant across the disciplines involved or it may be that nature "pushes back" and encourages researchers to consider these questions in order to get the solution to the overall problem to live up to specific disciplinary standards, or standards that have been imposed externally (e.g. by politicians requiring that a solution is judged from both an economical and environmental perspective). This notion of relevance requires a bit of unpacking. Different disciplines often consider the same phenomena from different perspectives without this leading to interdisciplinary collaboration, simply because the two perspectives are not, at that particular point in time, considered relevant to each other. A group of physicists may be attempting to model the folding of certain proteins without paying any attention to the evolution of these proteins studied by groups of biologists. What can make these two groups collaborate? One option is that one of the groups adopts (either out of curiosity or because they have been asked to by outsiders) a general problem which they want to solve. By 'solve' I mean that they want to identify at least one potential solution that satisfies the epistemological standards adopted by the group. In the case where the problem has been adopted out of curiosity the relevant epistemological standards will be the epistemological

[13] Complex systems theory does not capture all kinds of social complexity, and so this approximation is not appropriate when considering complex social problems. Interdisciplinary problem solving focusing on complex social problems will have to draw heavily on the social sciences and humanities. I focus on the problems addressed in interdisciplinary problem solving drawing mainly on the natural sciences. The literature indicates that the approximation is more appropriate for this subset of interdisciplinary problems.

[14] History shows that it is far from given which disciplines should study which properties of which objects. The case of the age of the Earth illustrates this. Starting out as a theological problem, Lord Kelvin turned it into a physical problem. Later it turned into a geological problem. Currently the age of the Earth is estimated through studies of radioactive samples, so it's back to physics. (For a concise treatment of this case in Danish see (Kragh 2004), for details in English see (Albritton 1980)). So, the argument that the study of specific properties of certain objects belong "naturally" within a particular discipline cannot be made without heavy metaphysical assumptions.

[15] The idea that general problems contain sub-problems is similar to the idea presented by Love (2008) that general problems structure a *problem agenda*.

standards of the discipline which the group belongs to. If the problem has been imposed from outside, there may be additional constraints. Adopting a general problem usually leads to the formulation of sub-questions that someone needs to answer before the general problem is solved. However, which sub-questions are formulated is not a given[16], as it will depend on the interests, knowledge and expertise of the individual researchers involved. (If you are very good at determining the structure of proteins, then you are probably more likely to ask sub-questions involving proteins and their structure). Furthermore, what sub-questions are formulated will depend on the epistemological standards relevant to the general problem. If the general problem, for instance, is to explain a given phenomenon, then researchers will know that certain elements are expected to be present in such an explanation. These constraints therefore guide the formulation of sub-questions[17]. Research collaborations from different disciplines may thus appear to address the same general problem, but ask and focus on different sub-questions. If we consider a number of research collaborations from different disciplines addressing a similar general problem, and look at the sub-questions they ask (not just the ones they focus on), there will most likely be some degree of overlap. This set of sub-questions that have been asked by a number of disciplines will be relevant sub-questions to the overall problem in the first sense described above.

There will be many sub-questions to a general problem that are not widely recognized as relevant and many more that are not even formulated. However, it may turn out, as attempts to solve the overall problem are made, that a solution that satisfies the epistemological standards of the collaboration cannot be reached if some of these questions are not addressed[18]. In this way, nature can 'push back' on researchers and encourages them to reconsider their priorities of sub-questions to a general problem. Whether sub-questions are relevant or not is thus not entirely a matter of convention.

Relevant sub-questions need not be attempted answered by all who deem them relevant. Researchers may well realize that they should not attempt to answer specific sub-questions themselves because they think others can do it better; in fact others may have proposed an answer already. This may lead them to pass the sub-question on to members of other disciplines[19]. For instance, atmospheric scientists accept that it is important to study the impact of the world's forests on the global climate even if they would, in principle, be happy to rely on other disciplines to perform and evaluate these studies. I say "in principle" because in reality this will only be possible when considering non-complex systems. In complex systems like the earth's climate individual parts cannot be considered in isolation and so a study of the forests impact on the climate will unavoidably involve the atmosphere, and so questions about the world's forests will at least to some extent

[16] Both Love and Brigandt seem to recognise this in their accounts but offer no alternative mechanism as to how problem agendas are generated (Brigandt, Love 2012; Brigandt 2010; Love 2008).
[17] For a similar point see (Kuhn 1977).
[18] The history of the search for magnetic monopoles after 1975 illustrates this. The continued failed efforts to detect monopoles forced researchers to ask why no (or only very few) monopoles exist although their existence is allowed by theory. This relates the particle physics puzzle of whether monopoles currently exist to more cosmological questions about the formation of monopoles in the early universe. This does not necessarily mean that the monopole problem is now an interdisciplinary problem, but it is not a puzzle either.
[19] This is one type of what Thóren and Persson (2013) have described as *problem feeding*.

have to be judged according to the standards in place for answers to questions about the atmosphere and vice versa. The set of possible solutions to an interdisciplinary problem - like question 2. which deals with a complex system about which relevant sub-questions have been asked by multiple disciplines - is thus constrained by the epistemological standards of more than one discipline (Hepburn, Thorén In preparation).

So, it is the combination of dealing with a complex system and asking questions that are relevant across disciplines that seems to be the factor that can “force” practitioners to seek some degree of integration (see section 3.3). The need for interdisciplinarity thus arises partly because of the complexity of nature and society itself and partly because practitioners from different disciplines have differing expertise and differing epistemological standards. Furthermore, by looking into what characterizes interdisciplinary problems we have also begun to see what exactly it is that researchers are being “forced” to do, when facing interdisciplinary problems: they must navigate a solution space that is constrained by the epistemological standards of multiple disciplines. This makes the process of solving an interdisciplinary problem different from the process of solving a disciplinary problem, and thus suggests that interdisciplinary problem solving also requires special skills.

Summing up, this section has explored the characteristics of interdisciplinary problems. We have seen that a useful characterization of interdisciplinary problems is:

> An **interdisciplinary problem** concerns a complex system, contains relevant sub-questions about the complex system and the answers to these questions that cannot be evaluated solely within one discipline.

I shall use the term in roughly this sense throughout this dissertation (except that I following the discussion in section 3.2 will speak of domains rather than disciplines). We have seen that establishing efficient communication among specialists from different disciplines is a common challenge in both interdisciplinary and disciplinary problem solving, and that interdisciplinary problem solving furthermore involves the challenge to find a solution to a problem that is acceptable according to the epistemological standards of more than one discipline.

3.1.1 Societal relevance of interdisciplinary problems

In many contexts, interdisciplinary problem solving has become almost synonymous with problem solving that addresses pressing societal problems like cancer and climate change (cf. the discussion of *Horizon 2020* above). Some authors (e.g. (Repko 2008) see also (NAS 2004)) include societal relevance as a necessary condition in their definition of an interdisciplinary problem. I have not included it in my characterization, mainly because it would not capture the diversity of interdisciplinary problem solving practice. To see this, consider the following comparison.

In section 3.1. I contrasted the problem

2. How can we counter the effects of the changing climate?

with the problem of whether magnetic monopoles exist. Now consider the differences between the problem about climate change and the following problem:

3. How did the vertebrate jaw evolve?

Problem 3. may have arisen within a particular discipline (evolutionary biology) but as explained in detail by Love (2008), questions about the development of “evolutionary novelties” such as the vertebrate jaw or avian flight are not easily explained within evolutionary biology - that is, relevant sub-questions arise that currently fall under the domain of other disciplines like paleontology or molecular biology (see also (Brigandt, Love 2012; Brigandt 2010)). Since vertebrates are complex systems we are here dealing with an interdisciplinary problem in the sense adopted here. However, it is not a problem of immediate societal relevance. Rather, it is a basic research question within biology. Because they arise within science, this type of research questions can provide a kind of epistemological motivation for going into interdisciplinary problem solving that is different from the motivations for taking on the societal challenges, so commonly emphasized by funding agencies (cf. sec 2).

It is important not to overlook this type of interdisciplinary problems as they seem to be the driving force behind much successful interdisciplinary problem solving, at least in the sense that it leads to the formation of new hybrid disciplines. Evolutionary novelties are now the focus of a new research field called evolutionary developmental biology – evo-devo for short – and previously other basic interdisciplinary research problems have given rise to new disciplines like biochemistry and cell biology even though they were not of immediate societal relevance (Bechtel 2006; Collins, Evans & Gorman 2007).

There are, thus, good reasons for allowing for the possibility of basic interdisciplinary research, even if it is also true that many questions of immediate societal concern are interdisciplinary as characterized above (see also (Klein 2000)). The downside to this choice is that it may require two separate arguments to argue for the value of interdisciplinary research; one for interdisciplinary research addressing problems of immediate societal concern, and one for basic interdisciplinary research. As mentioned I will not go into the debate over whether we should promote interdisciplinary problem solving, but only consider how interdisciplinary problem solving can be taught.

3.2 Intermezzo on disciplines

In the previous section I argued that both questions 2. and 3. are interdisciplinary problems. However, the two problems highlight an ambiguity in my discussion of interdisciplinary problems that has been left unaddressed until now: What is a discipline?

The ambiguity arises when we consider why the explanation of the evolution of the vertebrate jaw cannot be evaluated entirely within one discipline. Is this because paleontology – sometimes considered a part of the geological sciences – is involved or is it also because multiple *sub*-disciplines of biology are involved? In other words can interdisciplinary problem solving occur entirely within we often call disciplines – physics, biology, chemistry etc.?

My answer to this question is clearly yes. The science disciplines – especially biology – are so diverse that, in many contexts, collaborations within them should be analyzed as crossdisciplinary[20] collaborations (see also chapter 2 section 3 and chapter 3).

Due to this I have throughout this dissertation chosen to use the term *scientific domain* or simply domain (Collins 2011) rather than discipline when discussing interdisciplinary science[21]. Depending on the purpose of the study, 'domain' may refer to any social unit of science, be it a discipline, sub-discipline or even more specific communities within sub-disciplines. The challenge is to make sure that the concept does not become so plastic that it applies to any group of researchers. Some groupings in science should not qualify as a scientific domain.

Philosophers have historically employed two ways of distinguishing social units of science: Social units of science have distinguished by object of study or by epistemological standards. The former has roots back to ancient Greece; the latter is commonly attributed to Kuhn. None of them are unproblematic. In particular both run into vicious circularities.

Social units of science can be distinguished by looking at the objects and properties that are being studied. Aristotle distinguished metaphysics – the study of being *qua* being – from, for instance, biology – the study of being *qua* living (Metaphysics, Kappa 7). This way of dividing science into domains is also seen in the names we use to describe many research domains: biology, geology, zoology, organic vs inorganic chemistry etc.

Dividing science into smaller units in this way runs into a number of problems. For instance, when we attempt to use it on a discipline like physics that does not seem to have a particular category of objects that is being studied. Rather, the "physical properties" of any object may be studied. But what are physical properties? The best answer seems to be the properties that physicists study. And so the circle begins…

More generally, the idea that reality can be divided into stable well-defined spheres, including the biosphere and geosphere, each covered by a particular science is challenged in at least two ways. First, by history as the study of specific properties of certain objects shift from one sphere to another (cf. note 14), and second by the fact that specific properties of certain objects can simultaneously be the object of study for many different disciplines albeit from different perspectives (Maull 1977). Taken together, these issues challenge the idea that nature is independently divided into spheres, and suggest that what constitutes a sphere is at least partly defined by human research interests. So biology is the study of the biosphere, and the biosphere is defined referring to what is studied by biologists. And so the circle begins again…

---

[20] Following the convention I use *cross*disciplinary problem solving to refer to any kind of problem solving that crosses disciplinary boarders. Crossdisciplinary problem solving can be subdivided into interdisciplinary problem solving which involves a significant degree of integration and multidisciplinary problem solving which does not (see section 1.3.3).

[21] In philosophical ears this may sound odd as philosophers more commonly refer to communities (following Kuhn), disciplines or fields (e.g. (Darden, Maull 1977)). In the philosophical literature the term 'domain' is often associated with Shapere's (1984) influential account of domains of *theories.* When drawing on multiple, relatively distinct literatures, as I do, it is hard to avoid this kind of confusion. I have attempted to compensate for this in this introduction by stating very clearly what I take central terms to mean.

Kuhn (1996) took a different approach to the division of science into social units. He equated disciplines with scientific communities and defined a scientific community as a group of researchers that have come to share a *paradigm* because they have "undergone similar educations and professional initiations [and] in the process they have absorbed the same technical literature and drawn many of the same lessons from it" (p. 177). Kuhn saw the "technical literature" as conveying much more than factual knowledge. Science textbooks in particular present carefully selected examples of *good solutions to relevant problems* – what Kuhn called *exemplars* - that display certain epistemological standards to the student. As the student becomes enculturated into a specific paradigm she internalizes the epistemological standards displayed by the exemplars chosen by her teachers, and thus absorbs not only science content knowledge but also epistemological standards from the exemplars.

Like the Aristotelian account, Kuhn's original account runs into circularity problems. Kuhn initially provided no independent definition of what a paradigm is. It was simply characterized as that which members of a scientific community share. This is unfortunate, as Kuhn himself acknowledged (p. 176), since a scientific community was defined as a group of scientists, who share a paradigm.

Kuhn made an effort to break the circularity in the postscript that came out with the second edition of *The Structure of Scientific Revolutions*. Here, he characterized the *disciplinary matrix* in an effort to make the term 'paradigm' less vague. Furthermore, he argued (pp. 177) that scientific communities can be identified independently of paradigm considerations by looking, for instance, at citation networks or attendants lists of conferences and memberships of professional societies. This way, the claim that scientific communities share a paradigm becomes an empirical claim rather than a matter of definition[22].

Kuhn's method for identifying scientific communities gives the same plasticity to the term 'community' as can be found in the term 'domain'. For instance, a community of chemists may be identified by asking recent Nobel laureates in chemistry which chemistry journals they read and publish in, which recurring conferences they commonly attend and which professional societies they are members of and then identify people with similar combinations of reading and publication habits and societal memberships. Such a process will always result in boundary cases, but more or less connected clusters will appear (cf. below). Within this larger community more fine grained divisions may be drawn. National chemical societies may have sub-divisions, for instance in organic and inorganic chemistry, and the chemistry journals may be divided in the same way. Even these communities may be subdivided depending, for instance, on who regularly cite each other and meet regularly for specialized conferences. When I use the term 'scientific community' I use it in this plastic sense.

Today there is a vast and varied literature mapping aspects of scientific communities using a various methods (for a comprehensive recent review see (Morris, Van der Veer Martens 2008) see also (Chubin 1985)). Bibliometric methods are used to identify citation clusters and common topics, whereas other more qualitative sociological methods can be used to study shared beliefs and values within these clusters. Using such methods

[22] In one respect this claim has been contested. One part of the disciplinary matrix, according to Kuhn, is "symbolic generalizations", which are statements that are "readily formalized" using universal quantifiers. Several scholars have noted that such law-like statements do not have the same elevated importance in the life sciences as they do in the physical sciences (Mitchell 2000).

it is possible to identify a number of more or less overlapping “research specialties”. Members of these specialties “tend to work on a related set of problems, adopt a common paradigm, publish in the same set of journals, use a common technical jargon, attend the same technical conferences and cite the same set of core references in their papers”[23] (Morris, Van der Veer Martens 2008, p. 239). Within these research specialties a set of core members can be identified along with periphery members who may be members of a number of specialties simultaneously. Counting the members of a research specialty is difficult (when, for instance, is a person a periphery member rather than not a member of a research specialty?), but order of magnitude estimates put the size of research specialties at a few hundred up to a few thousand people (pp. 236-38). Groups larger than this are not considered sufficiently homogeneous with respect to all of the variables measured, to qualify as social units of science by the standards commonly employed in these studies. However, it is recognized that there are clusters of clusters of researchers that are homogeneous with respect to the variables mentioned relative to other, similar sized clusters, even if they are not homogeneous compared to the individual clusters. These are what we can call fields or even disciplines.

Education does not begin at the level of research specialties. Rather, it commonly begins at the higher level of discipline or field. If (as discussed in chapter 4) internalization of epistemological standards begins early in the education before specialization really begins, then we can expect that within these more general fields there is significant homogeneity with respect to epistemological standards even if there is not homogeneity with respect to other variables - such as technical conferences attended.

In this dissertation I have therefore adopted a the more flexible term ‘scientific domain’ rather than ‘discipline’ or ‘research specialty’ to refer to *any cluster of researchers that, relative to its size, is homogeneous with respect to relevant variables* - problems studied, exemplars, epistemological standards, journals published in, technical jargon and conferences attended.

Occasionally I refer to scientific communities as well. As indicated above, I take ‘scientific community’ and ‘scientific domain’ to be largely interchangeable.

3.3 Interdisciplinary solutions

With the clarification of scientific domains in place we can get back on track towards a characterization of interdisciplinary problem solving that can serve as the basis for answering the question of what special skills a required to engage in this process. Given the considerations presented in the previous section we can characterize interdisciplinary problems as follows:

> An **interdisciplinary problem** concerns a complex system, contains relevant sub-questions about the complex system and the answers to these questions that cannot be evaluated solely within one domain.

[23] This quote should not be read as a confirmation of Kuhn’s claim that scientific communities share a disciplinary matrix, as the authors’ use of the term ‘paradigm’ in this context covers mainly exemplars and “validation standards”, what I have called epistemological standards.

Interdisciplinary problem solving aims to find solutions to interdisciplinary problems. That much is clear. However, interdisciplinary problems may also be addressed in monodisciplinary research although the result might not be very satisfying. Interdisciplinary problem solving thus cannot simply be characterized by the kinds of problems it investigates. We need to look at the kinds of solutions that are sought after as well. This implies looking into the process of *integration.*

In broader discourse interdisciplinary problem solving is often used to refer to all kinds of problem solving that gets input from multiple domains. However, in the more specialized literature on interdisciplinary problem solving *cross*disciplinary problem solving commonly serves the role of the umbrella term, and *inter*disciplinary problem solving in then used to refer to a specific kind of crossdisciplinary problem solving. It is common to contrast *inter*disciplinary problem solving with *multi*disciplinary problem solving, by saying that contrary to multidisciplinary problem solving the aim of interdisciplinary problem solving is *integrated* solutions to interdisciplinary problems (Klein 2010). A number of different domains may contribute to multidisciplinary research, but no effort is made to integrate these contributions into a coherent solution. Interdisciplinary research aims to take this further step, and integrate the individual contributions. Learning to perform integration is thus a central aim for the student who is aiming to learn *inter*disciplinary problem solving. What integration more specifically consist in is not always clear, but looking into what has been said in the literature about integration will help us move towards a clearer idea about how integration can be taught.

Philosophers and educators tend to talk about integration in at least two different ways: Integration of *domains* and integration of *cognitive resources*. Often the two are conflated.

When tracing how new research domains like cell biology, biochemistry or nanoscience emerge from established domains, scholars often talk about the integration of *domains* (Collins, Evans & Gorman 2007, 2006; Grantham 2004; Bechtel 1986; Darden, Maull 1977). This first kind of integration implies mutual adaptation of epistemological standards within the domains, mutual exchange of methods and knowledge between the domains and more. The result of this process may eventually be the formation of an entirely new domain complete with journals, conferences and educational programs. As argued by Brigandt (2010) the majority of attempts to solve interdisciplinary problems that are generally deemed valuable do not result in this kind of integration. It can also seem unnecessarily ambitious to aim for this kind of integration and say that *all* interdisciplinary problem solving should contribute to the formation of new research domains in order to be deemed successful. Arguably, we should say that if an interdisciplinary research collaboration contributes to the solution of an interdisciplinary problem in an appropriate way, then this should be regarded as a success regardless of whether or not a new research domain is formed. At least this will be more in keeping with the judgments of practitioners.

The other way of talking about integration involves the integration of *cognitive resources*. With respect to a specific interdisciplinary problem, multiple domains will have standards in place to evaluate solutions to the overall problem and some of the relevant sub-questions. They will also most likely have contributed to the answer to some of these questions. The contributions from different domains may not be immediately combinable. For instance, models based on conflicting assumptions may have been used to reach the individual insights. Integrating these insights involves removing these tensions, contradictions even, among the available insights relevant to the specific interdisciplinary problem and in the process gain new insights. For instance by developing a new way of modeling the complex system, that to some extent, incorporates both the models initially supplied by the participating domains. The aim is not to keep every element provided by the individual domains, but to shape the individual elements and fit them into a new whole that satisfies all parties involved. Carp (2001) argues that the integrated solution should ideally "belong to no one" (p. 85), perhaps it is more appropriate to say that it should belong to everyone.

Integration in this sense is not an all or nothing concept. Rather, it is a matter of degree. Contributions to the solution of a given interdisciplinary problem may come from, say, four different domains. Various attempts to integrate these contributions may be presented. A solution that incorporates all contributions from all four domains is (*ceteris paribus*) *more integrated* than a solution that only manages to incorporate contributions from three of the four domains. Similarly, a solution that is considered *good* by the standards of *all* domains involved is more integrated, than an acceptable solutions that is only considered good by the standards of a few of the domains involved.

Integration of cognitive resources is, thus, a creative process that seeks to reframe and shape existing insights to create a solution to a specific interdisciplinary problem that as a minimum does not violate the epistemological standards of the relevant domains. Noticeably, this view does not argue that integration involves the development of a new perspective that can be employed beyond the specific interdisciplinary problem. Researchers from different domains may come together to solve a very specific problem, say to help a given city tackle its traffic problems, resulting in a concrete chance in the cityscape that is considered an acceptable solution to all parties involved, without necessarily changing anything fundamental in the domains represented in the collaboration or creating an approach that can be applied without fundamental modification elsewhere.

The two kinds of integration described above are intimately linked. On the one hand, it is easy to see how solving concrete interdisciplinary problems may lead to fruitful new perspectives that researchers want to apply to other problems. Successful integration of cognitive resources may thus start a process that leads to integration of domains.

On other occasions integration of domains may be necessary for achieving full integration of cognitive resources. For a solution to an interdisciplinary problem to be integrated, it must be acceptable when judged by a number of different epistemological standards. However, this may not always be possible, say if the epistemological standards of two domains are outright contradictory[24]. In such cases integration of cognitive

[24]An example of contradictory explanatory standards would be if one domain consistently prefers explanations that reduce higher level phenomena to lower level phenomena, and another domain rejects the very possibility of such explanations of certain phenomena e.g. mental phenomena.

resources cannot occur unless the epistemological standards of the domains involved bend towards one another[25].

In relation to interdisciplinary programs in higher education, and the teaching of skills needed in interdisciplinary problem solving it is mainly integration of cognitive resources that scholars have in mind when they talk about integration[26]. Arguably this is also the most relevant notion of integration in relation to education.

It is quite possible that those who run some interdisciplinary programs hope that the students will eventually contribute to the firm establishment of the emerging domain they are being enculturated into. If so, it is integration of domains that is the final aim. In this case solving concrete interdisciplinary problems (i.e. integration of cognitive resources) is a means to achieve the final aim of establishing a domain of significant size to allow *normal science* in a Kuhnian sense to function within it. Other interdisciplinary programs have interdisciplinary problem solving and thus integration of cognitive resources as the final aim, perhaps because the program focuses on interdisciplinary problems with immediate societal relevance. Either way, students in interdisciplinary programs should at least learn to find integrated solutions to interdisciplinary problems, where integration is understood as integration of cognitive resources.

With these remarks we have moved away from the more elusive notions of interdisciplinary problem solving, and towards a more detailed characterization of the central teaching goal in interdisciplinary education programs. By doing so we have also taken the first small steps towards identifying the skills needed to reach this goal. Further steps in this direction are taken in chapter 3.

[25] Some scholars seem to believe that this means that interdisciplinarity will break down the disciplines, leading to *trans*- or *post*disciplinary science finally fulfilling the vision of unity in science so forcefully held by the logical positivists. There are many reasons to think that this will not happen, and that we should not want it to happen (Andersen 2013; Hansson 1999; Bechtel 1986; Kuhn 1977). One of the reasons being, as we have seen, that interdisciplinary science can give birth to new research domains and thus create *more* domains not fewer.
[26] For a review of different models of integration see (Repko 2008, pp. 123).

4 Interdisciplinary education

The aim of section 3 was to exhibit some central characteristics of interdisciplinary problem solving. With this characterization in place I briefly return to the limits of this project.

An increasing number of higher education degree programs are described as interdisciplinary. The arguments in this dissertation will be more or less relevant for many of these but not all. This is partly because I have a narrower understanding of interdisciplinarity than can be found in some parts of discourse, and partly because I take the natural sciences as a starting point[27].

For obvious reasons most interdisciplinary programs focus on training students to solve interdisciplinary problems that involve specific domains, rather than interdisciplinary problem solving in general. My interest in this dissertation is interdisciplinary programs that focus primarily on domains within the natural sciences. Examples include nanoscience programs, programs in biotechnology, systems biology, environmental science bioinformatics, and many other programs that have an explicit emphasis on a broader kind of interdisciplinarity that integrates insights from what is commonly referred to as different disciplines within the natural sciences. Often these programs will to some extent draw on the engineering sciences as well. I do not see a sharp distinction in kind between natural and engineering science but rather a gradual transition, I include this entire transition zone when I speak of natural science.

Much of my analysis will focus on the process of integration and the skills that are valuable when trying to integrate cognitive resources from different domains. Integration is central to interdisciplinary problem solving in the characterization adopted here, but as we have seen there are many understandings of what interdisciplinarity is. I am thus not claiming that my arguments are equally relevant for all science programs that draw on multiple domains, particularly because some do not in practice aim to teach students to integrate cognitive resources from different domains[28]. However, I claim that *if* the aim of a program is to teach interdisciplinary problem solving in a sense similar to the one characterized above, then the analysis and suggestions presented in the following chapters will be of relevance, especially if the aim is to teach interdisciplinary problem solving drawing on domains with rather different epistemological standards (cf. chapter 3, sec. 4).

It may further be argued that all science programs in higher education, at least in the EU, are expected to teach interdisciplinary problem solving, even if some emphasize it more than others (Grüne-Yanoff 2014). The shared European definition of the skills required to earn a master's degree reads that the candidate should have:

> **[...] specialised problem-solving skills required in research and/or innovation in order to develop new knowledge and procedures and *to integrate knowledge from different fields.* (European Commission 2008, p. 13, emphasis added)**

[27]Much of general discussion so far has not depended on my focus on the natural sciences (but see note 13). However it will influence the discussion in chapters 2 and 3.
[28] And as mentioned I have refrained from attempting to argue that they *should* do so.

So it seems that the vision to train students to perform interdisciplinary problem solving does not pertain only to specific higher education programs but to all, although some programs may emphasize broader kinds of interdisciplinarity – i.e. interdisciplinary problem solving drawing on domains with rather different epistemological standards[29].

As argued in the next section (and in a different way in chapter 3), committing to the aim of enabling students to perform interdisciplinary problem solving, especially relatively broad interdisciplinary problem solving, effects the desired content of the curriculum.

5 Extra-disciplinary expertise in higher interdisciplinary education

I started this introduction with a quote from Larson and collaborators arguing that a "considerable body of literature" shows that not just anyone can perform interdisciplinary problem solving since it requires some special skills. Based on some of this considerable body of literature I have now provided some clarifications of the key concepts used to describe interdisciplinary problem solving. Through this analysis we have gained some more details on what exactly it is that is required in interdisciplinary problem solving. In section 3.3 we saw that that finding an integrated solution to an interdisciplinary problem requires not only *transfer* of cognitive resources from one domain to another, but also adaptation and extension of these cognitive resources into an integrated solution that is acceptable relative to the epistemological standards of a number of domains. Given that even practicing researchers find this process very difficult (Campbell 2005; Lélé, Norgaard 2005; Jakobsen, Hels & McLaughlin 2004), it is unreasonable to expect students to simply learn this by themselves. It is therefore essential that students in interdisciplinary programs are prepared for engaging in the process of integration of cognitive resources from various domains. The question is how to best prepare them for this.

Many current education programs described as interdisciplinary consist simply of sequential or parallel tracks, where students are taught subject matter from different science domains by specialists from these domains (Klein 2010). Other programs where students are given a basic introduction to the sciences can have a very similar structure, at least at the bachelor level, even if they are not described as interdisciplinary. A challenge for any program with this structure is that the separate tracks that the students go through remain largely separate in the heads of the students. As a result, students have problems transferring even basic knowledge and skills gained in one track to advanced courses in other tracks (Christiansen, Rump 2008)[30].

Given that it is difficult for many students to transfer cognitive resources from one domain onto puzzles from a different domain without explicit training, it is no surprise that interdisciplinary programs, where students are expected to learn to not only apply cognitive resources in different contexts, but also to adapt these in order to find integrated solutions to interdisciplinary problems, will have problems as well. The above discussion of the

[29] For more on the distinction between broad and narrow interdisciplinary problem solving see (Klein 2010).

[30] Presented in this way the problem is a special case of the more general problem of knowledge transfer from one context to another. This problem has been discussed extensively in psychology and science education. Problems of knowledge transfer have commonly been studied in relation to younger students who are tested in *far* transfer, that is, transfer of knowledge between very different contexts – the physics classroom and the supermarket for instance (Mestre 2005; Barnett, Ceci 2002). We see here that *near* transfer can also be a challenge even for older students.

specific characteristics of interdisciplinary problem solving gives further clues to what the special challenges facing students in interdisciplinary programs are. The need for integration arises when attempting to solve a problem addressing a complex system about which relevant sub-questions have been asked by various domains. The solution to such a problem is constrained by the epistemological standards of multiple domains, and practitioners must be able to navigate all of these different epistemological standards when searching for a good solution. Such skills are not taught in any standard science course. Rather, they are gained tough reflection on how science works, and how different sciences work differently. Such reflections are *meta*-reflections in the sense that they are not about nature, but rather reflection on the ways nature is reflected upon in different domains. Such meta-reflection is often thought to be the turf of philosophy of science and HPSS in general. Thus, Eigenbrode and collaborators (2007) argue that the special challenges in interdisciplinary problem solving imply addressing philosophical rather than scientific questions. If this is so, the question immediately arises whether there is something (results, concepts, methods?) in the 2500 year long history of philosophy that it would be relevant for the students or their teachers to know about? As illustrated in section 7 and 8, the answer is yes, but not as much as one might expect.

6 Teaching extra-disciplinary skills: The role of HPSS (mostly PS)

While insisting that interdisciplinary problem solving can be taught Clark and Wallace argue, based on their teaching experience, that:

> ***Many students find it difficult to confront their own epistemology, cognitive status, disciplinary prejudices, and conventional notions about policy processes […]. It is, in fact, possibly the most difficult part of learning interdisciplinarity (Clark, Wallace 2010, p. 178)***

Teaching skills in meta-reflection – or *meta-reflective skills* - is thus no small challenge. There seems to be a number of reasons for this. One tricky question is *how* to teach. Deciding how to teach ideally requires the explication of teaching goals, that is, an explication of *what* to teach them. This is where a framework like the one presented in chapter 3 become relevant. The expertise based account presented in chapter 3 shows that what we want to teach is the ability to do certain things which in turn requires a certain type of knowledge. The students should, in concrete cases, be able to explicate whether a proposed solution to an interdisciplinary problem is acceptable relative to the epistemological standards of the domains that she is most familiar with, be able to engage constructively in discussions about the quality of research with practitioners representing domains other than the ones she is most familiar with, and engage in the modification of the proposed solution in accordance with the conclusions drawn. The latter of these demands is of course the most challenging, and there is still much to be done in exploring how to train this process. Meeting these demands requires as a minimum some understanding of what domains and epistemological standards in general are, and relevant tools and knowledge to be able to continuously reflect on how solutions to concrete problems are constrained by the standards of the various domains.

The basic concepts for describing domains and epistemological standards come from HPSS, and it is one of the aims of HPSS to refine the existing vocabulary and use it to analyze how science develops. It thus seems unreasonable not, in one way or another, to draw on the insights gained in HPSS over the past decades when teaching meta-reflective skills. Some have advocated the direct study of HPSS as part of the curriculum, whereas others favor a more indirect approach, where HPSS informs the development of inquiry exercises and case material (for examples of the former see (Repko 2008; Spelt et al. 2009) for examples of the later see e.g. (Allchin, Andersen & Nielsen 2014) see also chapter 3 sec 5). Either way has its benefits, and the two approaches can of course be combined in the classroom. I am not going to enter a detailed discussion about how meta-reflective skills are best taught. Rather I will note that, regardless of whether HPSS insights are taught directly - through reading and discussion of say, philosophical theory - or whether they are taught more indirectly, the teaching of relevant HPSS insights and concepts requires that there are relevant insights and concepts from HPSS to be taught. That is, the need to teach meta-reflective skills creates a need for specific kinds of HPSS studies, for instance on the nature of scientific domains, the meaning, possibility and value of integration and detailed studies of the characteristics of the epistemological standards of various research domains. Some of these studies already exist in the literature, many are incomplete and some are completely absent.

As we have seen in the previous sections it is the *differences* in epistemological standards across scientific domains that create the need for integration. It is these differences that need to be negotiated in the process towards an integrated solution to an interdisciplinary problem, and it is these differences that students in interdisciplinary programs should learn to reflect upon and discuss constructively in relation to concrete cases. Philosophical studies - or other types of HPSS studies - of how epistemological standards differ across science domains would thus be of significant value to the teaching of meta-reflective skills. At least if they are sufficiently specific and context sensitive. Do such studies exist? I argue (in section 8) that they do not, at least when it comes to differences in explanatory standards. As mentioned in section 1.1, explanatory standards have been emphasized as an important subset of the epistemological standards of a domain in relation to interdisciplinary problem solving ((Green, Fagan & Jaeger 2014; O'Rourke, Crowley 2013; Love 2008; Hepburn, Thorén In preparation) see also appendix A).

Since explanatory standards form an essential subset of the overall epistemological standards of any domain our lack of knowledge about how they differ across domains makes our knowledge of differences in epistemological standards across science domains incomplete. Hence the second aim of this dissertation is to provide a general methodology that can be employed to begin filling this gap (chapter 4), and to provide a concrete example of how this general methodology can be put to use (chapter 5).

If the preceding pages have succeeded in clarifying and connecting the two aims of this dissertation, they have done so only by opening even bigger questions. One question arises out of the preceding paragraph: How does one conduct "sufficiently specific and context sensitive" philosophical studies? I will begin to grapple with this question in the next sections, and again in chapter 4 where I discuss methodological issues related to the study of explanatory standards.

As we follow the development of modern philosophy of science as it diversifies in its movement away from its positivistic origins, we also find a domain in constant fundamental debate over its own aims and methods. Focusing, as I do, on questions that arise in science education is not neutral with respect to this debate. It expresses the view that philosophers can legitimately consider questions that were once considered less relevant, and adopt methods that have commonly been looked down upon by traditional philosophers.

The process towards finding an integrated solution to the interdisciplinary problem "how are meta-reflective skills best taught" will thus require integration of domains – i.e. that the domains involved (science education, the HPSS disciplines and more) allow their epistemological standards to be shaped by other domains – in order to achieve integration of cognitive resources.

The influence goes both ways. Educators must also bend towards philosophers. In many ways this has already begun as the value of HPSS in general and history and philosophy of science in particular is increasingly being emphasized in education. The importance of some philosophical rigor and the usefulness of philosophical concepts is increasingly emphasized mainly in relation to the teaching of the nature of science (NoS) (see e.g. (Duschl 2008; Matthews 1994) for reviews)[31]. The continued exchange between the two fields will also highlight other relevant connections to be made to other neighboring domains. For instance, the connection between philosophy and the cognitive sciences which is already strongly developing gets a new dimension when related to teaching and learning.

7 Philosophy of science in practice: Relevant philosophy of science!

In the previous sections I argued that questions arising in interdisciplinary science education create a need for context sensitive philosophical studies. In this section I discuss a trend within parts of philosophy of science towards providing more practice oriented and context sensitive studies. This trend is sometimes called the *practice turn* in philosophy of science. My aim is not to show that the practice turn has (already) led to the formation of a new domain within philosophy of science (cf. the definition in section 3.2), as this would require an empirical investigation[32]. Rather, in order to give content to the idea of a practice oriented philosophy of science, I will to take a selective look at a trend within the debates over the aims and methods of philosophy of science that argues for an increased attention to scientific practice.

The practice turn in philosophy of science is a directed continuation of the historical turn in Anglophone philosophy of science in the 60's and 70's initiated by the works of Kuhn, Lakatos and Laudan and a variant within the general movement to naturalize philosophy of science (Giere 1985). One result of the practice turn is a partial reorientation of philosophy of science away from a fairly narrow focus on what Popper originally

[31] There are diverse understandings of what falls under the label NoS. Here I use it very broadly to include both the topics on the consensus list of topics (McComas 1998) and other topics such as the nature of scientific explanations. For a discussion of different conceptions of NoS see (Allchin 2013, ch. 1).

[32] One indication that such a subdomain is indeed taking shape is the recent formation of a society for philosophy of science in practice - the SPSP.

called the logic of research[33] and towards a much broader focus on understanding the practice of scientists. The practice turn not only broadens the focus of philosophy of science, it also takes a different angle on the traditional topics in philosophy of science (see section 7.2). A major benefit of the reorientation of the philosophical analysis of science is that it becomes increasingly relevant to others. The challenge for practice oriented philosophers is that once they broaden the range of questions that they are interested in, they also have to diversify the toolbox they use to investigate these questions.

I do not intend to give anything like a full historical account of how the transitions from positivistic philosophy of science to philosophy of science in practice transpired. Instead, I will present a more conceptual discussion illustrating different positions on the question of what the aim of philosophical analysis is and how this affects the required methodological toolbox of the philosopher. The discussion falls in two parts. In section 7.1 I discuss different views on what the aims of a philosophical analysis could and should be. The discussion is framed around responses to Feynman's controversial claim that "philosophy of science is about as useful to scientists as ornithology is to birds". In section 7.2 I discuss the methodological consequences of a broadening of the aim of philosophical analysis.

7.1 The aim of philosophical analysis

Physicist and Nobel laureate Richard Feynman is very famous for very many things. Among philosophers of science he is famous for claiming that "philosophy of science is about as useful to scientists as ornithology is to birds"[34]. There are many ways in which a philosopher may react to such a provocative claim, partly because it is not clear whether "useful" refers to all aspects of scientists work or specifically to the production of new scientific results. The philosopher's reaction to this claim will be telling for his view on what the aims of philosophy of science are, or should be. I will outline some characteristic responses in order to illustrate how the aim of practice oriented philosophy of science is different from other types of philosophy of science.

Maybe true, but not a problem. Some philosophers will accept that Feynman's claim may be absolutely true, but not be particularly worried by it, as they do not see it as the aim of philosophy of science to be relevant to scientists. There are plenty of people, including philosophers of science, who find philosophy of science useful and why should the judgment of scientists be privileged? The purest form of this reaction comes from philosophers who maintain that philosophical analysis, like basic research in theoretical physics, is valuable in itself even if it is not useful to anyone. Of course, if the results of philosophical analysis can be *applied* elsewhere then that's fine, but if scientists do not find philosophy of science useful, well, then that is at best an argument for not teaching philosophy of science to science students. Heather Douglas (2010) argues that this view was common in philosophy of science until recently and had been so since the 1960's. She writes that the prevalent view among philosophers of science was that philosophy of science was "its own field with its own internal standards, issues and agenda; whether it was of interest to scientists or society was not important" (p. 321). This, according to Douglas, made philosophy of science a rather narrow, but also a rather coherent, research domain. Research was structured around a few central problems - such as realism vs. anti-realism,

[33] The original German title of Popper's *The Logic of Scientific Discovery* was *Logik der Forschung*, which directly translated means "the logic of research".
[34] Although it is doubtful whether he actually ever said this.

rational theory choice, the nature of scientific explanation and how to account for progress in science - and the methods used to investigate problems were largely shared.

My response to such a view would be that I do not disagree that philosophical analysis can be valuable in itself. Nor do I wish to argue that every philosopher should aim to provide analyses that are useful beyond narrow philosophical circles. However, my aim and the aim of many other philosophers is to provide philosophical analysis that is in fact relevant beyond narrow philosophical circles and I see no reason why a philosopher should not have such ambitions.

Against my view, it might be argued that if the effort to make philosophical analysis more relevant takes the investigator away from the core problems and methods of traditional philosophy of science, then it also takes the investigator away from doing philosophy. It may be argued that while a trained philosopher may fruitfully take on problems that arise, for instance, within interdisciplinary science education, doing so will not be a *philosophical* effort[35].

A partial answer to this objection would be to question its importance. If a practice oriented philosopher contributes to the solution of an important problem, then what does not matter if her work is rightly called philosophy or something else? At the most, the label "philosophical" is important for where a contribution can be published, and whether the work qualifies the practice oriented philosopher for a degree or even a job in philosophy. The objection that practice oriented philosophy is not really philosophy thus works against the practice turn in philosophy of science at an institutional level. Therefore, practice oriented philosophers should be able to counter this objection. The objection seems to rest on the idea that there are certain problems that are inherently philosophical, and doing philosophy is to investigate these problems using philosophical methods. It thus rests on a version of the static 'sphere-view' of domains described in section 3.2. As we saw in section 3.2 this view is flawed. Problems drift in and out of research domains as the interests and cognitive resources of practitioners change. This is also the case in philosophy of science (Douglas 2010). Of course, at a given point in time some general problems are more suited for the trained philosopher than others, but this is not because of the philosophical nature of the problems. Rather, the reason is found in the (lack of) match between the skills needed to address (relevant) sub-questions of the problem (cf. section 3.1) and the skills of a person trained in the philosophical tradition. Like all other experts, philosophers should accept that their expertise is limited, and that they cannot solve every problem in the world, but they should also be able to recognize areas where they have something to contribute, even if it is outside the recent philosophical tradition.

Partly true, problematic for scientists. Other philosophers will respond to Feynman's provocative claim by claiming that it is to some extent true, and that this is indeed a problem since it is at least a partial aim of philosophy of science to provide results that are useful to scientists. However, they might argue, this has no effect on the way philosophy of science is practiced because large parts of the existing theoretical corpus of philosophy of science *is* in fact already potentially useful to scientists. Rather, the reason why many scientists

[35] Although this objection has not been raised against my work, versions of this objection that practice oriented philosophy of science is not real philosophy has been raised against some of my colleagues.

make little use of philosophy of science is the same reason that birds make little use of ornithology: They know very little about it! Continuing the animal analogies proponents of this view may borrow an expression from Lakatos and argue that the reason why scientists make little use of philosophy of science is that they "tend to know little more *about* science than fish about hydrodynamics" (1978, p. 62)[36]. On this view the problem for philosophers is not that they are asking irrelevant questions or providing irrelevant answers. Rather, they have simply not been good enough at educating scientists. The solution to the problem would thus be to start a campaign to teach students in high schools and universities general philosophy of science. An appropriate motto for the campaign would be: "Teaching ornithology to fish!" The aim of this campaign would be to teach science students how to apply general philosophical theories to the philosophical problems they face. If this does not work, an alternative would be to establish philosophical helpdesks on every campus, where scientists can bring their philosophical problems to the experts.

The view outlined above is perhaps not commonly defended explicitly within philosophy of science, but seems to underlie the view of proponents of applied general philosophy, be it applied epistemology or applied ethics - i.e. the view that philosophers can contribute to scientific debates by taking general epistemological or ethical theories "of the shelf" (Douglas 2010) and apply them to to specific problems.

I agree that learning some HPSS can be very valuable to science students, but that does not imply that all kinds of philosophy of science is relevant for science students. I am especially doubtful whether it is a good strategy to try and teach science students how to apply very general epistemological and ethical theories. The reason for this is that the philosophical problems that scientists face in practice are rarely *just* philosophical. Rather, they are relevant sub-questions to a more general problem. This means that the aim is not just to find a good answer to the sub-question relative to the epistemological standards of philosophy, but also to find an answer that can be integrated into a solution to the general problem. General epistemological theories like Bayesian theories can tell us how to act if our primary aim is to act rationally. Ethical theories can tell us how to act if our primary aim is to act in a morally defendable way. In practice, questions about epistemology and ethics are often entangled with questions about limited resources, policy and personal relations. Thus, the aims of practicing scientists are often more complex than the aims assumed in philosophical theories. This has significant consequences.

Philosophical theories, like all other general theories, can only provide us with an answer to a problem if we can model the problem in a language that the theory can handle (Cartwright 2003). As with all modelling this requires abstraction and idealization. The appropriateness of such abstractions and idealizations depend on the purpose of the modelling. If the aim of a model is to provide the basis for advice to scientists in their messy world of entangled ethical, epistemological and socioeconomic concerns, then it will not be appropriate to construct a model that focuses solely on one of these while ignoring the rest as this approach will fail to address the actual problem. However, we may have to construct a very simple model in order to handle the problem with our general philosophical theories. In such cases, we may be able to get an answer to a relevant

[36] Although Lakatos would probably not be supportive of this argument.

sub-question using a general philosophical theory, but the answer will be of little relevance, as it will be very hard to integrate into a solution to the overall problem.

If this argument holds then the “teaching ornithology to fish” campaign mentioned above should not simply aim to teach general philosophical theories, such as utilitarianism and deontology to science students as this will not help them very much in their daily practice. If philosophical theory should be taught directly, it should be philosophical theory that is (or can be made) sensitive to the aims and complexity of the specific practice that the students are becoming a part of (see chapter 3, see also (Grüne-Yanoff 2014; Eichinger, Abell & Dagher 1997)). If such philosophical theory does not exist then it is not the scientists’ fault that they do not know it, it is the fault of the philosophers who have failed to provide it.

Partly true, partly a problem for philosophers as well. The final reaction to Feynman’s claim that I wish to characterize here is the natural continuation from the discussion above. Many philosophers, including many in the more practice oriented part of philosophy of science, will argue that Feynman is right that learning philosophy of science is not useful to the research scientist in the sense that it will enable her to get more papers published in *Science* or *Nature*. However, this does not mean that practicing scientists and others working with the sciences cannot benefit from learning some relevant HPSS, as this will help them fulfill some of the other roles of the science professional and help them do research in a different way that may not be appreciated by the editors of *Science* and *Nature*, but is valuable nonetheless. Thus, if we interpret Feynman’s claim as referring very narrowly to “scientists” when working in the laboratory, and interpret “useful” as “raising the probability of getting scientific credit in the short run”, then there is probably some truth to Feynman’s claim, but if we interpret the claim as referring to actual practitioners with a science background doing what actual practitioners do, including, but not limited to, performing research, and interpret “useful” as something beyond raising the probability of getting scientific credit, then Feynman’s claim is arguably false.

Although some scientists may not be able see further than their next publication, there are (or at least should be) other things that matter when working in science. Science and technology is playing an increasingly important role in society in general. Scientists contribute (more or less directly) to the development of new technologies used in our everyday lives and appear regularly in the role of experts both in official contexts such as advisory boards and legal trials, and more informally in newspapers and other media. In spite of the objections from more traditionally minded philosophers, extensive HPSS research (too broad to be reviewed here) is being conducted in order to better understand, the interplay between science, technology and society. This research serves both to engage politicians and the public in more nuanced and reflective discussions of the role of science, technology and experts in society, but also to make scientists more aware of their roles, and help fulfill them in a more reflective manner. Furthermore, as we saw in section 2, there is a strong vision among politicians, funding agencies and some researchers that scientific practice should change in order to become more focused on interdisciplinary problem solving and societal challenges, although the current reward system in science does not reflect this vision. If the arguments presented in this dissertation go through, they show that learning HPSS can be a central part of gaining the skills necessary for performing interdisciplinary problem solving. HPSS can thus be relevant to all science students who go on to become researchers or in other ways engage in interdisciplinary problem solving.

Many science students do not engage directly in research and innovation after they finish their education, but go on to, for instance, become science teachers in other parts of the education system. As these students go on to teach science literacy to others they will also benefit from having reflected on the nature of science based on relevant HPSS material (see e.g. (Abd-El-Khalick 2013))[37].

It thus seems that if we look more broadly at what those educated in the sciences do, then there is an empirical argument to be made against Feynman's claim[38], although I have not provided the necessary data here. Philosophy is useful to some of those educated in the sciences and even to practicing researchers. Philosophers are fed (or take on) problems from other domains and provide valuable input. They identify problems in scientific practice or in the relation between science and society and feed them on to be investigated elsewhere if necessary. They engage the public, scientists and science students in discussions about scientific practice and the role of science in society and their research helps sophisticate these discussions. The really interesting question, therefore, is not whether philosophy and HPSS in general is useful or not. The interesting question is how the scope and methods of philosophy of science can be adapted to make philosophy of science even more useful and relevant.

7.2 Philosophy of science in practice: Normative force through descriptive accuracy

In the previous section I claimed that philosophy of science is regularly put into practice: what we may describe as philosophy-of-science in practice. I further claimed that not all philosophical theory is equally suited to be put into practice, partly because it is not sufficiently context sensitive. The question now is what kinds of philosophy are particularly suited to be put into practice? Promoters of the practice turn argue that a partial answer to this question is practice oriented philosophy of science: Philosophy of science-in-practice[39]. In this section I briefly discuss some of the meta-philosophical ideas behind practice oriented philosophy of science[40].

Philosophy of science in practice can be characterized through a combination of three aspects: 1) engaging in constructive problem feeding both as sender and receiver or 2) taking a new approach to traditional philosophical problems, combined with 3) a tight reflexive coupling between its methods and problems. Some of these characteristics are shared with other movements in philosophy of science, but the combination distinguishes practice oriented philosophy of science from other kinds of philosophy of science.

Consider first the willingness to engage in constructive problem feeding. Practice oriented philosophy of science is oriented towards the problems of scientific practice. What does this mean? It means that the practice oriented philosopher often (but not necessarily) takes as a starting point ongoing debates about science taking place in the individual sciences, in education, or in society in general, and submits them to

[37] This topic is further elaborated in chapter 5 of the upcoming book *Philosophy of Contemporary Science in Practice* co-authored by all practice oriented philosophers at the Centre for science studies in Aarhus (including myself).
[38] Provided of course, that we can persuade someone to spend time estimating how useful ornithology actually is to birds.
[39] The distinction between philosophy-of-science in practice and philosophy of science-in-practice was first introduced by John Dupré at the SPSP conference in Exeter in 2011. Subsequently it has been continuously discussed in the SPSP newsletter. For our group's contribution to this discussion see appendix B.
[40] The discussion is based partly on my contributions to the introduction chapter from the upcoming book *Philosophy of Contemporary Science in Practice* mentioned in note 37.

philosophical analysis. The aim of this analysis is to eventually be able to contribute to the ongoing debate. This dissertation exemplifies this approach[41]. The aim is not necessarily to resolve the debate but simply to contribute to the debate as one voice among many. How is this different from applied philosophy? To the extent that it is different from applied philosophy, it is so because the practice oriented philosopher takes a detailed contextualized description of the problem as his starting point. The practice oriented philosopher is not necessarily interested in learning which answer a given general philosophical theory would give to the problem unless the theory can be applied to a detailed and contextualized model.

Often, the practice oriented philosopher thus starts from a contextualized, accurate description of a problem discussed in the sciences. Although not unwilling, he thinks carefully before claiming that the concrete problem is an instance of a more general problem, and is attentive to the effects that such an abstraction has for the validity of the final result. In a way this is common sense, but it has significant consequences. To describe an ongoing scientific debate accurately requires that the philosopher understands the debate in significant detail. For instance, to be able to contribute to the debate in the biological sciences over the nature of genes (Rheinberger, Müller-Wille 2010), it is necessary to know quite a bit about biological theory, and the practice in which the concept is used. It may also be necessary to know the historical roots of this practice. To get a sufficiently detailed description of the problem the practice oriented philosopher often turns to detailed case studies of scientific practice. Because the practice oriented philosopher often needs deep knowledge of scientific practice she often needs to be rather specialized and often have a double background in philosophy and one of the sciences. It is thus not uncommon to hear practice oriented philosophers describe themselves as philosophers of some specific science rather than a philosopher of science in general. As philosophy of science turns increasingly to practice it therefore also risks becoming increasingly disunified and specialized.

Practice oriented philosophy of science is sometimes criticized for being "merely" descriptive[42]. The reason why is easy to see: Simply presenting a sufficiently nuanced picture of, for instance, an ongoing conceptual debate in science can fill a whole research paper. (Just consider the number of words in my coarse discussion of interdisciplinary problem solving in section 2-4.) Different positions must be identified and articulated, their underlying rationale presented, and their implications outlined. Often, it is not easy for the scientists themselves to do this clarificatory work, and performing it represents a valuable contribution in itself, even if it is not explicitly normative[43]. The philosopher may well realize that the best strategy is not to attempt to solve a newly clarified problem herself, but rather to feed the transformed problem back to the science in which it originated. Does this mean that clarificatory work is not philosophical? If one affirms this question, one must also affirm the claim that the theoretical physicist modelling catalysis on a platinum surface is not doing physics if he leaves the experimental testing of his work to a chemist. I think most would be reluctant to draw this

[41] Other examples include philosophical input to ongoing debates about misconduct in research (Andersen 2014) or the value of design heuristics in systems biology (Green 2014)

[42] Again, this objection has not been raised against my work, but has been raised at conferences against the work of my close colleagues.

[43] As argued in (Wagenknecht 2014, ch. 3) the idea of a purely descriptive study is an illusion. Any study involves making choices with respect to the quality and relevance of information. It may be that we chose simply to describe what we read and observe, but deciding what to read and observe still has a normative aspect.

conclusion. Rather, they would say that because the theoretical physicist is using some of the cognitive resources that are being transferred from one generation of researchers to another with in the domain of physics he is actually doing physics. If so, they would also have to agree that a person using the cognitive resources of philosophy to clarify an ongoing scientific debate without taking an explicit stance in the debate is actually performing philosophical work.

Practice oriented philosophy of science does not necessarily deal with problems identified outside philosophy of science. Classic philosophy of science problems like explanation, justification and causation may also be addressed in a practice oriented way. Addressing the classic problems in a practice oriented way implies a transformation of these problems. The classic philosophy of science problems focus primarily on identifying the relations between the theoretical products of science and a possible mind independent reality. Addressing these problems in a practice oriented way means focusing also on the processes through which scientific results are constructed and the aims and expertise of the people taking part in this process (Ankeny et al. 2011). Chang (2011) outlines the philosophical grammar for how to transform the classic philosophy of science problems. The transformation begins by formulating the classic problems increasingly in terms of verbs rather than nouns. The practice oriented philosopher of science deals with explaining, justifying and causing rather than explanation justification and causation. What does this difference imply? For one, it means that we must start to think about who, or what, is doing the explaining, justifying and causing. Science becomes a (human) activity again rather than a matter of theories T and T*. Explanation (discussed in detail in section 8) goes from being a question about the logical relations between a phenomenon P and a set of sentences E that are claimed to explain P, to a matter of one person, complete with interests, expectations and imperfect expertise, trying to explain something to someone else, who in turn has her own interests, expectations and imperfect expertise. Thus, turning the familiar nouns of philosophy back into verbs not only introduces the agent who explains and justifies, but also the other, whom the explanation or justification is given to. Of course, this complicates things quite a bit, but it also means that philosophical theories become less abstract and more related to the concrete instances of explaining, justifying and causing that are controversial in daily practice. I illustrate this for the case of explanation in section 8.

Finally, practice oriented philosophy of science is special with respect to the way it addresses the problems adapted from the pool of standard philosophical problems or identified from outside philosophy of science. As illustrated above, part of the problem is often to find out what the problem actually is. The practice oriented philosopher will often resort to detailed case studies in order to get a grip on an actual problem in scientific practice. Case studies are also valuable when studying how the classic philosophical problems change when real agents in context are reintroduced. On other occasions the practice oriented philosopher may combine traditional philosophical analysis, such as concept analysis, with historical and sociological methods, interviews and custom-made empirical methods (cf. chapter 4). Not only is it useful for the practice oriented philosopher to resort to these methods, he may in many cases be forced to do so[44]. Brunnander (2011) argues that if a philosopher's aim is to be descriptively accurate, for instance when describing how scientists explain, then

[44] In the "unforced" way that reason can force a researcher to adopt a certain method (cf. (Habermas 1996))

simply relying on personal intuitions or thought up scenarios will not do. If a philosopher makes a claim and purports it to be empirically accurate, then he had better provide empirical evidence for that claim[45].

Since practice oriented philosophy of science often involves a more empirical approach it is compatible with Brunnander's conclusion about providing empirical justification for empirical claims, and can be seen as part of the ongoing movement to naturalize philosophy in general and philosophy of science in particular (e.g. (Quine 1969; Giere 1985)). Philosophy of science is a way of studying science on a par with other ways of studying science, and its claims must be judged by similar standards as those that apply to other theories of science. If the methods of other domains are suitable for tackling the problems adopted by a philosopher then the philosopher should not hesitate to add these methods to his toolbox (or start a collaboration with an expert from another domain). Giere writes:

> ***If philosophy of science is naturalized, philosophers of science are on the same footing with historians, psychologists, sociologists, and others for whom the study of science is itself a scientific enterprise. The most philosophers of science could claim is to be the theoreticians of a developing science of science on the model of theoretical, as opposed to experimental, physics. Would that not be status enough? (Giere 1985, p. 343)***

To this, the practice oriented philosopher adds that just as theoretical and experimental work is not always delegated to different people in physics, so too can the philosopher on occasion take on the role as the empirical investigator.

Is practice oriented philosophy of science then simply the same as naturalized philosophy of science? Not quite. Not all naturalized philosophy of science is practice oriented. Naturalized philosophy of science as described by Giere deals mainly with the classic philosophy of science problems, albeit in a new way. Practice oriented philosophy of science on the other hand takes on new problems in addition to the classic philosophy of science problems, and sees the classic philosophy of science problems in a new light. Giere thus promotes a different way of addressing familiar problems, whereas practice oriented philosophy of science transforms the familiar problems and in addition promotes attention to new kinds of problems. It is the focus on these different problems that more closely reflect the problems faced outside philosophical circles that makes practice oriented philosophy of science potentially more directly relevant to others. But, focusing on relevant problems is not enough: if philosophical analysis of relevant problems is to carry any weight inside or outside philosophy, it must also be methodologically sound. The first step towards achieving this aim is to take the arguments of Giere and Brunnander seriously and move away from the heavy reliance on personal intuitions as justifications of empirical claims and towards the realization that some of the problems that practice oriented

[45] Brunnander further argues that philosophy of science need not strive for descriptive accuracy, because the aim of philosophy is to change practice not describe it. This argument is vulnerable to an objection similar to the one raised against applied philosophy in section 1.7.1. If philosophy of science is to carry normative force it needs to be sensitive to the actual aims and practice of researchers in order to be able to say what is wrong with the current situation, and how an alternative approach will help this situation. (Just as climatologists need to base their advice on how to change the way we affect the climate on a detailed description of how we currently affect the climate.)

philosophy aims to address are not just more complicated than the problems of traditional philosophy of science; they are complex interdisciplinary problems (cf. section 3.1), that require integration of cognitive resources from (at least) the philosophy, history, and sociology of science just to be properly spelled out (let alone solved).

The research presented in this dissertation is practice oriented in a number of different respects. It takes on a general problem identified in science education not in classic philosophy of science: How can we prepare students for interdisciplinary problem solving. The analysis of this problem provided so far and continued in chapter 3, focuses on the process of solving an interdisciplinary problem, a process involving real people with aims, expectations and imperfect expertise. The epistemological challenges related to interdisciplinary problem solving have been identified through practitioners' experiences, although I have not collected these experiences myself. As discussed in section 5, part of the epistemological challenges in interdisciplinary problem solving arise because different domains have different explanatory standards. Note that this problem relates to explaining – the action in which a group of scientists is involved – rather than explanation. The analysis of differences in explanatory standards presented in chapter 5 is based on an empirical analysis of textbook explanations, based on a method that explicitly aims to avoid reliance on the intuitions of the philosopher (cf. chapter 4). Furthermore, as I will outline in the next section, the empirical analysis of textbook explanations is based on a view of explanations that brings the agent and the other – the receiver of the explanation - back into the picture where logical relations between sentences have traditionally been in focus. Explanations and the act of explaining are thus intimately linked, providing an account of the one implies providing an account of the other.

8 Scientific explanations: A view from philosophy of science

One of the central concepts in this dissertation is *explanation*. In this section I present the inclusive pragmatic-functional account of explanations that underlies the later chapters. The details of my account differ from other available accounts of explanations, and may thus be seen as a contribution to the philosophical debate about the nature of explanations in general as well as an outline of the framework for the later chapters. In this section, I also discuss the main trends in the philosophical literature on explanations. This exercise serves both as part of the justification for my own account (as I argue that my view captures in a way that alternative accounts do not the few consensus points in the literature) and as a justification for the claim made previously, that philosophers have largely neglected to investigate the differences in explanatory standards that exist across scientific domains.

I draw mainly on the philosophical literature on explanations but sources from other domains are also drawn upon. Some of the non-philosophical studies of explanations are discussed further in chapter 8.4. Since I will defend a very inclusive view of explanations, I do not intend to criticize the details of individual accounts, but rather show how they can be combined in a more general framework. I therefore refrain from detailed discussions of individual accounts and focus primarily on the questions that have been asked.

The inclusive pragmatic-functional view on explanations that I defend characterizes explanations in the following way:

> An **explanation** is a non-irrelevant response to a question taken by the respondent to be explanation-seeking.

This is not meant to be a definition of the term ‘explanation’ as it does not state necessary conditions for something to be an explanation. Nor is it intended to be an answer to the question of what distinguishes good explanations from bad. Rather, I argue that it is a characterization that covers very many instances of explanations and is suited to serve as a basis for identifying explanations in practice.

The view defended here is *pragmatic* in the sense that it maintains that the quality of an explanation cannot be determined without considering the intended audience of the explanation and the context in which the explanation is given. The aim of the act of explaining is to provide understanding to the intended audience of the explanation. Whether or not this aim has been achieved cannot be judged without considering who the intended audience is. As will become clear this pragmatic dimension is hidden in the clause that the response must be “non-irrelevant”. I further characterized the view defended here as *functional* to underline that it maintains that it is impossible to adequately characterize explanations and explanation-seeking questions based on logical or linguistic characteristics. As elaborated in the following sections, explanation-seeking questions must be characterized with reference to their intended function. The quality of the answers to these questions must also be judged relative to this intended function. The view defended here is also very *inclusive*. My characterization is based on a rather permissive account of explanation-seeking questions compared to most other accounts. It therefore incorporates as special cases a number of alternative accounts of explanations rather than excluding them. Although the view defended here is rather permissive as to what should be counted as an explanation, it is not equally permissive as to what should be counted as a *good* explanation. Although the characterization of explanations provided above implies that there are very many very different answers to very many very different explanation-seeking questions that should be counted as explanations, it does not imply that there are equally many good explanations; it only implies that there are very many explanations that are very bad for very many different reasons.

In the remainder of this section 8 I review the existing philosophical literature on explanations in science and defend and elaborate on the view sketched above. The section has four parts. In 8.1 I discuss the very general philosophical studies of scientific explanations aiming to identify the differences between scientific explanations and other kinds of explanations and identify the few consensus points that have been reached. In 8.2 I review the recent philosophical literature on understanding - the product of explanations - and identify consensus points. Expanding on these results and the results from the previous section leads me to the pragmatic-functional characterization of explanations sketched above. This view serves as a framework for the discussion in section 8.3 of the diverse literature on various types of scientific explanations. In section 1.8.4 I look beyond philosophy of science and discuss some interesting studies on explanations from the cognitive sciences and science education.

8.1 Hempel and his heirs

The history of the philosophy of scientific explanations can be seen as a story about how the early account presented by Carl Hempel became first the received *correct* view of scientific explanations, then the received *wrong* view of scientific explanations, and recently it came to be cited as a prime example of an *irrelevant* view of explanations. Hempel's account has thus been tremendously influential ever since it was first introduced although the way in which it has been influential has changed over the years. In this section I discuss why Hempel's account of scientific explanations came to be seen as wrong, and discuss the most influential parts of the literature that aimed to succeed where Hempel failed. In section 8.3 I discuss why many philosophers of science are starting to view Hempel's account as not just wrong, but irrelevant.

But let me go back to the beginning[46]: In 1948 Hempel and Oppenheim published a paper titled *Studies in the Logic of Explanation* (reprinted along with later papers on the topic in (Hempel 1965)). The Deductive-Nomological (D-N) model of explanation, which Hempel and Oppenheim presented in their famous early paper and which Hempel continued to elaborate, was the standard view on scientific explanations in philosophy of science for many years. Even after the inadequacy of the D-N model was widely accepted, the research program outlined by Hempel lived on, as the questions about explanations raised by Hempel continued to shape the philosophical debate.

Hempel's aim was ambitious. He wanted to learn what scientific- as opposed to non-scientific explanations are, by identifying the logical structure specific to scientific explanations. Since Hempel's aim shaped the following

[46] Ruben (2012) traces philosophers' interest in explanations all the way back to Plato and Aristotle, but Hempel and Oppenheim's paper from 1948 is commonly considered to mark the beginning of the modern philosophy of scientific explanations.

decades of debate, it is worth pausing a moment to notice a few characteristics. First: It focuses on the *commonalities* between all scientific explanations that set them apart from non-scientific explanation. Relative to this aim the differences that may exist internally in science are, at best, of secondary interest. Second: It focuses primarily on the *structural* characteristics of scientific explanations; more specifically on the *logical* structure of scientific explanations. This is hardly surprising as the philosophy of science community at the time strongly favoured logical analysis as a means to understand science. However, explanations have other kinds of characteristics as well: In other domains the importance of the *functional* characteristics of explanations has, for instance, been emphasized along with the structural characteristics (Rowan 1988).

The D-N model is Hempel's answer to his ambitious aim. The model separates an explanation into two parts: explanandum and explanans. The explanandum is that which is being explained, and the explanans is that which is doing the explaining. So, if I ask you "why x?" and you answer "because y", x is the explanandum and y is the explanans. Hempel's answer to the question of what the structural characteristics of scientific explanations are is, briefly put, that we scientifically explain an explanandum statement *p*, by constructing a (sound) deductive argument, containing at least one law of nature as one of the premises, in which *p* is the conclusion. That is, we explain by deducing from the laws of nature.

Hempel supplemented his D-N model with another, very similar model to account for explanations based on statistical laws. Of course, we cannot *deduce* that John Jones was cured of pneumonia, by learning that he took penicillin, since penicillin is not 100% effective against pneumonia, but, we can at least know that the statement "John recovered from pneumonia" is *very likely* to be true given that he took penicillin. For the purposes of the present discussion the two models can be treated as one.

The D-N model has met criticism on many fronts[47], commonly in the form of more or less thought up counter examples[48]. Poor John Jones gets pneumonia, contracts paresis (but luckily for him and us, not syphilis) (Lipton 2004), he eats contraceptive pills without getting pregnant (Salmon 1971), he does experiments with hexed salt (Kyburg 1965) and contemplates the height of multiple flagpoles (Van Fraassen 1980; Godfrey-Smith 2003), all in the name of proving Hempel wrong. The exercise is useful; we learn that deductive arguments containing laws of nature are not always explanations (so Hempel does not describe sufficient conditions for having an explanation), nor are all explanations deductive arguments containing laws of nature (so Hempel fails to

[47] The D-N model is now standard textbook material in philosophy of science. It is used both to show what a fruitful philosophical theory can look like, and how one might go about attacking such a philosophical theory. Having presented Hempel's model(s), and defended it for a full paragraph, the standard philosophy textbooks usually turn to the philosophers' favourite exercise: Criticizing philosophical theories. The purpose of introducing Hempel's model in the philosophy textbook is namely also to show how various kinds of thought-up counter examples - some of the favourite argumentative tools in classic philosophy of science – can be used as a test of philosophical theories. I have on occasions asked science students to read chapters from such textbooks (e.g. (Godfrey-Smith 2003, ch. 13)) and can testify that they find this habit of introducing philosophical theories only to kill them shortly after rather bizarre. It has made me keenly aware of the need to be very explicit about why these texts are structured in this way, if I use them, and reflect on the appropriateness of using textbooks written for philosophy students in courses for scientists.
[48] For a more detailed discussion of the methodologies used in the classic studies of scientific explanation and the strategies employed to prove them wrong see chapter 4.

describe necessary conditions as well (Scriven 1962))[49]. Furthermore, the D-N model fails to account for what appears to be a very recognizable characteristic of explanations, namely that they are *asymmetric* (Bromberger 1966). This is the conclusion that can be drawn based on John Jones' contemplation of the height of a flagpole. Through a D-N-type argument it is possible to deduce the length of the shadow cast by a flagpole, if the height of the flagpole and the angle between the surface and the sun's rays are known. Thus, we can explain the length of the shadow by the height of the flagpole. So far so good. However, it is also possible to deduce the height of the flagpole from the length of its shadow (and the angle between the surface and the sun's rays). Thus, on Hempel's account, we can explain the length of the shadow cast by a flagpole with reference to the flagpole's height *and* the height of the flagpole with reference to the length of its shadow. This contradicts common judgment which says that the length of the shadow is explained by the height of the flagpole *but not vice versa* (but see (Van Fraassen 1980, pp. 130)). This is a serious problem for the D-N model. Godfrey-Smith (2003) goes so far as to call it "devastating" (p. 193). It seems that most philosophers have reached the same conclusion. And so, the D-N model, once the received view on explanations, became the received *wrong* view of explanations which everyone (including myself) use as an example of how *not* think about explanations. Noticeably, the D-N model did not (yet) become the received *irrelevant* view of explanations. Long after the D-N model was declared dead the questions that led Hempel to develop the model still framed philosophers' thinking about explanations.

Some of the most influential philosophers of explanations sought to succeed where Hempel failed. Scholars like Friedman (1974), Kitcher (1989) and Salmon[50] all tried in their way to answer the same question as Hempel did: What are the common structural characteristics of scientific explanations? Friedman and Kitcher found their answer in various notions of unification. According to Friedman and Kitcher, we explain by unifying our knowledge of nature. On Friedman's account this entails that we gain deeper understanding of nature when we develop more encompassing theories that reduce the number of facts that we have to accept as given. Salmon on the other hand eventually found his answer in his causal mechanical account. Explanations according to Salmon, explain why something is the case by stating what caused it. Explanation thus gets linked to causation (a prevalent intuition in philosophy (see sec. 8.3)). Unfortunately causation is an equally difficult concept to clarify. Salmon developed his own account of causation based on causal processes which I will not go into (see (Salmon 1998, Part III)).

All of these alternatives to Hempel's account have been criticized (Woodward 2011, 2003, ch. 8). As in Hempel's case, more or less thought up counter examples play a central role in this criticism alongside more abstract arguments. Contrary to the critique against Hempel, the critics of these later studies do not generally ague that unificationist or causal mechanical explanations do not exist in some form. What the counter examples are intended to show is simply that not all scientific explanations are either causal mechanical or

[49] One of the main issues in Hempel's model is that it relies on the notion of a law of nature. If one adopts a very strict notion of laws then the counter examples showing non-sufficiency may not work. The prize is that the arguments for non-necessity become even stronger.

[50] Of Salmon's extensive work on scientific explanation his (1990) is among the most cited as it provides a detailed discussion of the fall of Hempel's model and the efforts to correct it. His (1998) presents a collection of his later papers on the topic that I have used as the main reference for his view.

unificationist (see e.g. (Weatherall 2011; Berger 1998)). Thus, while Hempel was criticized for providing neither necessary nor sufficient conditions for explanation (let alone scientific explanation), the critics of these more recent accounts only show that necessary conditions have not been provided. A pluralist (like me) sees this as evidence that explanation is a diverse phenomenon that might not have any very interesting, essential characteristics; others see it as evidence that we have not yet found the best general theory of scientific explanations.

Besides a lot of disagreement, the debates over how to replace Hempel's model highlighted one structural feature of explanations that there is in fact consensus on: Explanations are best conceptualized as answers to questions. This consensus covers the claim that very many explanations are in fact given as answers to explicitly posed questions. Furthermore, it covers the claim that even when an explanation is not given as an answer to an explicitly posed question, the question is either implied in the context or the explicitly stated explanation-seeking uttering can be restated as a question without changing its meaning. As an example of the latter consider a teacher's requests to a student: "please explain how the heart works". The explanation given as a response to this request can be understood without significant distortion as an answer to the question "how does the heart work?". Questions that ask for an explanation are commonly referred to as *explanation-seeking questions*. What characterizes explanation-seeking questions has yet to be agreed upon, so the consensus that explanations are answers to these questions is not very deep. Explanation-seeking questions are commonly equated with why-questions, although some will be quick to point out that this is an approximation, as we equally often explain *how* (see (Faye 1999; Salmon 1998; Cross 1991; Scriven 1962)).

The realization that scientific explanations are answers to explanation-seeking questions is not generally seen as a satisfactory answer to the question of what the structural characteristics of scientific explanations are. First of all, this is because non-scientific explanations are also answers to explanation-seeking questions. Being an answer to an explanation-seeking question thus cannot be a sufficient condition for being a scientific explanation. Secondly, simply being an answer given to an explanation-seeking question does not necessarily make an answer an explanation. If I ask my teacher why the sky is blue, and he answers "yesterday it was cloudy", he has responded to my question, but not provided me with an explanation of why the sky is blue. This suggests that a certain *relation of relevance* between the explanation-seeking question and the answer must exist before the answer can be considered an explanation.

Identifying explanations with answers to explanation-seeking questions thus leaves an important question to be answered before the aim of knowing the structural characteristics of scientific explanations is reached:

1. What are the (structural) characteristics of *scientific* answers to explanation-seeking questions?

In light of the considerations on relevance the better question is:

2. How can the *relation of relevance* between an explanation-seeking question and their scientific answers be specified?

The latter question has been dealt with explicitly in the influential pragmatic strand within the Hempelian research program seeking to characterize scientific explanations in general. The pragmatic view of explanations is most noticeably represented by Bas van Fraassen ((Van Fraassen 1980, ch. 5), but see also (Achinstein 1983)).

Contrary to the accounts of scientific explanations discussed previously which all maintain that judgments of the quality of explanations can be made independently of any individual's knowledge, pragmatic accounts of explanations emphasize the *context and audience dependence* of explanations. I shall describe van Fraassen's account of what explanations are in some detail[51], as it forms the basis for my own account presented in sec 8.2.

On van Fraassen's account explanations are relevant answers to why-questions. What makes an answer relevant cannot be further specified in general terms due to the context and audience dependence of explanation. *Scientific* explanations are simply explanations that somehow draw on scientific knowledge (1980, p. 155). This is a rather bold, and to many a highly unsatisfying, answer to question 2. posed above. To see how van Fraassen arrives at this answer we must consider his arguments in more detail.

To reach his answer van Fraassen starts out by assuming that explanation-seeking questions can be equated with why-questions. On his view a theory of explanation can thus be reduced to a theory of why-questions (p. 141). As mentioned, this is arguably an approximation, and van Fraassen fails to recognize this explicitly. Before moving on, let us consider how appropriate this approximation is. If researchers ask relatively few explanation-seeking how-questions, then the approximation is appropriate, but if they do, then a lot of information would be lost if, for instance, an empirical investigation of explanations is performed by exclusively focusing on answers to why-questions. It turns out that scientists and educators are in fact often interested in explaining how, and that van Fraassen's approximation is therefore very significant (see sec 8.3, chapter 4 and (Dagher, Cossman 1992))[52]. This makes his theory of why-questions less useful for the purpose of characterizing explanations in science and science teaching and thus less useful for my purposes (chapter 4 and 5). In section 8.2 I therefore outline of a more inclusive pragmatic view of explanations which largely incorporates van Fraassen's theory of why questions as a special case. But for now I follow van Fraassen in making this approximation.

So, to get an explanation we ask "why x". Or rather, van Fraassen argues, we ask "why x and not (y,z,v, …)" where (y,z,v,…) is the *contrast class* related to the specific why-question posed. Thus, according to van Fraassen the question "why is the sky blue?" cannot be answered unless it is understood what the implicit contrast class is. Does the questioner want to know why the sky is *blue* (and not red, yellow or green), or does he want to know why the *sky* is blue (and not the trees or the clouds or ...)? The relevant answer to a given why-question

---

[51] Van Fraassen also gives first a rough and later more detailed account of what a *good* explanation is (1980, § 4.4). I will discuss and largely adopt the rough account that van Fraassen gives, but I do not adopt his more detailed account based on the notion of personal probabilities. For a critical discussion see (Achinstein 1984).

[52] Cross (1991) showed that it is relatively straightforward to generalize van Fraassen's theory to cover how-questions as well. Although a definite improvement of van Fraassen's theory, Cross' theory is still not sufficiently general for my purposes, and suffers from similar weaknesses as van Fraassen's account with respect to the detailed account of what makes an explanation good (cf. previous note).Therefore I have not adopted it here.

depends heavily on which contrast class is associated with it, which in turn depends on the knowledge and interests of the person(s) who pose the why-question. The context dependence of explanations does not end here according to van Fraassen. The task of the explainer is to provide a relevant answer to the person(s) posing the specific why-question, "why x and not (y,z,v, …)?". This implies that the answer provided must remove the need expressed by the person who posed the question to pose the specific question. The question was posed because the questioner believed that x was the case, but did not know why this particular possibility was realized rather than other - to the questioner - conceivable possibilities (y,z,v,…). The task of the explainer – given that she accepts that x is in fact the case - is thus to make it clear *to the one(s) who posed the question* why x rather than (y,z,v,…) is the case. This implies that the explainer must consider what the poser of the question already knows when constructing her answer to a given why-question, otherwise she risks providing an answer that does to help the receiver. Thus, even a scientist does not always give the same answer to the same explanation-seeking question. If the question is posed by her three year old daughter she will give a different answer than if the question was posed by a scientist from a different domain, which would again be different from the answer, she would give to a close colleague. If the structure of an explanation - scientific or not - depends heavily on who the explanation is given to, then it becomes rather difficult to say anything interesting in general terms about the structure of explanations beyond what has already been said. So, while van Fraassen starts out by considering Hempel's general question of what characterizes scientific explanations in general, his answer is at the same time a rejection of this question. If van Fraassen is right, then it is simply not very interesting to continue to consider common features of scientific explanations in general, as the structures of explanations given by scientists are so heavily context dependent. Van Fraassen draws the further conclusion that scientific explanations *per se* should not interest philosophers of science, but this does not follow from his account of explanations alone (Cross 1991). Even within van Fraassen's framework it makes perfect sense to argue that it is possible to say something relatively general about explanations given within groups that are fairly homogeneous with respect to interests and background knowledge – e.g. what I in section 3.2 defined as scientific domains. Thus, even if it is not possible to say very much interesting about the general structure of explanations given by scientists to any odd audience, this does not preclude that it is possible to say something about explanations given by scientists from a given domain to other (core or periphery) members of that domain (cf. sec 3.2). In fact, the literature discussed in sec. 8.3 shows that this is in fact possible.

Contrary to the other accounts of scientific explanations that have been discussed so far, van Fraassen's account has not met much fundamental criticism (but see (Salmon, Kitcher 1998))[53]. This does not mean that his view is now generally accepted. Rather, it should be seen as an expression of a shift in present day philosophy of explanation away from Hempel's general question and towards studies of *specific types* of scientific explanations. Once Hempel's question is no longer in focus the answers to the question provided by Hempel and others shift from being potentially wrong to being irrelevant relative to the current debates.

[53] Van Fraassen's version of anti-realism - constructive empiricism – which his pragmatic theory of explanation is part of has meet with much fundamental criticism, and is currently a favourite enemy in the realism vs. anti-realism debate (Chakravartty 2014).

Recent literature on explanations focuses increasingly on explanations from specific domains (noticeably exemplified by the vast literature on mechanistic explanations in the biological sciences), or on specific types of explanations across domains (noticeably the Salmon inspired literature on causal explanations), or on specific types of explanations within specific domains. I will return to these literatures in section 8.3. For now I would like to discuss some consequences of having made this move to study more specific kinds of explanations in science before the general debate on the general characteristics of explanations was settled.

In section 1.1 I claimed that the studies presented in chapter 4 and 5 of this dissertation fill a hole in the philosophical literature because differences in explanatory standards across scientific domains have not been addressed in detail. In the previous paragraph I claimed that philosophers have been interested in the specific characteristics of explanations from specific domains. If this is true, then it might be argued that philosophers have actually addressed differences in explanatory standards, albeit only indirectly, through their domain specific studies of explanations. It thus seems that differences in explanatory standards could conveniently be studied through a meta-analysis of existing studies of explanations from different domains.

The existing domain specific literature on explanations is definitely useful when studying differences in explanatory standards across scientific domains (see chapter 5), but simply performing meta-studies would not be a feasible approach. This is partly a consequence of the missing consensus in philosophy of science on what the general characteristics of explanations are, and partly because only a fraction of the existing science domains are covered by the existing literature. (How do geologists explain? To the best of my knowledge there have been no specific studies of this). Even meta-studies focusing only on the domains that have been covered would face severe methodological challenges. Individual studies of domain specific explanations are based on widely different definitions of what scientific explanations are. Some follow van Fraassen and take explanations to be answers to why-questions (e.g. (Goodwin 2003)). In other parts of the literature, particularly in the mechanism literature, how-questions are in focus (see sec. 8.3.2). Many studies do not even state what notion of explanations is employed, but given that other studies explicitly employ different definitions it would be too optimistic to hope that these studies all employ the same implicit definition. This observation raises two methodological questions when attempting to compare the existing studies: Are they in fact comparable? And if they are: how trustworthy will the results be? The first question can be further explicated as follows: Are all the different studies claiming to deal with explanations actually dealing with different types of the same general phenomenon, or are they dealing with distinct phenomena? The answer to this question depends on the general view of explanations adopted. If one follows the common trend within Hempel's research program and equates explanations with answers to why-questions then one would not be comparing explanations when comparing answers to why-questions from chemistry to answers to how-questions from biology. Rather, one would be comparing explanations from chemistry with descriptions from biology. Any differences identified would thus be differences between explanations and descriptions, not between types of explanations. To defend the claim that meta-studies actually give meaningful results *about explanations*, a more inclusive view of explanations is needed, than the one commonly defended. The view introduced in the beginning of this section 8 and defend in the next sub-section is one such view. Even if such an inclusive view is adopted, the differing views of explanations adopted across the literature mean that there are still severe methodological problems facing the meta-studies looking for differences in explanatory standards across domains. Comparing the existing studies will show that there are differences, but are these differences real or simply an artifact of the many different general definitions of explanations employed? This would be very difficult to discern from the meta-studies alone and would thus have to be tested through independent methods before robust conclusions can be drawn. This is part of the reason why I have adopted the more empirical methodology developed in chapter 4 and applied in an adapted form in chapter 5 to the case of molecular biology and polymer physics.

Given these considerations I turn to present a somewhat detailed discussion of the view of explanations introduced in the beginning of this section, before returning to the more specialized literature on explanations in section 8.3. In section 8.4 I take a selective view on the literature on explanations from the cognitive sciences and science education

8.2 A pragmatic-functional characterization of explanations

There is no deep consensus on what the structural characteristics of (scientific) explanations are. However, there is at least a superficial agreement on what the main *functional* characteristic of explanations is: We seek explanations to gain *understanding*[54]. *Good* explanations thus provide (a lot of or deep) understanding. All of the general studies of explanations mentioned until now have in one way or another tried to answer *how*, i.e. through what structural features, explanations provide understanding. This would be a lot easier if we had a good description of what it means to understand something (anything!). Unfortunately we do not. In fact, it has turned out to be just as difficult to characterize understanding as it is to characterize explanations. Therefore, Salmon concluded that:

> ***The use of such near-synonyms as "understanding" and "comprehension" [makes explanations] sound important and desirable, but helps not at all in the philosophical analysis of explanation. (Salmon 1998, p. 126)***

Rather, Salmon, like most of his predecessors and contemporaries, refrained from analyzing understanding independently, and focused on finding out *how* explanations provide us with understanding and, through this, finding an answer to *what* understanding is. Of course, since they all came up with different answers to the how-question they also gave different answers to the what-question.

The approach taken by Salmon and others has yielded fruitful lessons about explanations, but not an acceptable structural definition of a scientific explanation. On the other hand, understanding, as a topic of intrinsic interest, has received increased attention in philosophy of science and interesting results have appeared. So, perhaps it is time to reconsider Salmon's claim that the linking of explanations to understanding is of no help at all in the philosophical analysis of explanations.

In the following I shall therefore briefly review recent philosophical studies on understanding and use this as a way into answering the question *what is an explanation*?

[54] There is even empirical evidence that both practicing researchers and lay people are very reluctant to consider something an explanation if it does not provide understanding (Waskan et al. 2014).

8.2.1 Understanding

Direct philosophical analysis of scientific understanding (as opposed to the indirect study of the Hempelian research program) has only recently received much attention in philosophy of science. Basic clarifications are still being made, as philosophers discuss what kind of thing understanding is. One thing is settled: The *feeling* of understanding - the aha-experience - should be separated from genuine understanding (Trout 2002). The feeling of understanding is (a not always reliable) *indicator* of understanding, but it is not the core phenomenon we are interested in.

Traditionally understanding is viewed as *some kind of knowledge*. The knowledge of the person who understands is different from that of the person who does not understand, and this is the crucial difference. The difference need not be in quantity, but rather in quality. On Hempel's account we understand a phenomenon when we know that it was to be expected given the laws of nature. On the unificationists accounts we gain understanding by unifying our knowledge of nature (see also (Kosso 2007)), and Salmon discusses how Pascal's demon, who knows every regularity in the universe, still does not understand the world, because it has neither causal nor deductive knowledge (both would be preferable) about why these regularities hold (Salmon 1998, ch. 8).

A recent account by Ylikosky (2009) challenges the traditional view and argues (following Wittgenstein) that understanding is an *ability*. Ylikosky's account is a somewhat extreme example of a more general trend in recent philosophy of understanding to focus on the practical virtues of understanding. Similar to trends in education[55], philosophers increasingly focus on what the person can do with the knowledge she has gained rather than simply focusing on what the person knows. Ylikosky's account is radical in the sense that it *reduces* understanding to the ability to act, whereas others take the ability to act as another *indicator* of some cognitive state that can be called understanding (van Camp 2014)[56].

The current focus on the practical side of understanding is largely due to the influential account by de Regt and Dieks ((De Regt, Dieks 2005) see also (de Regt 2009)). In their account it is a little unclear whether understanding is a kind of knowledge or an ability, but it focuses very much on the fact that understanding – whatever it is - allows people to *do* certain things. In the vocabulary of Dieks and de Regt a theory is intelligible to scientists if they can "recognize characteristic consequences of [the theory] without performing exact calculations" (De Regt, Dieks 2005, p. 151)[57]. A phenomenon, P, is in turn understood if an intelligible theory of

[55] There is some discrepancy in vocabulary between the science education and philosophy of science literatures. What philosophers characterise as understanding is probably best compared to what educators call learning. Recent taxonomies of learning (e.g. (Anderson, et. al. 2001)), emphasise that high levels of learning imply the ability to act. In these taxonomies understanding tends to have a more specific meaning as passive understanding (as in "I understand what you say").

[56] Wilkenfeld (2013) also argues that understanding is a cognitive phenomenon, but goes on to argue that understanding some phenomenon means having an adequate mental representation of that phenomenon. What the relation between metal representations and knowledge is on Wilkenfeld's account is unclear.

[57] It is striking to note that Collins and Evans argue that this is exactly what an interactional expert is able to do (cf. Chapter 2). It seems that the major difference is that according to Collins and Evans this is all that the interactional expert is able to do, whereas the contributory expert (the person who presumably has the deeper understanding) *could* also perform the exact calculations but need not do so.

P exists, and the explanation of P is acceptable relative to the relevant explanatory standards. Thus, according to de Regt and Dieks, understanding is primarily about being able to do certain things in a certain way. Whether these abilities are to be equated with understanding or whether they are a consequence of some cognitive state that should be called understanding is unclear from de Regt and Dieks account.

I follow van Camp (2014) to the general conclusion that understanding and knowledge are the same kind of thing (or kinds of things)[58]. Furthermore, I agree that the most reliable indicator of understanding is a person's ability to act in ways that are deemed valuable by relevant experts, and that we often seek understanding in order to gain this specific ability to act.

The notion that knowledge is justified true belief is so deeply entrenched in many parts of philosophy that many philosophers automatically assume that understanding is a kind of justified true belief when they learn that understanding and knowledge is the same kind of thing. This is a mistake, and makes the difference between viewing understanding as an ability and as a kind of knowledge look more significant than it is.

Regardless of whether or not explicit knowledge can be defined as justified true belief, *tacit* knowledge cannot. Often, performing the tasks that express understanding (and are used to measure it) requires tacit knowledge. The account by de Regt and Dieks illustrates this. To recognize characteristic qualitative consequences of a given scientific theory a person must reason with the theory in relation to a concrete problem. This creative process cannot be understood as a rule following process, but must be seen as a process involving skills and tacit knowledge ((de Regt 2009) see also (Kuhn 1996, pp. 187)). Thus, on the account by de Regt and Dieks having understanding implies having both explicit and tacit knowledge. The definition of understanding presented by de Regt and Dieks does not cover all kinds of scientific understanding, but should rather be seen as an important explication of a way in which understanding is expressed in physics and other parts of science where exact calculations are commonly used as a tool of reasoning (Wilkenfeld 2013). Still, their account illustrates that tacit knowledge can be a very important part of understanding. (For other illustrations see (Leonelli 2009; Galison 1997)). Adopting this broader view of understanding-as-knowledge makes it look very similar to the understanding-as-ability view defended by Ylikosky. After all, gaining an ability involves mainly the gaining of tacit and explicit knowledge. And thus, what initially looked like an ontological disagreement may turn out to be two sides of the same coin[59].

If understanding something means having relevant knowledge related to that thing which allows the person to act in certain ways, two questions arise:

[58] van Camp goes on to argue that knowledge is justified true belief, and adds a rather naive realism to construct a decontextualized account of understanding. I do not follow him in this respect.

[59] Of course, one might go a step further and argue that a person is not really able, for instance, to perform experiments without a laboratory and limbs to operate it, but this does not seem to magnify the difference between the understanding-as-knowledge and understanding-as-ability accounts. Ylikosky does maintain that *persons* understand which seems to imply that he does not want to claim that you need to have a lab in order to have the ability (in his technical sense) to experiment. Intuitively, a person loses her ability to experiment if she loses her arms, and thus it seems that on Ylikosky's account a person who loses a limb loses understanding along with the limb. Whether a person also loses tacit knowledge - and thus potentially understanding on the understanding-as-knowledge account - is not clear, but it does not seem to be precluded.

1. Is it possible to say something more about what knowledge is relevant?
2. Is it possible to say something more about the way in which the relevant tasks must be performed?

When it comes to understanding in science it seems that the latter question must be answered with reference to the aims and epistemological standards of the relevant science domains. Ultimately, a person is said to have understanding of something used in science (theories, models, physical objects etc.), if she can use it to perform the tasks that are valued in the community in a way that is satisfactory according to the epistemological standards of the community, for instance constructing novel explanations, building devices, performing experiments or simply satisfying personal or peer curiosity on related topics[60]. A theory is understood if it can be used to solve relevant problems in appropriate ways; experimental settings are understood if they can be used to perform good experiments etc. This means that understanding, linked as it is to the changing epistemological standards of science domains, does not have a static meaning, but is subject to change over the course of history.

Adopting the view that understanding of science subject matter is always relative to a set of epistemological standards has a number of consequences. For one it implies that it is possible for two people to have a relatively complete understanding of the same class of objects even though they cannot do exactly the same with these objects, because their understanding is complete relative to the epistemological standards of different domains. This fits well with the discourse on interdisciplinarity, where domains contributing to an interdisciplinary problem solving are generally considered to be of equal value. The view of understanding defended here also means that gaining understanding has an important social component. In many cases it will not be possible to gain sufficient tacit knowledge through solitary reading and problem solving. Rather, skills must be learned in a social process of practice and feedback from experienced peers ((Leonelli 2009; Collins, Evans 2007) see also chapter 2)[61].

Furthermore, this view of understanding means that it is possible to gain some understanding from false theories. Using phlogiston theory scientists in the 15^{th} and 16^{th} century were able to perform very many tasks that were deemed valuable. They constructed explanations of multiple phenomena, carried out experiments, and found solutions to concrete problems that passed as good research according to the epistemological standards of the relevant part of science at the time, partly because some of the solutions made an actual difference to people's lives. Of course they were not able to do everything they wanted to do in a satisfactory way, and phlogiston theory has been abandoned partly because of this, but that does not mean that researchers did not gain *some* understanding of nature through phlogiston theory. This claim contradicts a

[60] In this way explanations are both an end and a means to an end in science. We commonly seek explanations in order to be able to construct other explanations or do other things. Occasionally, we will seek explanations simply to satisfy our curiosity, but at least in science explanations seem to be most valued when they also enable either the posing of new interesting questions or construction of novel explanations, experiments, devices etc. (cf. (Lakatos 1978)).

[61] In this respect, as well as in many others, there is a close relation between understanding and expertise. Although the philosophical literature on understanding reviewed in this section has not so far been connected to the more sociology dominated literature on expertise (discussed in chapter 2), it seems clear that the two literatures are largely dealing with the same phenomenon (which, cf. note 55, is again similar to the 'learning' phenomena discussed in the science education literature).

prevalent intuition in parts of philosophy of science (partly coinciding with the parts that are strongly committed to the claim that the ultimate aim of science is truth) which holds that although it may be possible to understand a false scientific theory, it is not possible to understand *nature* through a false scientific theory (e.g. (Trout 2002; Strevens 2013; van Camp 2014)). This more platonic way of thinking about understanding, as a kind of deep knowledge of nature may be relevant when contemplating the ultimate aims of science. Indeed, it does seem to deal with a more abstract kind of understanding that "science" has achieved, rather than the more common notion of understanding as something a person can gain. In analogy to the distinction by Wagenknecht (2014) between 'knowledge' as the impersonal knowledge found in books and encyclopedias and 'knowing' as a personal belief state, we might thus distinguish between Understanding (with capital U) as the impersonal complement to knowledge and understanding (with lower case u) as the personal complement to knowing. For discussions about everyday science and teaching practice the more abstract notion of Understanding is not very useful as it is a very different kind of understanding than what is commonly taught and tested. My concern is therefore primarily with the latter type of understanding. This kind of understanding comes in degrees (we do grade student performance), and is commonly measured by the person's ability to perform various tasks in accordance with the epistemological standards of the relevant domain.

Returning to question 1 posed above, about whether it is possible to say something more about the knowledge relevant for understanding, it may be argued (in analogy to van Fraassen's arguments on explanations) that given the diversity of things that it is possible to understand – processes, methods, theories etc. – it seems unlikely that it is possible to say something very interesting and at the same time very general about the knowledge needed for understanding. Rather it seems prudent to assume that there are *many kinds* of understanding. The best we can hope for in terms of general principles is some general notion analogous to van Fraassens relevance relation, but even such a vague principle will be useful for my purposes.

Philosophers who take understanding to be a kind of knowledge do seem to agree on the point that for understanding of something to be achieved, the explicit knowledge that one has related to this thing cannot simply be a collection of facts. (Learning a chemistry data book by heart does not provide understanding). Rather, the knowledge on the given topic must show some degree of *connectedness*[62] (e.g. (van Camp 2014; Kosso 2007; Friedman 1974; Scriven 1962))[63]. There is no consensus on the nature of this connectedness, and probably a multitude of types exist. In some cases a unified knowledge a la Friedman will be valuable, whereas it may be other kinds of connections, e.g. connections through categorization as cause and effect that are valuable in other contexts. For my current purposes these details do not really matter. Suffices to say that understanding can be understood as some kind of knowledge (explicit and tacit combined), which is characterized not only by consisting of relevant pieces of knowledge, but also of exhibiting relevant types and degrees of connectedness, at least for the explicit knowledge component.

[62] Those who in addition emphasize that understanding also has a tacit component further emphasize that the tacit and explicit components must show some degree of coordination (Leonelli 2009; Galison 1997).

[63] Cf. note 56 Wilkenfeld (2013) argued that understanding means having a mental representation that can be manipulated. Regardless of whether this amounts to a claim that understanding is a kind of knowledge, it does at least seem to imply that a person's knowledge about a given topic is connected through the metal representation when understanding is present.

With these general considerations on understanding - the aim of our quest for explanations - we can return to a characterization of what explanations are.

8.2.2 Explanation-seeking questions and explanatory answers

The preceding discussions on the philosophical literature on understanding and the studies in the Hempelian research program on explanations have revealed some important characteristics of explanations that there seems to be general consensus on:

- Explanations answer explanation-seeking questions in order to provide understanding.
- The explicit component of understanding is characterized by some kind and degree of connectedness.
- We often seek understanding in order to be able to perform certain tasks well (relative to the relevant epistemological standards).

Combining these insights we can move towards an answer to the question: what is an explanation?

Explanation-seeking questions are most commonly posed out of a perceived lack of understanding. In teaching practice students pose many explanation-seeking questions for this reason. But part of a science education is also to learn to ask the right kinds of questions, and the teacher (or the textbook) will therefore in addition pose a series of explanation-seeking question for the student, knowing that students lack understanding of a given matter, and that the answer provided will provide at least some of the lacking understanding[64]. When posing an explanation-seeking question the hope is of course, that the person(s) responding to the question can provide this lacking understanding. Understanding is not just a collection of facts, but a connected (and coordinated) body of knowledge. And so, a characteristic of relevant answers to explanation-seeking questions – i.e. explanations – is that they must provide not just facts but also relevant connections between new knowledge provided and the existing knowledge of the person(s) the answer is given to[65]. Conversely, explanation-seeking questions are characterized by being questions that ask for connections in addition to facts.

In this way we can distinguish between two general types of questions: explanation-seeking questions and fact-seeking questions. Explanation-seeking questions were discussed above. Fact-seeking questions are requests to the respondent for a piece of knowledge that will fill in a hole in the question-poser's knowledge. We may characterize answers to fact-seeking questions, or questions perceived to be of this kind by the respondent, as *descriptions*.

From this distinction it follows that it is not the question itself, i.e. the wording of the question, but the function of the question that decides whether a question is explanation-seeking or not. The function of a question phrased in a certain way may change depending on the context. Some questions, in particular what-questions like 'what is the diameter of the earth?', may be explanation-seeking when posed in one context, but

[64] Alternatively explanation-seeking questions may be posed by teachers as a test of the understanding of the person(s) the question is posed to.
[65] Importantly this connection may in some cases be a replacement. The explanation provides new knowledge and shows how it rejects previously held beliefs and alters the knowledge system of the individual.

fact-seeking when posed in another[66]. However, since we do not have direct access to what the question poser intended when posing the question, we need to have some *reliable linguistic indicators* in order to avoid constant misunderstandings. Indeed, it seems that certain question structures, like ‘what is x’ and ‘how much (of unit y) is x?’ are used mainly for fact-seeking questions, while other structures like ‘why x’ and ‘how does x work’ are reserved mainly for explanation-seeking questions[67]. These conventions hold to such an extent that it is sometimes necessary to clarify that a question that looks like a fact-seeking question is in fact explanation-seeking. In textbooks, for instance, exercise questions often have the structure ‘how much (of unit y) is x?’, which is characteristic of fact-seeking questions. So in order to convey to the student that (s)he should answer this as an explanation-seeking question the authors will often offer a clarification like “Explain your reasoning!” or just “Explain!” right after the exercise question (see chapter 4). So although we cannot *define* explanations as relevant answers to questions involving certain interrogatives or questions with certain structures, we can view interrogatives like ‘why’ and some question structures, such as ‘how does X work?’, as reliable indicators of explanation-seeking questions. (For further discussion see chapter 4.)

Explanations are commonly conceptualized as answers to explanation-seeking questions. However, as van Fraassen points out (see sec 8.1), not all responses to explanation-seeking questions can be considered explanations. A response that in the given context does not in any way address the lack of understanding that lead to the posing of the explanation-seeking question is irrelevant and should, following van Fraassen, not be considered an explanation. From this consideration, and the observation that a respondent cannot know for sure whether a question posed is explanation-seeking or fact-seeking, we reach the general characterization of explanations presented earlier[68]:

> An **explanation** is a non-irrelevant response to a question taken by the respondent to be explanation-seeking.

The above characterization should not be read as a definition, as I have not argued that it provides necessary conditions for what an explanation is[69]. Nor is it a characterization of *good* explanations. Rather it is a very general characterization that captures the consensus of the literature and can serve as an inclusive framework for analyzing both the more specific literature on types of scientific explanations and explanations in practice.

Comparing my characterization to van Fraassen’s characterization of explanations as relevant answers to why questions reveals that I have not just replaced “why-questions” with “explanation-seeking questions”. I also speak of “responses” rather than “answers” and use “non-irrelevant” rather than “relevant”. Why? First of all

[66] If the question “what is the diameter of the Earth?” is simply a request for a number then it is fact-seeking, but it may also be posed by a person who has realized that the earth is not a perfect sphere and is therefore wondering what the convention behind this phrase is.
[67] Interestingly, why-questions seem to be almost exclusively explanation-seeking.
[68] Faye (1999) reaches a very similar characterization independently of an analysis of understanding. The characterization sketched by Scriven (1962) is also similar, although it does not conceptualize explanations as answers to questions.
[69] I do agree with Faye, that explanations are necessarily results of acts of explaining which seems to exclude the various ontic accounts of explanation, most explicitly defended by Craver (2007), where explanations exist independently of human actions. On the view defended here, explanations are *constructed*, not discovered.

the differences do not express deep disagreements with van Fraassen, rather a wish to clarify further what I take van Fraassen to mean and to highlight another common approximation made in the philosophical literature on explanations.

I take an answer to a question to be a sentence or a set of sentences, and thus by defining explanations as relevant answers they are defined as linguistic entities. This is very common in philosophy (e.g. (Lipton 2009; Achinstein 1984; Hempel 1965)). Perhaps this assumption is a relic of the early attempts to characterize explanations as *arguments* with specific characteristics. This is most easily done if explanations are thought of as being sets of propositions if not sentences. In practice however, even printed explanations are not always made up entirely of sentences: Pictures, diagrams and graphs take up much space in science textbooks and are frequently referred to in explanations. As we shall see in chapter 5 some of the most interesting differences in textbook explanations from physics and biology are exactly in the way graphical representations are used in explanations. In classroom and laboratory practice, actions (other than speech-acts) serve as important mediators of tacit knowledge from master to apprentice and thus also play important roles in the transfer of understanding. Some of these actions are performed as (parts of) a response to an explanation-seeking question and should thus be counted as part of the explanation[70]. Therefore, I speak more generally of responses to explanation-seeking questions while at the same time acknowledging that very many explanations are, to a good approximation, linguistic entities.

When van Fraassen claims that an answer to a why question must be relevant to be an explanation the question arises: how relevant? My reading of van Fraassen is that his answer to this question would in principle be that as long as the answer to the why-question is not irrelevant then it is an explanation. I have transferred this reading to my characterization. This reading means that there are, in a given context, very many explanations of a given explanandum. This does not mean that there are equally many *good* explanations; in fact it probably means only that there very many very bad explanations.

'What makes an explanation good?' is a question at the core of this dissertation in general and chapters 4 and 5 in particular. In general, the quality of an explanation depends of whether it fulfills its function: to provide understanding. A good explanation provides new knowledge and alters connections between pieces of knowledge in a way that alters, in a positive way, the way in which the receiver of the explanation can perform the tasks that are taken to be indicative of understanding in the relevant social context. In the context of science we can interpret "in a positive way" as meaning that the receiver of the explanation becomes better equipped for producing explanations, experiments, arguments etc. that comply to the epistemological standards of the relevant community. If science had just one aim and one set of epistemological standards,

[70] Can an explanation be completely silent? I recall once having problems opening a bottle of acetic acid. I asked my father: "how do I open this?". The response was swift but silent: He took the lid off slowly showing me how he opened it by pressing the sides of the lid and turning it at the same time. Then he put the lid back on and gave me the bottle back. Afterwards I knew how to open bottles with that type of lid. By my characterization, my father's response is an explanation (it is even a good one). Had I claimed that I provided a definition of an explanation, this would have been a problem. My characterization does not capture the distinction made in English between an explanation and a *demonstration* which seems to be what my father performed. However, since this distinction does not matter for my purposes I shall not address it further.

then there would be only one kind of scientific understanding. But it appears that science is not as homogeneous as was once thought, and with the heterogeneity of the aims and epistemological standards of scientific practice comes the heterogeneity of explanation and understanding. The diversity in explanatory practice in science has meant that many philosophers have chosen to narrow the scope of their studies of explanations. Many recent studies of explanations focus not on scientific explanations in general, but more specifically on particular types of explanations. At least this is the most natural reading of the recent literature given the inclusive general view of explanations outlined above.

To complete the tale of how Hempels D-N model went from being received correct view of explanations over being the received wrong view of explanations to become cited as a good example of a irrelevant view of explanations the following section provides a selective survey of the recent philosophical literature of different types of scientific explanations. Section 8.4 then points to some interesting studies of explanations originating from outside philosophy of science. The sections do not add significantly my own account of explanations presented above, and are not drawn upon in later chapters. Readers who are familiar with the recent literature of types of scientific explanations can thus jump directly to section 9.

8.3 Hempel becomes irrelevant: Types of scientific explanations

The aim of Hempel's research program was to characterize a particular type of explanations: Scientific explanations. This research program yielded significant results, but is no longer the dominant research program in philosophy of explanation. New research programs are evolving, aiming to answer slightly different questions, and to some extent using different methods than were used within Hempel's program. A research program on causal explanations has developed in the wake of Salmon's work on explanations in general. In parallel to this program, a research program on mechanistic explanations in the biological sciences, a particular type of causal explanations, has developed over the past 20 years. Mathematical explanations, both in the sciences and within "pure" mathematics, are receiving increasing attention[71], and interest has been developing in characterizing explanations in chemistry (see sec. 8.3.3). Large parts of the programs focusing on explanations within particular domains are more practice oriented (sec. 7) than traditional philosophy of explanation. Although some of these research programs have been around for quite some time and some are only just emerging, I shall refer to them collectively as *The New Programs*.

While there are important differences between the aims and methods both among The New Programs and between The New Programs and Hempel's program, there are also similarities. One important similarity between The New Programs and Hempel's program is that they all aim to characterize a specific type of explanations. While Hempel was interested in scientific explanations in general, the more recent programs aim to characterize specific types of scientific explanations. As noted in section 8.1 Hempel's program focuses almost exclusively on commonalities among scientific explanations. Similarly, we find that The New Programs focus mainly on characterizing commonalities within the type of explanations that are of interest to the particular program. So, as with Hempel's program, The New Programs only find differences within the

[71] I will not discuss these types of mathematical explanations further. For an interesting example of mathematical explanations in science see (Baker 2005), for a review of the literature on explanations in mathematics see (Mancosu 2011).

individual types and across types of scientific explanations of secondary interest. The program on causal explanation, for instance, focuses mainly on what causal explanations across the sciences have in common and not, for instance, on detailed comparisons of causal explanations from different sciences. However, one does find some contrastive discussions within some of The New Programs, mainly with the aim of distinguishing the individual programs from Hempel's program. Again, Hempel himself is the favorite enemy, although Salmon also gets his share of criticism. As might be expected, the strategy, when discussing research from another research program is not so much to argue that the general characterizations of scientific explanations are *wrong*, but more that they are *irrelevant* relative to the aims of The New Programs[72].

To illustrate, consider Bechtel and Abrahamsen's (2005) influential paper on mechanistic explanations. The paper starts out by citing Hempel's D-N model (discussed in section 8.1) as "the received view of scientific explanation in philosophy" (p. 421). As argued in section 8.1, this claim had not been true for many years in 2005 (cf. (Salmon 1990)), but it serves to highlight the mechanistic program as something new and different. Typically, a paper on in the Hempelian tradition would now go on to a somewhat detailed discussion aiming to show the weaknesses of the D-N model as a general model of scientific explanation, most likely listing one of the many counter examples discussed in section 8.1. Bechtel and Abrahamsen do not adopt this strategy, as their primary aim is not to succeed where Hempel failed but rather to characterize *actual explanations in the biological sciences*. Their aim is thus different from Hempel's and there is therefore no need to argue extensively for why Hempel does not reach his own aim. Instead, Bechtel and Abrahamsen argue briefly that the D-N model is irrelevant for their purposes as "most actual explanations in biology do not appeal to laws in the manner specified in the D-N model" (p. 422). Since the D-N model cannot be used as an account of actual explanations in the biological sciences, Bechtel and Abrahamsen find that there is a need to supplement with an account of mechanistic explanations[73].

Often the contrastive discussions between the New Programs and Hempel's program draws on a more general line of criticism against general philosophy of science that much of what is presented as general philosophy of *science* is in fact philosophy of *physics*. In the context of explanation the argument is that the purported general literature captures, at best, only what is going on in physics, partly because the general literature draws exclusively on physics for scientific examples, and completely ignores the differences in practice across the sciences. As illustrated above, Hempel is often criticized for giving the laws of nature such a central role in his D-N model. The argument is not so much that Hempel's account of lawhood is wrong, but rather that general statements that could count as laws on Hempel's account play no essential roles in actual explanations outside physics and physical chemistry (see also (Mitchell 2000)). So again, Hempel's model is not so much proven wrong, as it is proven irrelevant to the study of explanations in the special sciences.

[72]As a corollary, these arguments do show that the general studies cannot have captured the essence of a scientific explanation and thus must be wrong on their own terms in addition to being irrelevant, but this is not the main focus of the discussions.
[73] For a similar argument to the effect that the D-N-model is irrelevant when studying explanations in chemistry see (Akeroyd 2008).

Salmon's causal mechanical account gets a similar treatment by Glennan (2002) in another of the influential papers within the mechanism program. Glennan justifies the need for an account of mechanistic explanations by arguing that while Salmon's general account of causal explanations, and especially his reliance of the notion of a causal process, may fit well with classical mechanics it does not capture the causal explanations for instance in the life sciences. Glennan acknowledges that Salmon's account has certain advantages from a logical perspective as it "avoid[s] certain objections to the counterfactual approach" (p. S346), but since it does not serve Glennan's purpose of characterizing explanations in practice (outside classical mechanics) Salmon's account must be deemed irrelevant for Glennan's purposes.

So, The New Programs distance themselves from Hempel's program and focus increasingly on giving more practice oriented accounts of specific kinds of scientific explanations. The literature on scientific explanations is thus growing increasingly diverse and difficult to concisely review. In the following sections I discuss some of The New Programs and mention the results and discussions that are most relevant for the discussions in chapter 4 and 5 of this dissertation.

8.3.1 Causal explanations

It is quite simple to state the basic intuition behind the research program on causal explanations: To causally explain something is to state what caused it. This intuition seems to capture many of the explanations that we give in everyday life and even in scientific practice. Furthermore, there is hope that the link between causation and explanation can account for some of the other characteristics of explanations. In section 8.1 we saw that a major problem for the D-N model was that it could not account for the *asymmetry* of explanations. The flagpole example illustrated this. On the D-N model we can explain both the length of the shadow cast by a flagpole with reference to the height of the flagpole *and* we can explain the height of the flagpole with reference to the length of its shadow. This, to most philosophers, is an absurd consequence and must lead us to abandon the D-N model. The flagpole example also gives a hint as to *why* there is asymmetry in many explanations: We can explain the length of shadow with reference to the height of the flagpole, because the flagpole *causes* the presence and shape of the shadow and not vice versa. The cause explains the effect, not vice versa. So linking explanation to causation can provide a reason for why explanations display the characteristic asymmetry that Hempel's model could not account for. The basic idea behind the research program on causal explanations thus looks promising, but as is often the case in philosophy, the problems come when we try to explicate in detail what this idea entails.

Aristotle (Metaphysics, delta 2) identified four kinds of causes – material, formal efficient and final causes - all of which can be relevant in order to understand a given explanandum. Descartes argued that he found the "customary search for final causes to be totally useless in physics" (Descartes 1996, p. 39), illustrating the gradual turn in science to focus mainly on efficient causes. Hume famously challenged philosophers to point out something in our sensory experiences that shows that there is anything like causal connections in the world, not just conjunctions (or correlations), raising concern that talk of causation is just superficial metaphysics (Hume 1969, book I, sec II-III). Russell further argued that the word 'cause' "is so inextricably bound up with misleading associations" that a "complete extrusion" of the word from our vocabulary would be desirable (1913, p. 1). And still, scientists continue to talk about causes, also when they explain. So "if you can't beat them join them" says the proverb, and philosophers of explanation did indeed start to spell out their theories of explanation through a further analysis of causation. The literature on causality and causal explanations is too vast to be reviewed here, and I will only mention a few influential and illustrative examples.

Wesley Salmon first presented a causal account of scientific explanations based on a on statistically relevant factors (1971). A factor C is statistically relevant for another factor E iff $P(E) \neq P(E|C)$. This account was criticized among others by Kitcher (1989). Salmon therefore later changed his view to the causal mechanical view mentioned in section 8.1 which is based on a notion of *causal processes* transporting conserved quantities.

Peter Lipton (2004) based his influential account of *Inference to the Best Explanation* on a causal account of explanations relying on John Stuart Mill's account of causes as difference makers. Although Lipton recognizes that not all explanations, not even all scientific explanations, are causal, especially his argument for why inference to the best explanation as a method actually works relies on the assumption that explanations state causes.

In 2005, James Woodward was awarded the prestigious Lakatos Award for an outstanding contribution to philosophy of science for his *Making Things Happen*: *A theory of Causal Explanation* (Woodward 2003). The book has had tremendous influence, including a tremendous influence on the literature on mechanistic explanations (see next section). Woodward's account is based on an interventionist account of causation. According to Woodward causal connections are best understood as a connection among *variables* not events or processes. Very roughly two variables X and Y are causally connected if there is a possible intervention (a technical term which Woodward takes great care to define (2003, ch. 3)) on X that would change the value or probability distribution of Y. Woodward uses this account to argue that we causally explain by referring to such counterfactual dependencies among variables, and these explanations have the main virtue that they allow us to answer questions of the kind: "what if things had been different?"[74]. I return to Woodward's account in chapter 5.

The existing accounts of causal explanations are thus very diverse partly because they rely on rather different accounts of causation. While the accounts mentioned above are explicitly intended to be very general, and cover instances of causal explanations across many parts of science the studies of mechanistic explanations to which I now turn are often very careful to stress that they only intend to capture explanations in particular parts of science.

8.3.2 Mechanistic explanations in the biological sciences

Mechanistic explanations are a particular type of causal explanations that are abundant in the biological sciences. As Glennan (2002) argues, the causal explanations that biologists prefer often refer to relatively stable and localized arrangements of parts: what they themselves call *mechanisms*. This makes mechanistic explanations more constrained than causal explanations in general and different from some causal mechanical explanations in physics. There is thus a need to further characterize these biological mechanisms that feature so prominently in explanations. Again the literature on this type of explanations is too vast to be reviewed

[74] This illustrates again that what-questions can be explanation-seeking.

here, but I will mention a few important ideas and publications and return to the topic when I, in detail, discuss the differences in explanatory standards in molecular biology and polymer physics in chapter 5.

The most cited paper within the past three years from the prestigious journal *Philosophy of Science*[75] is *Thinking about Mechanisms* (Machamer, Darden & Craver 2000). The paper provides a definition of a mechanism as consisting of entities and activities that produce regular changes from set-up to termination conditions, and discusses a case study of a mechanistic explanation and the importance of diagrams in the representation of mechanisms. I have based my discussion of mechanisms in chapter 5 primarily on this definition.

Alternative definitions of mechanisms exist. Noticeably Bechtel and Abrahamsen (2005) stress that mechanisms not only produce regular changes, they perform a *function*. Researchers in the biological sciences are often interested in knowing how certain functions of biological systems are performed, and investigate this by identifying the parts of a mechanism and investigating how they interact to perform the functions of interest (see also (Bechtel, Richardson 2010)).

The research program on mechanistic explanations has been criticized for stretching the mechanism concept so far that it has become empty[76]. It seems that just about any explanation deals in one way or another with "entities and activities" - after all we do construct our sentences using nouns (referring to entities) and verbs (referring to activities) – so for there to be some content left in the mechanisms term we would probably want to further specify *what kinds* of entities and activities are used in mechanistic explanations, although this would probably mean that some of the generality of the framework is lost. In chapter 5 I compare explanations on polymers from molecular biology and polymer physics. One question that interests me is whether there is a difference in the entities and activities that are referred to in the explanations. I conclude that there are indeed differences. In the molecular biology books the entities referred to are generally *material objects* – i.e. objects with mass and extension – whereas explanations in polymer physics refer mainly to *variables*, some of which represent *physical quantities* – e.g. energy and entropy – which are not material objects with mass and extension. It is therefore fair to say that these explanations are not typical mechanistic explanations.

8.3.3 Explanations from other (natural) sciences

The literature on explanations in the biological sciences is by far the biggest of the literatures dealing with explanations in a specific part of science. Except, of course if one counts the entire Hempelian program as being a program investigating the nature of explanations in physics.

Explanations in chemistry have received some attention[77], for instance in a special issue of the *Annals of the New York Academy of Sciences* (vol. 988, no. 1). One of the major contributors to the understanding of explanations in chemistry comes from the highly Kuhn inspired work of Andrea Woody (2003, 2004a, 2004b).

[75]According to the journal's own web-site: http://www.jstor.org/action/showMostCitedArticles?journalCode=philscie (accessed May 2014).
[76] John Dupré has argued to this effect at various conferences, but is yet to put his criticism to print.
[77] This seems to be part of a broader trend among philosophers of science to finally start paying attention to chemistry (for an introduction see (Woody, Hendry & Needham 2012)).

Woody's work is interesting both because of its contents which details the importance and role of diagrams, equations and tables – including of course the periodic table of elements – in chemical explanations, and because of the empirical methodology that is employed[78]. Woody shows how diagrams like diagrammatic orbital schemes help the chemist relate quantum theory and spatial reasoning and are therefore indispensable to the chemist's efforts to construct explanations that satisfy the explanatory standards of chemistry (Woody 2004a). In a similar vein she argues that although the ideal gas law is not the most accurate equation of state for the gasses that a chemist can use, it serves an important role in the chemist's thinking about gasses as gasses, which is why it remains at the core of the curriculum in introductory chemistry (Woody 2013, Woody 2004a)[79].

Explanations in the remaining special sciences – e.g. the earth sciences - have received almost no attention so far, although Grantham (2004) does draw on a case from palaeontology in his discussion of the challenges to construct integrated explanations.

Explanations in physics as a special kind of explanations – as opposed to representative examples of scientific explanations - have not been receiving much attention either. One exception already mentioned is Weatherall's discussion of an explanation of why gravitational mass and inertial mass are identical[80]. The paper is discussed further in chapter 4, but it illustrates that there is important diversity in the way physicists explain that is not captured by the more general accounts of scientific explanations.

The puzzling world of quantum physics has long been of interest to philosophers. When discussed from the angle of explanation the question is not so much *how* quantum mechanics explains but rather *whether* quantum mechanics explains. Salmon emphasised that his causal mechanical account did not capture what was going on in quantum physics, but maintained that quantum mechanics was indeed explanatory (Salmon 1998, pp. 58-60). Others, like Cushing (1991) and Fine (1982), argued that quantum mechanics does not provide any understanding, and is therefore not an explanatory theory. Fine writes:

> ***[…] although quantum mechanics can be said to predict correctly the observed experimental statistics (if we simply bracket the measurement problem), it affords no understanding whatsoever as to how the statistics arise […]. It is rather the blackest of black-box theories; a marvellous predictor but an incompetent explainer. (Fine 1982, p. 740)***

[78] A detailed discussion of the methodology is presented in chapter 4.
[79] The status and role of laws in chemistry - often in comparison to their role and status in physics - is subject to a separate debate which to some extent relates to the discussions about what is special about explanation in chemistry (se e.g. (Christie, Christie 2000)).
[80] The inertial mass of a body is the "m" in Newton's second law, F=ma, and is a measure of how difficult it is to accelerate the body. The gravitational mass is the "m" in Newton's law of universal gravitation and is a measure of how much the body attracts and is attracted by other heavy bodies – in this respect it is a kind of 'gravitational charge'. It is not obvious why the 'gravitational charge' of a body is *always* numerically identical to the inertial mass while, for instance, the electrical charge of a body can be changed without changing its inertial mass.

There is no doubt that quantum mechanics raises many explanation-seeking questions that it cannot answer, and that the explanations that can be constructed using quantum mechanics are rather different from the causal mechanical explanations of classical mechanics[81], but the claim that quantum mechanics is not at all explanatory seems much too strong, if one adopts a characterization of explanation similar to mine.

8.4 Outlook: Teachers' explanations and the cognitive science of explanations

The preceding sections focused on how research on explanations within philosophy of science has developed over the past six decades. I argued that the quest to characterize the structural (even logical) characteristics of scientific explanations in general is no longer the primary focus of philosophy of science. Instead causal- and mechanistic explanations have recently been of primary interest to philosophers of science. This section takes a brief and highly selective look at the study of explanations within other research fields, especially science education and cognitive science. The outlook will show that we have much to learn from each other both when it comes to the conceptualization of explanations and the methods for studying them.

8.4.1 Explanations in the cognitive sciences

The pragmatic account of explanation defended here maintains that there is a cognitive side to explanations that is worth exploring. Explanations provide understanding, which in turn is a cognitive phenomenon. For instance, the connectedness that is claimed to be a characteristic of the knowledge of a person who has understood a given topic could be investigated through the methods of cognitive science. What kinds of connections exist? How do different kinds of responses to explanation-seeking questions alter these connections? How does the physiological setup of the brain constrain the kind and degree of understanding that is possible for the individual, and how do these change through the individual's life? These are some of the questions that the cognitive sciences could contribute interesting perspectives on. Explanation is a relatively new topic in the cognitive sciences and there is still much to explore[82].

Many philosophers reject the value of cognitive studies of explanations because the cognitive aspects of explanation are not epistemic aspects and are thus not interesting to the philosopher (this argument goes back to Hempel (1965)). As argued by Waskan and collaborators (2014), this argument rests on the false assumption that the only thing remotely related to explanation that a cognitive scientist can investigate is the feeling of understanding. Given this assumption it is easy to argue that the feeling of understanding may be interesting to the psychologist but is not of interest to the philosopher since it is subjective and at best an unreliable indicator of an epistemically interesting phenomenon. However, this argument breaks down once we accept that understanding is a cognitive phenomenon different from the feeling of understanding[83]. Instead, new and interesting questions arise, one of which is why the feeling of understanding exists at all? Gopnik provided an evolution-based (and rather speculative) answer to this question in her *Explanation as Orgasm* (2000). Here she suggests that because explanations improve our ability to act in a way that benefits us, they potentially

[81] As mentioned in the preface this was actually what got me interested in explanations in the first place.
[82] Related topics such as theories and prediction have been in focus for much longer (Wilson, Keil 2000).
[83] And those who do not accept this are challenged to counter the empirical argument presented by Waskan and collaborators showing that both scientists and lay people generally identify explanations as something that makes the explanandum intelligible.

improve our ability to survive and pass on our genes. So, just as the orgasm is an emotional reward for having done something that benefits the chance of passing on our genes so too, according to Gopnik, could the aha-experience be an emotional reward for having improved you and your offspring's chances of survival. Gopnik's arguments are theoretical, but the majority of interesting studies on explanations from the cognitive sciences are empirical[84].

Lombrozo (2006) reviewed recent cognitive studies on explanations. She noted that it is mainly causal explanations that are being studied, partly because these seem to be the most abundant in everyday life. Explanations are thus often conceptualized as a kind of causal reasoning that can be investigated further. Are specific kinds of causal reasoning preferred? When do subjects start to seek causal explanations? And are there other properties of causal explanations that will make explanations more convincing? In response to the latter question Lombrozo concludes that relatively simple causal explanations – i.e. explanations that refer to relatively few causes - and general explanations that can be used to explain multiple phenomena are often more likely to be accepted.

Another important topic in the cognitive science of explanation is the role of explanation in learning[85]. For instance, Lombrozo highlights studies on the effectiveness of self-explanation – i.e. explaining the content of a text or example to one self – as a method for learning compared to reading the material multiple times, receiving feedback or thinking out loud.

Although there is much potential in investigating questions related to explanations using the methods developed within the cognitive sciences, the current literature also indicates that further collaboration among philosophers and cognitive scientists is needed. Just as philosophers have something to learn from cognitive scientists, so too do cognitive scientists have something to learn from philosophers. Especially when it comes to characterizing the phenomenon they are investigating. If anything, studies in cognitive science demonstrate the pragmatic character of explanations by pointing to the importance of understanding (Waskan et al. 2014), and the importance of relating to the prior knowledge of the audience when explaining (Lombrozo 2006). And still, pragmatic philosophical accounts of explanations have very little influence in the cognitive sciences where, as mentioned, it is the research program on causal explanations that is most influential. The pragmatic accounts of explanations do not deny that causal explanations are abundant but, unlike some of the causal accounts of explanations, they also explicitly acknowledge the context dependence of explanations. It thus seems that the pragmatic accounts of explanations could serve as a promising framework for the cognitive studies of explanations[86].

8.4.2 Explanations in the classroom

Lombrozo (2006, p. 468) notes that a significant limitation to the cognitive studies of the role of explanations in learning is that they all take place in the controlled environment of the laboratory. How to generalize the

[84] Keil and Wilson edited a volume containing both empirical, theoretical and more programmatic papers on the cognitive science on explanations (Keil, Wilson 2000).
[85] Cf. note 55 this could also be seen as the role of explanation in gaining understanding.
[86] Antaki similarly criticizes the heavy reliance by cognitive scientists on causal accounts of explanation, and advocates the more pragmatic accounts (1988, introduction).

results to real world situations is therefore not entirely clear. In this section I focus on a different kind of study on explanations in education, namely the ones that focus on the explanations given to students either through their teaching material or by teachers. The aim in these studies is to characterize the explanations given and to some extent judge the quality of them. Explanations given to students are interesting in many respects. As argued by Treagust and Harrison (1999) the teacher is constrained in multiple ways when constructing explanations for students. In order to provide understanding, the explanations must relate to the skills and knowledge already possessed by the students in a relevant way. The teacher must thus consider this background knowledge of the students when constructing explanations, implying that it is probably not advisable to give the same explanation to a 9th grader and a 4th grader. Furthermore, the teachers' construction of explanations is constrained by the explanatory standards of the subject that (s)he is teaching. Learning science is partly learning to explain in the right way, referring to the right entities (gluons and grasshoppers not gods or goblins) doing the right activities (charges repulse and attract they do not 'like each other'). The students learn this largely through the explanations they are presented with in their teaching, and so if they are going to learn to explain in the right way then the examples they are given better live up to these standards.

One question that has interested educators is therefore whether explanations given in an educational setting are sufficiently 'scientific'. Given the discussion in section 8.1 on the difficulties related to characterizing scientific explanations in general one would expect that this is a rather difficult thing to measure, since it is rather unclear what is being measured. So the methodologies used in these studies are of interest as well.

Expecting to receive a number of D-N explanations, Edginton and Barufaldi (1995) asked science teachers and practicing scientists to scientifically explain two specific features of Newton's cradle (a row of swinging spheres commonly used to illustrate the conservation of energy and momentum). The study showed that there was indeed consensus on the ideal that explanations in mechanics should take the form of a derivation from general principles (whether or not the D-N model best captures this concept is discussed in chapter 5). In this respect the teacher's explanations were no less scientific than the practicing scientists. As could be expected given the discussion in section 8.1, the participants in the study complained that the task they were given was ill defined. Which audience should the explanation be aimed at? The answer to this question is crucial for determining the appropriate level of sophistication and detail to include in the explanation.

Chambliss and collaborators (2003) set out to evaluate explanations given by students rather than teachers. After explicit training in reading and constructing explanations they asked fourth graders to construct and write an explanation intended for third graders of the effects of a pollutant on a specific ecosystem. In preparing for this task the students had among other things read explanations of the effect of pollutants on ecosystems written by scientists. These explanations served as exemplars, both to the students and the investigators in the following parts of the study. The explanations written by the students were analyzed to see whether and to what extent the students used a "scientific causal model" in their explanations – i.e. a model similar to the one in the exemplar explanation. The study concluded that just under half of the students did use a scientific model in their explanation.

Other attempts to judge the quality of school explanations against a more or less explicit ideal of a scientific explanation include Peker and Wallace's (2011) discussion of the quality of explanations given by students in high school biology, and Unsworth's (2001) critical analysis of the language of explanations in textbooks aimed at junior high school students. Both of these studies seem to rely on a more tacit ideal when evaluating the quality of explanations.

Except for (Edgington, Barufaldi 1995), all of the studies cited above conclude that explanation as a genre of discourse has been overlooked in education and that it is necessary to pay more attention to it both in teaching practice and in research (see also (Unsworth 1997; Solomon 1995; Rowan 1988)). One might add that the situation is similar in philosophy of science. Philosophers have been so preoccupied with the logical and structural characteristics of explanations that they have neglected to study explanations as a genre of discourse (exceptions include (Faye 1999)).

In a more descriptive vein, Dagher and Cossman identified and categorized explanations given by teachers to 7th and 8th graders. I provide a detailed discussion of the methodology of this study in chapter 4. The study identifies ten basic categories of explanations that were given in class. These are again grouped into three very general categories: Practical-, theoretical- and tautological explanations. Practical and tautological explanations have no subcategories whereas theoretical explanations are subdivided into genuine and spurious explanations (clearly these names indicate that the authors found some theoretical explanations better than others, but they explicitly refrain from judging whether any of the explanations identified are scientific). The spurious and genuine explanations are then further subdivided into two and six subcategories respectively. The study illustrates the diversity in both the kinds and amount of verbal explanations that teachers provide in the classroom and it would be interesting to compare to similar studies focusing on different educational levels. Such studies could serve as a valuable empirical basis for further discussion of the importance of paying attention to how the explanations that teachers provide in the classroom develop over the course of an education.

9 Where are we now?

This introduction started out in section 1.1 by identifying the specific aims of the dissertation and sketching its argument structure. I stated two main aims of the dissertation. One was to *highlight specific expertises that, while not essential for engaging in traditional disciplinary problem solving, are nevertheless highly valuable when engaging in interdisciplinary problem solving.* The other main aim of the thesis was to *exemplify how a practice oriented approach to philosophy of science can be used to characterize differences in explanatory standards across scientific domains.*

In philosophy it is often both necessary and valuable to analyze the specific meaning of a question before trying to answer it. Therefore I set out to clarify and discuss the key concepts mentioned in the general aims of the thesis: *interdisciplinary problem solving, practice oriented philosophy of science, explanation* and *expertise*. I started out in sections 2-6 by discussing the topics related to interdisciplinary problem solving and interdisciplinary education. These sections clarified important concepts, including *integration*, and further connected the main aims of the thesis. In sections 7-8 I took a step back and looked more broadly at recent developments in philosophy of science in general and philosophy of explanations in particular. These sections positioned the thesis in a broader context, gave content to the idea of a practice oriented philosophy of science, and outlined the pragmatic view of explanations that underlies the thesis. What remains to be clarified is the concept *expertise* in general, and in particular how expertise can be divided into different categories enabling talk of expertises (in plural). I turn to this topic in the next chapter.

Getting to know what the aim actually is can sometimes take as much effort as it then takes to reach the aim. At least this is the case in this dissertation. Once we get to chapter 3, which fulfills the first main aim, we will already be more than half way on our journey. What remain are two chapters and a short section on “Future perspectives”. Chapter 4 outlines a general methodology for identifying explanations in textbooks, and chapter 5 applies a modified version of it to the case of polymer physics and molecular biology.

References

Abd-El-Khalick, F. (2013): Teaching With and About Nature of Science, and Science Teacher Knowledge Domains. *Science & Education, 22*(9), 2087-2107.

Achinstein, P. (1984): The Pragmatic Character of Explanation. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, 1984*, 275-292.

Achinstein, P. (1983): *The nature of explanation.* New York: Oxford University Press.

Akeroyd, F.M. (2008): Mechanistic Explanation versus Deductive-Nomological Explanation. *Foundations of Chemistry, 10*(1), 39-48.

Albritton, C.C. (1980): *The Abyss of Time.* New York: Dover Books.

Allchin, D. (2013): *Teaching the Nature of Science: Perspectives and Resources.* Saint Paul: SHiPS Education Press.

Allchin, D., Andersen, H.M. & Nielsen, K. (2014): Complementary approaches to Teaching Nature of Science: Integrating Student Inquiry, Historical Cases, and Contemporary Cases in Classroom Practice. *Science Education, 98*, 461-486.

Andersen, H. (2014): Epistemic dependence in contemporary science: Practices and Malpractises. In L. Soler, S. Zwart, M. Lynch & V. Israel-Jost (eds.), *Science after the Practice Turn in the Philosophy, History, and Social Studies of Science* (pp. 161-174). New York: Routledge.

Andersen, H. (2013): The Second Essential Tension: on Tradition and Innovation in Interdisciplinary Research. *Topoi,* 32(1), 3-8.

Anderson, L.W. & et. al. (2001): *A taxonomy for learning, teaching, and assessing: a revision of Bloom's taxonomy of educational objectives* (Complete ed.). New York: Longman.

Ankeny, R., Chang, H., Boumans, M. & Boon, M. (2011): Introduction: philosophy of science in practice. *European Journal for Philosophy of Science,* 1(3), 303-307.

Antaki, C. (1988): *Analysing everyday explanation: a casebook of methods.* London: Sage.

Aristotle (2004): *The Metaphysics.* London: Penguin.

AU (2011): *Den faglige udviklingsprocess*, Aarhus: Aarhus University.

Baker, A. (2005): Are there Genuine Mathematical Explanations of Physical Phenomena? *Mind, 114*(454), 223-238.

Barnett, S.M. & Ceci, S.J. (2002): When and where do we apply what we learn? A taxonomy for far transfer. *Psychological bulletin, 128*(4), 612-637.

Bechtel, W. & Richardson, R. (2010): *Discovering Complexity: Decomposition and localization as strategies in scientific research.* Cambridge MA: MIT Press.

Bechtel, W. (2006): *Discovering cell mechanisms.* Cambridge University Press.

Bechtel, W. (ed) (1986): *Integrating scientific disciplines.* Dordrecht: Martinus Nijhoff.

Bechtel, W. & Abrahamsen, A. (2005): Explanation: a mechanist alternative. *Studies in history and philosophy of biological and biomedical sciences,* 36(2), 421-441.

Berger, R. (1998): Understanding Science: Why Causes Are Not Enough. *Philosophy of Science, 65*(2), 306-332.

Boix Mansilla, V. & Duraisingh, E.D. (2007): Targeted Assessment of Students' Interdisciplinary Work: An Empirically Grounded Framework Proposed. *The Journal of Higher Education,* 78(2), 215-237.

Brigandt, I. (2010): Beyond Reduction and Pluralism: Toward an Epistemology of Explanatory Integration in Biology. *Erkenntnis, 73*(3), 295-311.

Brigandt, I. & Love, A.C. (2012): Conceptualizing Evolutionary Novelty: Moving Beyond Definitional Debates. *Journal of Experimental Zoology, 318B*, 417-427.

Bromberger, S. (1966): Why-Questions. In R. Colodny (ed.), *Mind and Cosmos.* Pittsburgh: University of Pittsburgh Press.

Brunnander, B. (2011): Philosophy and Default Descriptivism: The Functions Debate. *Metaphilosophy, 42*(4), 417-430.

Campbell, L.M. (2005): Overcoming Obstacles to Interdisciplinary Research. *Conservation Biology, 19*(2), 574-577.

Carp, R.M. (2001): Integrative Praxes: Learning from Multiple Knowledge Formations. *Issues in Integrative Studies, 19*, 71-121.

Cartwright, N. (2003): *The dappled world: A study of the boundaries of science.* Cambridge, UK: Cambridge University Press.

Chakravartty, A. (2014): Scientific Realism. In E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy,* (Spring 2014 Edition), URL= <http://plato.stanford.edu/archives/spr2014/entries/scientific-realism/>.

Chambliss, M.J., Christenson, L.A. & Parker, C. (2003): Fourth graders composing scientific explanations about the effects of pollutants: writing to understand. *Written Communication,* 20(4), 426-454.

Chang, H. (2011): The Philosophical Grammar of Scientific Practice. *International Studies in the Philosophy of Science, 25*(3), 205-221.

Christiansen, F.V. & Rump, C. (2008): Three Conceptions of Thermodynamics: Technical Matrices in Science and Engineering. *Research in Science Education,* 38(5), 545-564.

Christie, M. & Christie, J.R. (2000): Laws and Theories in Chemistry Do Not Obey the Rules. In N. Bhushan & S.M. Rosenfeld (eds.), *Of minds and molecules: new philosophical perspectives on chemistry* (ch 3). New York: Oxford University Press.

Chubin, D.E. (1985): Beyond invisible colleges: Inspirations and aspirations of post-1972 social studies of science. *Scientometrics, 7*(3-6), 221-254.

Clark, S.G. & Wallace, R.L. (2010): Learning Interdisciplinary Problem Solving and Leadership Skills: A Comparison of Four Designs. In S.G. Clark, A. Hohl, C. Picard & D. Newsome (eds.), *Large Scale*

Conservation: Integrating Science, Management, and Policy in the Common Interest (pp. 171-206). Yale School of Forestry & Environmental Studies.

Collins, H. (2011): Language and practice. *Social Studies of Science, 41*(2), 271-300.

Collins, H.M. & Evans, R. (2007): *Rethinking expertise.* Chicago: University of Chicago Press.

Collins, H.M. & Evans, R. (2002): The Third Wave of Science Studies. *Social Studies of Science, 32*(2), 235-296.

Collins, H., Evans, R. & Gorman, M. (2007): Trading zones and interactional expertise. *Studies in History and Philosophy of Science, 38*(4), 657-666.

Craver, C.F. (2007): *Explaining the brain: mechanisms and the mosaic unity of neuroscience.* Oxford: Clarendon Press.

Cross, C. (1991): Explanation and the Theory of Questions. *Erkenntnis, 34*, 237-260.

Cushing, J.T. (1991): Quantum Theory and Explanatory Discourse: Endgame for Understanding? *Philosophy of Science,* 58(3), 337-358.

Dagher, Z. & Cossman, G. (1992): Verbal Explanations Given by Science Teachers: Their Nature and Implications. *Journal of Research in Science Teaching, 29*(4), 361-374.

Darden, L. & Maull, N. (1977): Interfield Theories. *Philosophy of Science, 44*(1), 43-64.

de Regt, H.W. (2009): The Epistemic Value of Understanding. *Philosophy of Science, 76*(5), 585-597.

de Regt, H.W. & Dieks, D. (2005): A Contextual Approach to Scientific Understanding. *Synthese, 144*(1), 137-170.

Descartes, R. (1996): *Meditation on First Philosophy, With Selections from the Objections and Replies.* Cambridge: Cambridge University Press.

Douglas, H. (2010): Engagement for progress: applied philosophy of science in context. *Synthese, 177*(3), 317-335.

Duschl, R. (2008): Science Education in Three-Part Harmony: Balancing Conceptual, Epistemic, and Social Learning Goals. *Review of Research in Education, 32*(1), 268-291.

Edgington, J.R. & Barufaldi, J.P. (1995): How Research physicists and High-School Physics Teachers Deal with the Scientific explanation of a Physical Phenomenon. *Paper presented at the Annual Meeting of the National Association for Research in Science Teaching*.

Eichinger, D.C., Abell, S.K. & Dagher, Z.R. (1997): Developing a Graduate Level Science Education Course on the Nature of Science. *Science & Education, 6*(4), 417-429.

Eigenbrode, S. et al. (2007): Employing Philosophical Dialogue in Collaborative Science. *Bioscience, 57*(1), 55-65.

European Commision (2014): *Horizon 2020 in brief.* Luxembourg: Publications Office of the European Union.

European Commission (2008): *The European framework for lifelong learning (EQF)*. Luxembourg: Office for Official Publications of the European Communities.

Faye, J. (1999): Explanation Explained. *Synthese, 120*(1), 61-75.

Felt, U., Igelsböck, J., Schikowitz, A. & Völker, T. (2013): Growing into what? The (un-)disciplined socialisation of early stage researchers in transdisciplinary research. *Higher Education, 65*(4), 511-524.

Fine, A. (1982): Antinomies of Entanglement: The Puzzling Case of the Tangled Statistics. *The Journal of Philosophy,* 79(12), 733-747.

Fraenkel, J.R., Wallen, N.E. & Hyun, H.H. (2012): *How to design and evaluate research in education* (8.th ed.). New York: McGraw-Hill.

Friedman, M. (1974): Explanation and Scientific Understanding. *The Journal of Philosophy, 71*(1), 5-19.

Galison, P. (1997): *Image and logic: A material culture of microphysics.* Chicago: University of Chicago Press.

Gerson, E. (2009): Specialty boundaries, compound problems, and collaborative complexity. *Biological theory, 4*, 247-252.

Giere, R.N. (1985): Philosophy of Science Naturalized. *Philosophy of Science, 52*(3), 331-356.

Glennan, S. (2002): Rethinking Mechanistic Explanation. *Philosophy of Science, 69*(S3), S342-S353.

Godfrey-Smith, P. (2003): *Theory and reality: an introduction to the philosophy of science.* Chicago: University of Chicago Press.

Goodwin, W. (2003): Explanation in Organic Chemistry. *Annals of the New York Academy of Sciences, 988*(1), 141-153.

Gopnik, A. (2000): Explanation as Orgasm and the Drive for Causal Knowledge: The Function, Evolution, and Phenomenology of the Theory Formation System. In F.C. Keil & R.A. Wilson (eds.), *Explanation and Cognition* (pp. 299-325). Cambridge, Massachusetts: The MIT Press.

Grantham, T.A. (2004): Conceptualizing the (Dis)unity of Science. *Philosophy of Science,71*(2), 133-155.

Green, S. (2014): *Systems Biology and the Quest for General Principles.* PhD dissertation, Aarhus University.

Green, S., Fagan, M. & Jaeger, J. (2014): Explanatory Integration Challenges in Evolutionary Systems Biology. *Biological theory,* Online ahead of print. URL= http://link.springer.com/article/10.1007%2Fs13752-014-0185-8#page-1.

Grüne-Yanoff, T. (2014): Teaching philosophy of science to scientists: why, what and how. *European Journal for Philosophy of Science, 4*(1), 115-134.

Habermas, J. (1996): *Between facts and norms: contributions to a discourse theory of law and democracy.* Cambridge, Mass: MIT Press.

Hansson, B. (1999): Interdisciplinarity: For what Purpose. *Policy Sciences, 32*, 339-343.

Hempel, C.G. (1965): *Aspects of scientific explanation.* New York: The Free Press.

Hepburn, B. & Thorén, H. (in preparation): *Integrating interdisciplinary problem solving through process*.

Hume, D. (1969): *A Treatise of Human Nature.* London: Penguin.

Jakobsen, C.H., Hels, T. & McLaughlin, W.J. (2004): Barriers and facilitators to integration among scientists in transdisciplinary landscape analyses: a cross-country comparison. *Forest Policy and Economics, 6*(1), 15-31.

Jantsch, E. (1972): Inter- and transdisciplinary university: A systems approach to education and innovation. *Higher Education, 1*, 7-37.

Keil, F.C. & Wilson, R.A. (eds.) (2000): *Explanation and Cognition.* Cambridge, Massachusetts: The MIT Press.

Kitcher, P. (1989): Explanatory Unification and the Causal Structure of the World. In P. Kitcher & W.C. Salmon (eds.), *Scientific explanation* (pp. 410-505). Minneapolis: University of Minnesota Press.

Kitcher, P. & Salmon, W.C. (1989): *Scientific explanation.* Minneapolis: University of Minnesota Press.

Klein, J.T. (2000): A Conceptual Vocabulary of Interdisciplinary Science. In P. Weingart & N. Stehr (eds.), *Practicing Interdisciplinarity* (pp. 3-25). Toronto: University of Toronto Press.

Klein, J.T. (2010): A taxonomy of interdisciplinary research. In R. Frodeman, J.T. Klein & C. Mitcham (eds.), *The Oxford Handbook of Interdisciplinarity* (pp. 15-30). Oxford University Press.

Klein, J.T. (1990): *Interdisciplinarity: history, theory, and practice.* Detroit: Wayne State University Press.

Kosso, P. (2007): Scientific Understanding. *Foundations of Science, 12*(2), 173-188.

Kragh, H. (2004): *Naturerkendelse og videnskabsteori: de uorganiske videnskabers filosofi og historie.* Aarhus: Aarhus Universitetsforlag.

Kragh, H. (1981): The concept of the monopole. A historical and analytical case-study. *Studies in History and Philosophy of Science, 12*(2), 141-172.

Kuhn, T.S. (1996): *The structure of scientific revolutions,* (3rd ed.). Chicago: University of Chicago Press.

Kuhn, T.S. (1977): The essential tension: Tradition and Innovation in Scientific Research. In *The essential tension: Selected studies in scientific tradition and change.* Chicago: University of Chicago Press.

Kyburg, H. (1965): Comment. *Philosophy of Science, 32*, 147-151.

Lakatos, I. (1978): Falsification and the Methodology of Scientific Research Programmes. In J. Worral & G. Currie (eds.), *The Methodology of Scientific Research Programmes : Philosophical Papers : Vol.: 1.* Cambridge: Cambridge University Press.

Larson, E.L., Landers, T.F. & Begg, M.D. (2011): Building interdisciplinary research models: a didactic course to prepare interdisciplinary scholars and faculty. *Clinical and translational science, 4*(1), 38.

Lélé, S. & Norgaard, R.B. (2005): Practicing Interdisciplinarity. *Bioscience, 55*(11), 967-975.

Leonelli, S. (2009): Understanding in Biology: The impure Nature of Biological Knowledge. In H.W. de Regt, S. Leonelli & K. Eigner (eds.), *Scientific Understanding*, (pp. 189-210). Pittsburgh: University of Pittsburgh Press.

Lipton, P. (2009): Understanding without Explanation. In H.W. de Regt, S. Leonelli & K. Eigner (eds.), *Scientific Understanding,* (pp. 43-64). Pittsburgh: University of Pittsburgh Press.

Lipton, P. (2004): *Inference to the best explanation,* (2nd ed.). London: Routledge/Taylor and Francis Group.

Lombrozo, T. (2006): The structure and function of explanations. *Trends in cognitive sciences, 10*(10), 464-470.

Love, A.C. (2008): Explaining Evolutionary Innovations and Novelties: Criteria of Explanatory Adequacy and Epistemological Prerequisites. *Philosophy of Science, 75*(5), 874-886.

Machamer, P., Darden, L. & Craver, C.F. (2000): Thinking about Mechanisms. *Philosophy of Science, 67*(1), 1-25.

Mancosu, P. (2011): Explanation in Mathematics. In E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy* (Summer 2011 ed.). URL= <http://plato.stanford.edu/archives/sum2011/entries/mathematics-explanation/>.

Matthews, M. (1994): *Science Teaching: The Role of History and Philosophy of Science.* New York: Routledge.

Maull, N.L. (1977): Unifying science without reduction. *Studies in History and Philosophy of Science, 8*(2), 143-162.

McComas, W.F. (1998): *The nature of science in science education: rationales and strategies.* Dordrecht: Kluwer Academic Publishers.

Mestre, J.P. (2005): *Transfer of learning from a modern multidisciplinary perspective.* Greenwich: IAP - Information Age.

Mitchell, S.D. (2000): Dimensions of Scientific Law. *Philosophy of Science, 67*(2), 242-265.

Morris, S.A. & Van der Veer Martens, B. (2008): Mapping research specialties. *Annual Review of Information Science and Technology, 42*(1), 213-295.

Morrison, M. (2009): Understanding in Physics and Biology: From the Abstract to the Concrete. In H.W. de Regt, S. Leonelli & K. Eigner (eds.), *Scientific Understanding*, (pp. 123-146). Pittsburgh: University of Pittsburgh Press.

NAS (2004). *Facilitating interdisciplinary research.* Washington, D.C: The National Academies Press.

O'Rourke, M. & Crowley, S.J. (2013): Philosophical intervention and cross-disciplinary science: the story of the Toolbox Project. *Synthese, 190*(11), 1937-1954.

Peker, D. & Wallace, C. (2011): Characterizing High School Students' Written Explanations in Biology Laboratories. *Research in Science Education, 41*(2), 169-191.

Quine, W.V. (1969): Epistemology Naturalized. In *Ontological relativity and other essays* New York: Columbia Univ. Press.

Repko, A. (2008): *Interdisciplinary Research: Process and Theory.* USA: Sage.

Rheinberger, H. & Müller-Wille, S. (2010): Gene. In E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy* (Spring 2010 Edition). URL= <http://plato.stanford.edu/archives/spr2010/entries/gene/>

Rhoten, D. & Parker, A. (2004): Risks and Rewards of an Interdisciplinary Research Path. *Science, 306*(5704), 2046.

Roco, M.C. & Bainbridge, W.S. (eds.) (2003): *Converging Technologies for Improving Human Performance.* Dordrecht: Kluwer Academic Publishers.

Rohrer, H. (1993): Limits and Possibilities of Miniaturization. *Jpn. J. Appl. Phys., 32*, 1335-1341.

Rowan, K.E. (1988): A Contemporary Theory of Explanatory Writing. *Written Communication, 5*(1), 23-56.

Ruben, D. (2012), *Explaining explanation,* (Updated and expanded 2. ed.). Boulder: Paradigm Publishers.

Russell, B. (1913): On the Notion of Cause. *Proceedings of the Aristotelian Society, New Series, 13*, 1-26.

Salmon, W.C. (1998): *Causality and explanation.* New York: Oxford University Press.

Salmon, W.C. (1990): *Four decades of scientific explanation.* Minneapolis: University of Minnesota Press.

Salmon, W. (1971): *Statistical Explanation and Statistical Relevance.* Pittsburgh: University of Pittsburgh Press.

Salmon, W.C. & Kitcher, P. (1998): Van Fraassen on Explanation. In W.C. Salmon (ed.), *Causality and explanation*, (pp. 178-191). New York: Oxford University Press.

Scriven, M. (1962): Explanations, Predictions, and Laws. In H. Feigl & G. Maxwell (eds.), *Scientific Explanation, Space, and Time*, (pp. 170-230). Minneapolis: University of Minnesota Press.

Shapere, D. (1984): *Reason and the search for knowledge: investigations in the philosophy of science.* Dordrecht: D. Reidel.

Solomon, J. (1995): Higher level understanding of the nature of science. *School Science Review, 76*(276), 15-22.

Spelt, E.J.H., Biemans, H.J.A., Tobi, H., Luning, P.A. & Mulder, M. (2009): Teaching and Learning in Interdisciplinary Higher Education: A Systematic Review. *Educational Psychology Review, 21*(4), 365-378.

Strevens, M. (2013): No understanding without explanation. *Studies in History and Philosophy of Science, 44*(3), 510-516.

Thorén, H. & Persson, J. (2013): The Philosophy of Interdisciplinarity: Sustainability Science and Problem-Feeding. *Journal for General Philosophy of Science, 44*(2), 337-355.

Treagust, D. & Harrison, A. (1999): The genesis of Effective Scientific Explanations for the Classroom. In J.L. Loughran (ed.), *Researching Teaching: Methodologies and Practices for Understanding Pedagogy* (pp. 28-43). London: Routledge.

Trout, J.D. (2002): Scientific Explanation and the Sense of Understanding. *Philosophy of Science, 69*(2), 212-233.

Unsworth, L. (2001): Evaluating the language of different types of explanations in junior high school science texts. *International Journal of Science Education,23*(6), 585-609.

Unsworth, L. (1997): Explaining Explanations: Enhancing Science Learning and Literacy Development. *Australian Science Teachers Journal, 43*(1), 34-50.

van Camp, W. (2014): Explaining understanding (or understanding explanation). *European Journal for Philosophy of Science, 4*(1), 95-114.

Van Fraassen, B.C. (1980): *The scientific image.* Oxford:Clarendon.

Wagenknecht, S. (2014): *Collaboration in scientific practice: A social epistemology of research groups.* Ph.D thesis, Aarhus University.

Waskan, J., Harmon, I., Horne, Z., Spino, J. & Clevenger, J. (2014). Explanatory anti-psychologism overturned by lay and scientific case classifications. *Synthese, 191*(5), 1013-1035.

Weatherall, J.O. (2011): On (Some) Explanations in Physics. *Philosophy of Science, 78*(3), 421-447.

Weingart, P. (2000): Interdisciplinarity: The Paradoxical Discourse. In P. Weingart & N. Stehr (eds.), *Practicing Interdisciplinarity*, (pp. 25-43). Toronto: University of Toronto Press.

Wilkenfeld, D.A. (2013): Understanding as representation manipulability. *Synthese, 190*(6), 997-1016.

Wilson, R.A. & Keil, F.C. (2000): The Shadows and Shallows of Explanation. In F.C. Keil & R.A. Wilson (eds.) *Explanation and Cognition,* (pp. 87-115). Cambridge, Massachusetts: The MIT Press.

Woodward, J. (2011): Scientific Explanation. In E.N. Zalta (ed.) *Stanford Encyclopedia of Philosophy*, (Winter 2011 ed.). URL= <http://plato.stanford.edu/archives/win2011/entries/scientific-explanation/>

Woodward, J. (2003): *Making things happen: A theory of causal explanation.* New York: Oxford University Press.

Woody, A. (2013): How is the Ideal Gas Law Explanatory? *Science & Education, 22*(7), 1563-1580.

Woody, A. (2004a): Telltale Signs: What Common Explanatory Strategies in Chemistry Reveal about Explanation Itself. *Foundations of Chemistry, 6*(1), 13-43.

Woody, A. (2004b): More Telltale Signs: What Attention to Representation Reveals about Scientific Explanation. *Philosophy of Science, 71*(5, Proceedings of the 2002 Biennial Meeting of The Philosophy of Science Association Part II: Symposia Papers Edited by Sandra D. Mitchell), 780-793.

Woody, A. (2003): On Explanatory Practice and Disciplinary Identity. *Annals of the New York Academy of Sciences, 988*(1), 22-29.

Woody, A., Hendry, R.F. & Needham, P. (2012): *Philosophy of chemistry.* Amsterdam: Elsevier/North Holland.

Ylikoski, P. (2009): The Illusion of Depth of Understanding in Science. In H.W. de Regt, S. Leonelli & K. Eigner (eds.), *Scientific Understanding* (pp. 100-123). Pittsburgh, Pa.: University of Pittsburgh Press.

Recent RePoSS issues

#14 (2011-2) **Helge Kragh**: "*Quantenspringerei*: Schrödinger vs. Bohr" (Talk)

#15 (2011-7) **Helge Kragh**: "Conventions and the Order of Nature: Some Historical Perspectives" (Talk)

#16 (2011-8) **Kristian Peder Moesgaard**: "Lunar Eclipse Astronomy" (Article)

#17 (2011-12) **Helge Kragh**: "The Periodic Table in a National Context: Denmark, 1880-1923" (Book chapter)

#18 (2012-1) **Henrik Kragh Sørensen**: "Making philosophy of science relevant for science students" (Talk)

#19 (2012-8) **Helge Kragh**: "Newtonianism in the Scandinavian Countries, 1690–1790" (Book chapter)

#20 (2012-9) **Helge Kragh**: "Har Naturvidenskaben en Fremtid?" (Book chapter)

#21 (2013-6) **Helge Kragh**: "Empty space or ethereal plenum? Early ideas from Aristotle to Einstein" (Article)

#22 (2014-7) **Helge Kragh**: "Jørgen From (1605–1651): Et Billede fra Dansk Astronomi- og Lærdomshistorie" (Article)

#23 (2014-8) **Nils Randlev Hundebøl**: "Big ambitions and supercomputing: A case of industry engagement in climate modeling" (Preprint)

#24 (2014-8) **Julie Louise Hørberg Kristensen**: "Jagten på en god oplevelse i zoologisk have" (Masters thesis)

#25 (2014-9) **Allan Lyngs Kjærgaard**: "Offentlig støtte til forskning og udvikling i industrien fra 1940'erne til 1990'erne" (Masters thesis)

#26 (2014-10) **Julie Nørgaard Hostrup**: "Når discipliner mødes: Et casestudie af faglig integration i nanoforskning" (Masters thesis)

#27 (2014-10) **Kristian Danielsen, Morten Misfeldt and Henrik Kragh Sørensen**: "Picks Sætning: Inspiration til undervisning om eksperimenterende udforskning i matematik med et videnskabsteoretisk sigte" (Preprint)

#28 (2014-11) **Mads Goddiksen**: "Philosophical Perspectives on Interdisciplinary Science Education: An Introduction" (Book chapter)

This list is up-to-date as of November 17, 2014. The full list of RePoSS issues can be found at `www.css.au.dk/reposs`.

RePoSS (Research Publications on Science Studies) is a series of electronic publications originating in research done at the Centre for Science Studies, University of Aarhus.

The publications are available from the Centre homepage (www.css.au.dk/reposs).

ISSN 1903-413X

Centre for Science Studies
University of Aarhus
Ny Munkegade 120, building 1520
DK-8000 Aarhus C
DENMARK
www.css.au.dk

Faculty of Science
Department of Science Studies
AARHUS UNIVERSITET