

Teaching Philosophy of Science to Scientists: Why, What and How

Till Grüne-Yanoff

Royal Institute of Technology (KTH), Stockholm
Finnish Centre of Excellence in the Philosophy of the Social Sciences, Helsinki (TINT)

Forthcoming in *European Journal for Philosophy of Science*

ABSTRACT This paper provides arguments to philosophers, scientists, administrators and students for why science students should be instructed in a mandatory, custom-designed, interdisciplinary course in the philosophy of science. The argument begins by diagnosing that most science students are taught only *conventional* methodology: a fixed set of methods whose justification is rarely addressed. It proceeds by identifying seven benefits that scientists incur from going beyond these conventions and from acquiring abilities to analyse and evaluate justifications of scientific methods. It concludes that teaching science students these skills makes them better scientists. Based on this argument, the paper then analyses the standard philosophy of science curriculum, and in particular its adequacy for teaching science students. It is argued that the standard curriculum on the one hand lacks important analytic tools relevant for going beyond conventional methodology – especially with respect to non-epistemic normative aspects of scientific practice – while on the other hand contains many topics and tools that are not relevant for the instruction of science students. Consequently, the optimal way of training science students in the analysis and evaluation of scientific methods requires a revision of the standard curriculum. Finally, the paper addresses five common characteristics of students taking such a course, which often clash with typical teaching approaches in philosophy. Strategies how best to deal with these constraints are offered for each of these characteristics.

1. Introduction

Most philosophers of science would agree that their work is about science.¹ But to whom do we mostly teach it? To philosophers, not to scientists.² This is an unfortunate state of affairs. Science is a systematic and self-reflexive enterprise, thus it is – or at least should be – interested in relevant insights about itself. Philosophers of science who are serious about their research should be eager to share their results with budding scientists. And of course most of them are – but a number of obstacles in the current academic landscape make this kind of sharing difficult. This paper seeks to help philosophers of science to overcome these obstacles, by providing arguments for such a course, and by deriving from these arguments proposals for a suitable curriculum and appropriate teaching strategies.

First, there are disciplinary boundaries to overcome. We have first and foremost obligations towards our own students; our heads of department or our deans might not consider teaching students from other departments equally important, might not count it equally towards our teaching duties or might not compensate the department adequately.

¹ Philosophers of science tend to treat science as a broader category than standard English usage – more akin to the scope of the Latin *scientiae* or the German *Wissenschaften*, including the social and engineering sciences, and (to a lesser degree) the humanities. This will also be my usage in this paper.

² I know of no reliable data about course numbers for any country. But anecdotal evidence collected from friends and colleagues in the philosophy profession shows that most of them are regularly involved in teaching PoS to philosophers, while only few teach it to science students.

Being able to teach philosophy of science (PoS) to scientists thus requires convincing university administrators.

Second, there is the scepticism of many scientists to overcome. The stereotype of philosophers as scientifically illiterate but prescriptively omnipotent seems rather common amongst scientists, perhaps not without reason. They are often concerned that philosophers fail to appreciate the subtleties of scientific practices, are proponents of some general and sterile scepticism, seek to impose constraints detrimental to science (Medawar 1963), or simply produce work irrelevant to science (as in Feynman's infamous if apocryphal ornithology comparison)³. Consequently, they often are not thrilled if their students should spend their precious time on a PoS course. Being able to teach PoS to scientists thus requires convincing scientists.

Finally, there is student inertia to overcome. Science students often do not know what to expect of PoS course, believe that philosophy cannot help them in becoming successful scientists, or consider philosophy to be 'waffle'. Consequently, they often choose not to attend voluntary courses, and are irked when courses are mandatory. Being able to teach PoS to scientists thus requires convincing students.

Convincing administration, scientists and students requires good arguments. In the second section of this paper I therefore spell out answers to *why* PoS is important for training scientist.

Yet scientists' scepticism and students' disinterest is not entirely to blame. At least sometimes, their attitudes are justified, because – as I will argue – not the entire standard PoS curriculum is relevant for scientists. Consequently, we need not only convince others that teaching PoS to scientists is a worthy endeavour; we also have to revise our own teaching program. In the third section, I analyse the standard PoS curriculum, propounded in PoS textbooks of the last 15 years. I then spell out *what* parts of this standard curriculum are not relevant for training scientists, and what relevant parts are missing.

Finally, despite all our pedagogical idiosyncrasies, most philosophers share certain modes of teaching their students. These include a tendency towards generalising conclusions, using historical categorisations, engaging in conceptual analysis, use of thought experiments, and focus on exact reading of dense texts. Many of the items on this list we cherish as valuable tools for philosophical instruction, and with good reason. Yet to science students, many of our teaching modes will be alien to them – and will alienate them further from the topic that we're trying so hard to bring them closer to! In order to avoid this, we should not only revise what we teach, but also *how* we teach it. So in the fourth part, I propose various strategies how to make the right PoS curriculum more accessible to science students.

³ Cf. also: "Most scientists receive no tuition in scientific method, but those who have been instructed perform no better as scientists than those who have not. Of what other branch of learning can it be said that it gives its proficient no advantage; that it need not be taught or, if taught, need not be learned?" (Medawar 1969).

The paper's intended audiences are both philosophers of science and practicing scientists. For each group, a small and fragmented literature has discussed some aspects of the arguments presented here. In philosophy of science, authors have argued that the extant science textbooks give inadequate analyses of scientific methods (Martin 1976, Blachowicz 2009), implying the need for a proper philosophy of science analysis. Others have criticized science instruction for its stilted and often confused structure and have advocated using the resources of the history and philosophy of science to improve them (Ennis 1979, Hodson 1991, Matthews 1994). Yet others have stressed the centrality of rationality or critical thinking as a fundamental educational ideal, and have from that derived the importance of philosophy of science for science education (Siegel 1989). Amongst scientists, authors have occasionally called for a more thoroughgoing philosophy of science education for their students, arguing that a better philosophy education greatly facilitates a better science education (Grayson 2006) and suggesting to put "the 'Ph' back into 'PhD'" (Prather et al. 2009). All these authors agree that philosophy of science should play a more prominent role in the education of scientists. This paper systematically develops these arguments and strategies further in the light of extensive teaching experience of such a course.

2. Why?

In order to obtain a Masters degree at a European university, a student should have acquired "highly specialised knowledge, some of which is at the forefront of knowledge in [her] field" and show "critical awareness of knowledge issues in a field and at the interface between different fields" (European Commission 2008). Furthermore, the student should master "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" (ibid.).⁴

In my experience, in the science disciplines these requirements are commonly achieved by educating students in three broad domains. First, by teaching them the main theories and some exemplary models and experiments of a field. Second, by giving them instructions in some basic relevant methods of these fields, e.g. how to collect data, how to build a model or how to design an experiment. Third, by tutoring them in some basic skills in applying these methods, e.g. running experiments, programming simulations, or performing measurements.

A common observation about these educational programs across scientific disciplines is that they teach a lot of *conventional methodology*. That is, students are instructed about

⁴ For Bachelor degrees, the respective requirements are to have "advanced knowledge of a field ... involving a critical understanding of theories and principles" and the ability "to solve complex and unpredictable problems in [her] specialised field" (ibid.). These definitions by the European Qualifications Framework (EQF) aim to relate different countries' national qualifications systems to a common European reference framework, and thus are a good predictor of the near future standards at European universities. Similar standards exist outside of the EU as well (e.g. Ontario Council of Academic Vice Presidents 2008).

basic methods and tutored in their use, but they are not given an explanation or a justification for this choice of methods, or why these methods are designed the way they are.

To support this claim, let me discuss a few cases. In economics, for example, students are taught to value simplicity in theoretical model building. “Write down the simplest possible model you can think of, and see if it still exhibits some interesting behaviour. If it does, then make it even simpler” (Varian 1997, 4-5) suggests Hal Varian, the author of the most popular microeconomics textbooks.⁵ This preference for simplicity is further expressed in economists’ preference for analytically solvable models over computational simulation models (Lehtinen and Kuorikoski 2007). Consequently, economics students are rarely taught about or trained in computational methods. This focus on simplicity and analytic solvability contrasts economic education to other disciplines in the social sciences, e.g. analytical sociologists. So the question arises, why does such a difference exist? But that is a question usually not discussed with economics students.

Another example concerns the widespread use of 0.05 as a sufficient significance level for statistical tests. Of course, students are taught at least one of the meanings of statistical significance, which imply that significance is a continuous concept, and – at least under the Fisherian interpretation – the smaller is the better. Yet informally, students are taught that rejecting the null at a p-value lower than 0.05 is sufficient and constitutes a “statistically significant test”. Students are not normally taught why this number is sufficient, but rather are told that this is a convention (cf. Stigler 2008).

A third example concerns the calibration of novel measurement instruments. Calibration requires the selection of a standard, i.e. a measurement device with known or assigned correctness. When making this selection, the 4:1 accuracy ration is often cited as a rule of thumb: the standard should be four times more accurate than the instrument being checked. The reason for ensuring a 4:1 ratio is to minimize the effect of the accuracy of the standard on the overall calibration accuracy, given that standards used are likely themselves calibrated with higher-level standards (Cable 2005,3). Thus, a conventional rule of thumb replaces a more precise but also more involved acceptability judgment, based on the specific uncertainty of the instrument calibration.

A fourth example concerns the way forecasts are derived from a data set. In many disciplines, including economics, students are taught that they must always derive the forecasting model from some background theory, and the practice of identifying patterns from the data set alone is denigrated as “data mining”. Yet there is no obvious epistemic argument that would prohibit all kinds of such “data mining” practices altogether, and indeed, the attitudes towards such practices differ between disciplines (Hoover & Perez 2000).⁶ Yet again, students are taught disciplinary conventions rather than being exposed

⁵ Note, however, that he makes this recommendation not in any of his textbooks, but in some separately published paper.

⁶ Roughly, their argument is that analyzing the data might be a fruitful way of developing theory, even if no theoretical model of the relationship between endogenous and exogenous variables has been explicitly

to the reasons for choosing different kind of forecasting methods.

A fifth example concerns how potentially compounding factors are identified in experiments. The behavioural sciences are deeply divided in their opinion on how the lack of material incentives might compound an experiment. Economists virtually never accept experiments where choice alternatives are not properly incentivised, while psychologists generally consider the lack of material incentives to have little experimental effect (Hertwig and Ortmann 2001, 391). They teach their respective convictions to their students, thus entrenching already deep-running conventional differences.

These examples are only meant to illustrate my claim that students are instructed about basic methods and tutored in their use, but are not given an explanation or a justification for method choices. I have chosen these examples for their clarity and generality, but I do not mean to suggest that they are isolated occurrences. To the contrary, I presume that most readers can remember many more instances of such practical advice, given by the lecturer momentarily veering from her script, the seminar leader criticising a student's thesis, the tutor suggesting how to improve your model, or the lab assistant pointing out what you missed in the experimental set-up. I certainly remember many such instances from my own education in economics. They are usually not printed in textbooks or written down in the script, and there usually is no time or desire to discuss the "why?".⁷ Instead the students learn to accept them as little helpful hints in their way to become trained and acculturated in the discipline of their choice.

Philosophers of science are well qualified to address this "why?". This is at least for two reasons: their specific competences and their specific perspectives. Concerning the former reason, philosophers of science are trained in describing and analysing scientific procedures of evaluating and credentialing scientific knowledge, and that's what they pursue in their daily work. A typical procedure in this work starts by analysing a method or practice, by describing it in the context in which it is applied, and by then identifying the features that such a kind of method exhibits beyond the specific context of application. It then proceeds to analyse the scientific study or project in which this method was employed, its scientific objective, its theoretical and conceptual base, and the available and prospective evidence. An evaluation is pursued by relating the methods to the identified objectives, given theoretical and evidential background, and by reconstructing the arguments that purportedly justify this relation.

formulated. Such an approach would still presuppose theory in various ways, but not in the way mainstream economists claim is required.

⁷ For an investigation of 70 textbooks from the main scientific disciplines, see Blachowicz (2009). He concludes (i) that textbooks tend to present a simple empiricist view of science that inaccurately downplays theoretical and pragmatic considerations, (ii) that they overstate the demarcation between scientific and non-scientific inquiry, providing a stereotyped view of the latter, and (iii) that they tend to downplay the prevalence of controversy of science, eliciting the inaccurate picture of methodological harmony.

Of course, many practicing scientists also perform this kind of work, and they also have the competence for it. So what makes philosophers especially qualified to address this “why?” is not their competence alone, but also their particular perspective onto science as a whole, and their position outside of specific (sub-)disciplines, in contrast to virtually all practicing scientists. Philosophers of science study science, and often study very specific parts of specific disciplines, but they do not normally position themselves within these disciplines. Rather, they assume a broader perspective, in which they compare the concepts and practices of different disciplines, or in which they compare the concepts and practices of a specific discipline with a more abstract model of science or scientific practice. Such a perspective allows them, for example, to compare similar concepts (e.g. “operationalization”) in different disciplines and work out its discipline-independent aspects. Or it allows them to consider possible alternatives to a given method choice, perhaps by drawing on projects in other areas or disciplines, and evaluating the actually chosen method in the light of these possible alternatives.

It is through these comparative and abstracting perspectives that philosophers contrast with practicing scientists. Through them, they (i) make explicit the justifications for certain methods, (ii) render comparable the strengths of justifications of different methods in different disciplines (at least for similar applications) and (iii) uncover possible gaps or inconsistencies in the arguments that are supposed to justify these methods. This makes them ideally qualified in filling the lacunae that were left by the conventional methodology taught in the sciences – and it makes them more qualified to do so than practicing scientists.

The difference between scientist and philosophers being determined more by perspective than competence, Feynman’s alleged claim that scientists are the birds that need not take any interest in the ornithologists’ studies is doubly wrong. Not only should they be interested for their own good, but neither is there such a clear division between scientists and philosophers. Scientists often consider the justifications of their method choices themselves, they often see the need to do so and they often have the required competence. So they are already bird-ornithologist hybrids. The philosophers, on the other hand, need to be sufficiently familiar with scientific practice to pursue their work – they also must be hybrids. Between these two hybrids, competences, positions and perspectives are weighted differently, but no clear-cut binary distinction emerges. Instead, the different weighings are justified by a division of labour beneficial to both. Consequently teaching students how to explain and justify their method choices must be a collaborative effort. Achieving this aim requires a lot from both sides and cannot be achieved by philosophers alone.

So far, I have only argued that typical science education does not address the explanation or justification of conventional methodology, and that philosophers of science are ideally qualified to do precisely that. Sceptics might argue however that there is no need to do so – that it is irrelevant for students to reflect on these issues, and that it distracts from the facts, methods and skills they really are supposed to learn.

I argue against such an irrelevance claim. The procedures that philosophers of science use in order to identify, compare and evaluate justifications of methods are highly relevant for scientists and therefore should be a mandatory part of their training. First of all, they provide scientists with a *better understanding*: by learning what justifies the methods they use, scientists increase their understanding of these methods, and their scope, purpose and relation to other methods. Second, they provide scientists with a greater capacity of *critical reflection*: by improving their understanding of the scope and purpose of their methods, and by comparing them to others, scientists improve their ability to discern their respective advantages and disadvantages.

Better understanding and greater capacity of critical reflection clearly are intellectually rewarding and desirable. Furthermore, they are considered to belong centrally to *any* scientific education. The European Quality Framework, for example, requires “a critical understanding of theories and principles” for Bachelor degrees and “critical awareness of knowledge issues in a field” for Master degrees (European Commission 2008). From this requirement and my arguments above, it follows that science students should be instructed in PoS by philosophers.

Nevertheless, some critics seem dissatisfied with these “broader” requirements of critical understanding or awareness. They argue that such capacities, although of merit on their own, are not directly relevant for a proper functioning as a scientist, and therefore need not be part of a scientist’s training. Against such a claim, I show that philosophy of science not only produces intellectually rewarding capacities, but also directly contributes to the proper functioning of scientists.

Students who obtain a university degree in science today are expected to apply their knowledge *flexibly* and in new and unexpected situations. Bachelor degree holders are required “to solve complex and unpredictable problems in a specialised field” (European Commission 2008) and take “responsibility for decision-making in unpredictable work or study contexts” (ibid.). Master degree holders “manage and transform work or study contexts that are complex, unpredictable and require new strategic approaches” (ibid.). These requirements stress the need for applying knowledge and skills outside of the contexts in which they were acquired. Scientists thus are confronted with the question whether the specific methods they acquired as part of these skills can be applied to the novel contexts, or whether adjustments or even alternative methods are needed. A conventional methodology cannot answer this question. Instead, it demands that scientists have thought about the justifications for the methods they are familiar with, are therefore aware of the scope, purpose and relation to other methods, and therefore can judge competently whether these methods are applicable in these novel and unexpected contexts, or whether other methods would do better.

A related requirement is that graduates “demonstrate innovation”, “develop new knowledge”, implement “new strategic approaches” and produce “original thinking and/or research” (European Commission 2008). Yet the methods justified by the conventional methodology support such *innovative, bold and risk-taking* explorations only to a limited extent. The conventions were developed and credentialed for the

existing theories, not for novel and anomalous problem settings. To pursue bold and innovate research thus often requires freeing oneself from the confines of these conventions. The best way to equip students for such bold and innovative moves thus is to teach them about the extant justifications of their methods and make them reflect about their potentials and limitations.

Another requirement on graduates is that they are capable of working *interdisciplinarily*. Master degree holders are supposed to have “critical awareness of knowledge issues ... at the interface between different fields” and “to integrate knowledge from different fields” (ibid.). Many grant-giving institutions strongly favour proposals with strong interdisciplinary aspects. Yet, as becomes clear from many of the cases I discussed above, conventional methodologies between disciplines often differ and hence present a serious obstacle for interdisciplinary collaboration and exchange. In order to overcome this obstacle, scientists seeking interdisciplinary collaboration or exchange must be able to identify the conventional aspects of their respective methods’ justifications, must be able to compare their respective methods, and ultimately must agree on new methods acceptable to all of them and their epistemic needs. This requires that each of them is able to analyse, compare and evaluate their own and their collaborative partners’ methods.

Another requirement on scientists is *science communication*. The European FP7 framework, for example, states that “the beneficiaries shall, throughout the duration of the project, take appropriate measures to engage with the public and the media about the project aims and results” (FP7 Grant agreement II.2). A lot of effort is put into developing and funding elaborative dissemination strategies. Yet scientists themselves often have a difficult time communicating not just their results to a general public, but also the epistemic uncertainty associated with these results.

This issue of epistemic uncertainty is directly connected to the dominance of a conventional methodology in the sciences. Scientists are usually very good at specifying the scientific uncertainties of their results – e.g. by identifying the variance of a distribution of a data set. These specifications however rely on certain methods themselves – methods that are conventionally accepted within a discipline. Within this discipline, it is implicitly understood how confident one should be about the accuracy of these methods. Yet these “higher-order beliefs” about the confidence in the underlying theory, in the methods employed or in the researcher who carried out the study are seldom quantified, and therefore difficult to communicate (Hillerbrand and Ghil 2008, 2136). The repeated public confusion about the IPCC’s climate forecasts is a case in point. The IPCC’s *summary for policy makers* gives “best estimates” for a number of scenarios, and specifies “likely ranges” for each of these estimated scenarios (IPCC 2007). The ambiguity of the uncertainty qualifiers “best” and “likely” has contributed to (possibly overtly strong) interpretations of certain data. For example, in the so-called “Antarctica Cooling Controversy”, various advocacy groups took the finding that “between 1966 and 2000, 58% of Antarctica was cooling” as strong evidence against the IPCC’s model projections (Doran 2006). Yet in the scientific community, no one interpreted the evidence this way. The divergence in this interpretation thus can be plausibly attributed to the scientists’ difficulty in communicating the epistemic

uncertainty of their forecasting methods. Teaching science students PoS can help here: by improving scientists' understanding of the justifications of their methods, by increasing their potential to critically reflect on the scope and purpose of their methods, and by comparing them to others, we also improve their science communication abilities.⁸

Finally, scientists are supposed to consider non-epistemic values and norms in their choices how to conduct research. These values not only determine what bad must be avoided in research, but also identify the moral or social relevance of research decisions, and help scientists to align research questions with problems of public importance (Kitcher 2003). Yet neither the content of such considerations, nor the way how to incorporate them in one's professional choices, are commonly taught in scientific training. Philosophy of science here has two important roles to play: on the one hand, it draws on extensive research in normative philosophy that helps structure and critically reflect on moral, societal, prudential and instrumental values. On the other hand, by analysing and evaluating the justifications of scientific methods, it opens up a normative dimension in science that allows for an incorporation of non-epistemic values and norms in choosing how to conduct research. By instructing students how to reflect on methods, we thus also increase their competences in incorporating non-epistemic normative considerations.

In this section, I have argued that there typically is a lacuna in the education of science Bachelor and Master students: students learn a conventional methodology, without being taught how to critically analyse, compare and justify their and others' methods. I further showed that philosophers of science are ideally qualified to fill this gap, both through their competences, and through their particular position outside of any discipline. Finally, I argued that this lacuna is indeed a deficit, and that philosophical instruction is beneficial for scientists: it increases their ability to apply these methods more flexibly, equips them for bolder and more innovative research projects, increases their capacity to work interdisciplinarily, improves their science communication skills and increases their competences to make responsible research choices. In other words, teaching philosophy of science to science students makes them better scientists.⁹

3. What?

Philosophy contributes to scientific work in multiple ways. Many individual disciplines collaborate with philosophers on foundational issues. Philosophers of physics, for example, discuss questions of quantum physics, philosophers of biology investigate conceptual issues in evolution and philosophers of economics address problems of

⁸ And, as a side effect, improving their communication skills will also improve students' ability to write scientific papers!

⁹ Some European countries have acknowledged this role and have introduced legislation requiring philosophy education for scientists. Examples are the Swedish "högskoleförordningen", that mandates courses in "Theory of Science" (*Vetenskapsteori*) for most teacher examinations. The Finnish national graduate program requires philosophy of science courses for PhD students. The Danish re-instated (by decree of parliament, no less) the *examen philosophicum* which requires a course of philosophy of science for all undergraduate students.

rationality in decisions. Clearly, philosophers can contribute in these ways when teaching students from these respective disciplines.

Yet this is not my focus here. What I stress in this argument is not the relevance of philosophy for foundations of special disciplines, but rather the relevance of philosophy for science instruction generally. This has a practical and a conceptual motivation. Practically, courses of the kind I am arguing for here often have a highly interdisciplinary audience. Take my own case (at the Royal Institute of Technology (KTH) in Stockholm), where I teach a PoS course for science students from disciplines as far apart as nuclear physics, bioengineering and urban planning. Foundational issues – which are usually tightly linked to a discipline or even sub-discipline – will not be of interest to such an audience. Conceptually, I want to argue that philosophy *sui generis* is relevant for science education. Foundational issues are usually also addressed by the respective sciences, and are often claimed to fall into their respective domains. Methodology – the analysis and critical evaluation of scientific methods with the tools of the philosophy of science – on the other hand, is a core part of the philosophy of science. To argue that methodology is relevant for science education thus implies arguing that philosophy is – and this is what I intend to do.

So, for such an interdisciplinary methodology course, what should be taught? My starting point is the *standard PoS curriculum*. To clarify what I mean by that, I surveyed the nine main PoS textbooks published in the last 15 years, identifying chapter headings and main topics addressed in sub-chapter sections. Table 1 shows the result of this survey, ordered alphabetically. The first result of this survey is that textbooks show a lot of variation – of the 24 identified topics, the average textbook covers less than 8. Only three topics (explanation, induction & confirmation, realism) are shared by almost all textbooks, and 6 topics (causation, experiments, ethics, logical empiricism, natural kinds, observation, PoS as philosophy) are addressed (as chapter or subsection heading) by only one textbook each.¹⁰ Incorporating this variety, I will take all topics mentioned by one textbook or another as belonging to the standard PoS curriculum.

¹⁰ A quick survey of the relevant prospecti reveals that most of the top general PoS programs (according to the *Leiter report 2011*) follow this pluralist line. In LSE's Master program, "you can learn about both general philosophical problems raised by the sciences and particular philosophical-foundational problems that emerge in specific sciences". University of Pittsburgh's HPS1653 "provides a broad survey of a number of important issues in philosophy of science". Carnegie Mellon's 80-220 "examines some historical case studies (...) against which we will assess views pertaining to the significance, justification, and production of scientific knowledge". Cambridge's HPS undergraduate program investigates "how the sciences achieved their position in our society ... the processes of scientific knowledge, technological projects and medical strategies ... how and why these enterprises exert their powers and how they are trusted, contested and changed".

	Okasha 2002	Rosenberg 2005	PGS 2003	Curd & Cover 2012	Chalmers 1999	Ladyman 2002	Bird 1998	Salmon et al. 1999	Bortolotti 2008
Bayesianism			x	x	x				
Causation		x							
Demarcation	x			x	x				x
Experiment					x				
Explanation	x	x	x	x		X	x	x	
Ethics									x
Foundational issues in special sciences	x							x	
Heterodox views	x	x	x						
Hypothesis testing		x		x	x	X			x
Induction & confirmation		x	x	x	x	X	x	x	x
Kuhn			x		x				
Lakatos, Laudan, Feyerabend			x		x				
Laws		x		x			x		
Logical empiricism			x						
Modelling		x							x
Natural Kinds							x		
Naturalism		x	x						
Observation					x				
Popper			x		x				
PoS as philosophy		x							
Realism	x	x	x	x		X		x	x
Reduction		x		x					
Scientific reasoning	x			x			x		x
Scientific progress	x	x				X	x	x	x

Table 1

A lot of the standard curriculum is highly relevant for a course of the kind that I am proposing here. In particular, the standard PoS curriculum offers many important analytical tools for the analysis and evaluation of scientific practices. To name but a few, philosophers have developed tools like the HD model, the falsification concept or the model of inference to the best explanation, which help analyse and evaluate scientific hypothesis testing. Models are analysed with the help of notions like abstraction and idealisation, which philosophers helped to clarify. Scientific experiments are evaluated with the concepts of internal and external validity, which were developed by psychologists and philosophers. Philosophers, supporting the evaluation of scientific projects, have proposed numerous models of explanations. These conceptual tools and models are taught in most standard philosophy of science courses, and greatly facilitate

the analysis of scientific articles and practices.¹¹ Specifically, familiarity with these tools facilitate scientists' understanding, critical reflection, flexibility, innovativeness, interdisciplinarity and communicative skills – the very reasons, as I argued in the previous section, for teaching PoS to scientists.

However, the standard curriculum almost exclusively offers tools for an epistemic analysis of scientific practices. Yet there are important aspects of scientific practices that are non-epistemic in character and require corresponding analytical tools. Perhaps considerations of *ethical aspects* are the most obvious non-epistemic analyses of scientific practices. These comprise topics as diverse as honesty in research, ethical conduct of experiments, and considering the consequences of one's research.

Issues of *scientific honesty* arise when researchers misreport their empirical findings (e.g. by “cooking” or straightforwardly creating the data), when they plagiarise, denigrate the involvement of co-researchers or otherwise distort the research process. *Ethical conduct of experiments* concerns the way human and animal experiments ought to be conducted, and whether they are permissible for certain purposes at all; how experiments with human tissue are permissible; as well as how experiments that bear other risk (pollution, potential disasters, potential political or ethical complications) should be handled. Today, there are many codes of conduct, usually in the form of a catalogue of specific rules of behaviour, which address these issues (e.g. COPE's *Code of Conduct and Best Practice Guidelines for Journal Editors* for scientific honesty; the Nuremberg Code, the Declaration of Helsinki and the US Code of Federal Regulations Title 45 Part 46 for human experiments).

Students are commonly taught (if at all) only these rules, and no effort is made to instruct students in understanding, critically reflecting and applying these rules – an understanding that would make it easier for scientists to develop an intuitive comprehension of these issues and that would make it easier for them to judge these issues in novel contexts. The situation here is quite similar to that of a conventional methodology. Scientists always are moral agents who make choices based on their moral values and norms. Instructing them about codes of conduct might affect their value and norm commitments, but does typically not make them critically reflect on or seek to justify these commitments. Philosophers, with the analytic tools of normative and applied ethics, are capable of instructing students in analysing, comparing and evaluating such bodies of rules, and thus greatly help in a better understanding and a more flexible and comprehensive application of these rules.

But what are the right tools that philosophers should teach in such a course? Many textbooks on ethical issues in science commence by giving a basic introduction into the main ethical theories, e.g. utilitarianism, Kantian deontology and virtue ethics (cf. Macrina 2005, Israel & Hay 2008). I believe that such a foundational approach is

¹¹ It is noteworthy in this context that philosophers of science are rarely trained in applying these tools and models to scientific texts. At best, specific (excerpts of) texts are cited as illustrations for these tools. Hence as a side effect of my argument here, it seems recommendable to develop formal training methods for philosophy of science students to practice the analysis of scientific texts.

misplaced in such a course. One reason to avoid it is time limitation: at best, such a course can afford two lectures on non-epistemic topics. Teaching basic ethics would already take up half of that. Another reason is that the gain in understanding and critical reflection will be minimal from teaching foundations in such a clipped style. Instead of trying to analyse rules in terms of underlying principles, philosophers should address concrete cases, elicit intuitive judgments from students, and analyse these judgments as expressions of specific values (like e.g. honesty, openness, transparency, confidentiality). Philosophers' contribution to such a course thus mainly lies in their ability to analyse and assess arguments, not in their knowledge of ethical theories.¹²

Philosophers play a similar role when discussing the ethically relevant consequences of research, both in terms of what bad can be avoided and what good can be done. Such exercises are quite controversial, already from the claim that scientists are at all responsible for the consequences of their research. Historical examples, as well as thought experiments exemplify what such consequences might be. They also might help eliciting students' normative intuitions about such scenarios. Philosophy here facilitates the reflection about such questions, rather than helping to find a solution. As many students are concerned about such questions, but do not find them addressed in their courses, philosophy should provide instruction here as well.

Besides ethical discussions, science also is subject to other non-epistemic analyses. The *deliberational domain* concerns itself with the connection between scientific research and policy making. One important aspect concerns the possible conflict between epistemic and deliberational values. Take for example the design of a statistical test, in particular the setting of the type I and II error margins. Epistemically, a minimal type I error margin most often is desirable. Yet for many policy contexts, a type II error can have a much greater weight than a type I error (e.g. when testing for the toxic effect of a substance or the defect of a machine). Strict standards of proof in science and valid policy reasons may thus conflict (Hansson 2012). This discrepancy between judging evidence for epistemic or for policy purposes is rarely discussed in science courses, but is a serious problem many scientists will encounter in their professional career. Philosophy, through the tools of decision theory, can help instructing students how to analyse and tackle such problems.

Finally, scientists often find themselves involved in the *political domain* through their role as *experts*. This role is not easy to navigate, requiring weighting the epistemic authority of expertise against democratic principles, yet there is little instruction in science courses how to handle it. Political philosophy, by reflecting on the role of scientists in a liberal democracy, and by analysing this role as an expression of a democratic decision (Turner 2001), contributes to the understanding the function of experts and thus can contribute to instructing scientists in this role.

To conclude, there are at least five non-epistemic analyses that are relevant for science students and that philosophers are experts in: (i) scientific conduct, (ii) experimentation

¹² Recent publications like Macrina (2005) and Resnik (2009) offer plenty of case material for such analyses.

with animals and humans (iii) reflecting on consequences of research, (iv) risk analysis, and (v) the scientist's role as expert. None of these topics are part of the standard PoS curriculum – although they should be for a PoS course to scientists. Specifically, the ability to engage in these analyses improves scientists' understanding, their abilities to critically reflect, communicate to the public and to incorporate normative non-epistemic considerations in their professional judgments – the very reasons, as I argued in the previous section, for teaching PoS to scientists.

The standard PoS curriculum not only needs expansion, however. It also contains many topics that are not relevant for an interdisciplinary audience of science students. Specifically, topics that require too much background knowledge in either philosophical theories or analytical tools, and topics that are mainly interesting for the philosopher, should be excluded.

Some topics require more knowledge about philosophical theories than science students commonly have (or could acquire in the limiting time of such a course). For example, how is a student to appreciate some of the main aspects of scientific realism (beyond the very basic idea of the belief in the reality of something) if he lacks a basic understanding of what metaphysics and semantics are about? It is rather difficult to convey the meaning of “a mind-independent existence of entities” or of “having truth values” without anticipating a minimal understanding in these domains. Yet instructing students in these domains will quickly go beyond the scope and capacity of such a course.

Additionally, there are analytic tools that have proven very helpful to the philosopher of science in her reconstruction of scientific arguments and practices, but which are too involved to teach within the confines of a one-semester course to science students. Examples of such tools include First Order Logic and Bayesian decision theory. In contrast, the tools and models that we teach must be simple and easy to handle, so that the students can actually reap the benefits from them during the time of the course.

My first argument against the relevance of certain topics and methods of the standard PoS curriculum was based on the time constraints imposed on such a course. In principle, it would be desirable to teach the students Bayesian decision theory and expose them to an intense discussion of the problem of induction from Sextus to Goodman. It's just that there isn't time, and that we have to compromise in order to teach an optimal course. A second, stronger argument claims that there are some topics that are primarily of philosophical interest and irrelevant for teaching PoS to scientists *per se*.

Where “mere” philosophical interest ends and relevance for analysis, comparison and evaluation of scientific practices begins will be difficult to determine in any general way. Let me instead give a few examples that illustrate this divide, without drawing a very specific line. A first instance is the *historical and comparative perspective* on philosophy that philosophers of science often assume. Too often, PoS courses consist of going through a sequence of schools and “isms”. This might be instructive when teaching philosophers of science – but what relevance has it to the science students that she knows which program Carnap followed or in which ways Lakatos disagreed with Kuhn? The

useful concepts and models from these different authors and schools should of course be taught, and it might be even useful to contrast them with each other in their respective functions and shortcomings, but there is no reason to teach science students how to historically classify these models and tools, and how they connect to a larger program.¹³

Of course, the philosopher's ability to analyse these models and tools from a historical perspective might be very helpful for teaching such a course. The difference I insist on here, however, is between knowledge derived from historical analysis that might be relevant for such a course, and instruction in such historical analysis itself, which I claim is not relevant. Similar arguments can be made about many standard topics in PoS. When philosophers discuss, for example, (i) scientific inferences as cases for epistemological discussions, e.g. about the relation between knowledge and truth; or (ii) realism, reductionism, natural kinds issues as material for metaphysics, they investigate science mainly out of philosophical interest, which arises from long-held positions, discussions and conflicts within the discipline of philosophy. While knowledge of these debates might help the instructor in optimising such a course, I doubt that instructing science students about these debates will contribute to making them better scientists in any of the senses I discussed in section 2 of this paper. Of course, science may develop an interest in these debates (they are, after all, fascinating). But philosophers should be carefully demarcating those topics that scientists need to be instructed in (in order to become better scientists) from those that might be of interest to them. The latter interest can be satisfied with (non-mandatory) courses, but these topics should not be taught in a course mandatory for all science students.

To conclude this section, "being part of the standard PoS curriculum" is neither a sufficient nor a necessary condition for including such material in a PoS course for scientists. It is not sufficient because many topics, especially those relating to non-epistemic analyses, are missing from table 1. It is not necessary, because many entries of table 1 are too demanding (e.g. topics like "Natural kinds" or methods like "Bayesianism") or not relevant (e.g. the many historical entries, or topics like "Realism" or "Reduction") for such a course. Consequently, there are good reasons to consider revising our own teaching program, if we want to create a successful PoS course for scientists.

4. How?

A mandatory course in PoS for students from different scientific disciplines will likely have the following five characteristics. First, most students have no training in the humanities beyond high school level and don't know what to expect from a philosophy course. Second, many students – according to my impression – have no intrinsic

¹³ Note that I am making an argument against assuming a historical and comparative perspective on *philosophy*. Clearly, to teach a successful PoS course requires using many examples from science, and hence requires relying on the history of *science*. But these two historical perspectives are clearly distinct, and one can extensively employ the latter without having to make use of the former.

motivation to take such a course. Third, many students don't know about and have not engaged in scientific research. Fourth, students will come from widely differing fields. Fifth, the student body will be large.

These student and course characteristics are likely to clash with typical teaching strategies of most philosophers of science, who usually expect familiarity with the humanities (and, more specifically, with basic philosophical ways of thinking and instructing) and a relatively high level of motivation. I will therefore take the above five characteristics as constraints on how one can teach such course, and suggest strategies how to deal with these constraints successfully.¹⁴

The first constraint comes from the fact that most science students (in countries without a liberal arts college education system like the UK, Germany or Sweden) have no training in the humanities beyond high school level and don't know what to expect from a philosophy course. This means that philosophy teachers must keep their conceptual and method presumptions in check. *Conceptually*, it is obvious that we cannot expect that science students will perk up if we promise them a solution to the Problem of Induction or the Raven's Paradox. These concepts need explanation, and it is a justified question whether it is worth the time to explain them in such a course. But science students will also have a problem with much more basic concepts – e.g. deduction or causation – that philosophers tend to assume as known. Worst, science students often will not see the importance (and sometimes not comprehend the concept of) of notions as basic to philosophical work as argumentation or justification. From my own experience, I know how easy it is to presume these as known, and how detrimental for a course such assumptions are.

To avoid such mistakes, the instructor should be as clear as possible about the motivations for such a course (where the discussion from section 1 might come handy). Why are we focussing on argument and justification? Giving examples from real science that show where arguments are problematic and where a justification is wanting will be crucial in answering this question.

A related problem is that most science students will not be familiar with the *methods* of philosophy of science. Science students tend to learn science ahistorically – they focus on only those theories and models that are part of the current scientific corpus. This sense of progress, combined with a large set of discarded theories, is much less prevalent in philosophy. Philosophers therefore tend to relativize: a model might be better in some respects than another model, yet there is almost always a counterargument that needs to be considered and can rarely be conclusively refuted. Some philosophers have developed this practice into a veritable case of *counterexampelitis*: pointing out the failures of the concept or model tends to take up more space than its advantages.¹⁵ This will not impress

¹⁴ I will restrict myself here to how arguments from section 2 affect optimal teaching strategies, and also take into account what philosophers of science typically consider optimal teaching. For more general suggestion on optimal teaching strategies, see Biggs & Tang (2011) or Mazur (1997).

¹⁵ As the attentive reader will notice, I am here – on the philosophical level – explicitly *not* practicing what I preach for the science level: while I argue that the ability to reflect and analyse scientific methods is a

most science students, whose own training will compel them to discount theories with too many counterexamples. While such a clash of methods cannot be fully avoided, I believe that philosophers would do well to mitigate it – by showing what works, by training the application of tools, and by checking (perhaps even testing) what works and what does not.

The second constraint is of course related to the first – non-familiarity often yields a lack of motivation. This might express itself as a general mistrust against the humanities, and sometimes as a “hard science” hyperbole. Such sentiments are often fed by the students’ previous science education, unfortunately. Students are not exposed enough to real scientific research. Instead, they often base their knowledge of science on the textbooks of their discipline, and these textbooks tend to downplay the prevalence and importance of methodological problems.

For a successful course, it is of crucial importance that this scepticism and hyperbole is early and directly addressed. One way to do so is by discussing *scientific errors* extensively. Analysing such historical instances of errors serves three purposes. First, it shows that scientific error is common, and often related to methodological problems. This should help motivate why methodology is important to scientists. Second, one can learn from an analysis of error: by seeing where something went wrong, the students get a feeling for the difficulties and subtleties of actual scientific research. Third, students should learn that error not only is ubiquitous in research, but also necessary. Because science is in its core dealing with uncertainty, any attempt to safeguard against error completely will make research sterile. Error thus holds a crucial role in teaching science students (Allchin 2001).

Another important strategy to address scepticism or hyperbole is through discussing the *demarcation problem*. I am rather sceptical about demarcation *criteria* proposed by philosophers, as they inevitably seem to do injustice to real science. But the *problem* how to distinguish genuine science from junk science or “merchants of doubt” is very much alive and should motivate students to think about their methodological preconceptions and conventions.

A final strategy to deal with scepticism or hyperbole is to pick up students at points in their scientific development where they are themselves faced with method choices and their justification. This often happens rather late in their academic career – e.g. when writing their Master’s thesis or the first paper of their PhD – but my own experience is that these are key moments when students become much more open to the kind of problems philosophers of science want to discuss with them.¹⁶

crucial part of a scientists’ education, I want to avoid too much reflection or analysis of the philosophical methods used in this. The reason for this is that philosophy in this contexts provides a service for the scientists: the methods for analysis and reflection, and it needs to be confident of its methods to render this service well.

¹⁶ A related option – that avoids some of the time-sensitive issues of this approach – is to let them write a criticism of a research paper they had written at an earlier time, applying the concepts and tools that they hopefully acquired during their PoS course.

The third constraint on how to teach such a course lies in the fact that many science students know surprisingly little about science. These are often good students by the standards of textbooks and even lab examinations, but they lack knowledge even about basic research practices, for example how to develop a research question, how to formulate a hypothesis, how to design an experiment, or how to build a model.

The first step to deal with this problem is to clarify to the students both the distinction between textbook and actual research papers on the one hand, and research papers and research practices on the other (Schickore 2008). This can be supported by engaging the students with current research articles, requiring them to summarise and critically discuss the research methods used therein. To engage them with actual research practices, one can have students reflect on their own first steps in this direction – e.g. by having them write methodological discussions of their lab class or their lab log, or by having them write a methodological essay about their Master or (parts of) their PhD thesis.

Another strategy to deal with limited familiarity with research practices is to *simulate* them. Because of the complexity of current science, this requires either finding a fictional “toy” situation, or going back into the history of science. For a number of years, we have used the toy situation exercise “what’s in the box?” to teach students about dealing with uncertain evidence.¹⁷ Alternatively, Hasok Chang has proposed making history of science useful for this purpose. Students are required to ‘Practice 18th-Century Science Today’, engaging with 18th-Century science texts as if the knowledge level then was that of today, and writing responses and critics to these texts.¹⁸

The fourth constraint consists in the interdisciplinary diversity that such a course will likely have.¹⁹ Strategies to deal with this require both diversity in illustrative cases and examples, so that most students’ discipline at least at some point will be addressed and choosing cases in such a way that they have maximal relevance for the lecture audience.

The final constraint derives from the large size that such a course will typically have. Many of the problems addressed in this section and in section 3 could be easily solved if the group were small, and more tutorial- or seminar-style interaction was possible. However, as I am proposing the arguments of this paper partly so that philosophers of science will have large audiences, there is no reason to bewail here the lack of intimacy. Instead, lectures in front of large audiences will be inevitable. By increasing interactions, e.g. through introductory quizzes or concept tests (Mazur 1997), such lectures can be

¹⁷ This method was invented in the 1990s by Sören Törnkvist of the Royal Institute of Technology (KTH) in Stockholm (communication with Lars-Göran Johansson of Uppsala University). A similar technique was presented by Hardcastle and Slater at PSA2012 in San Diego.

¹⁸ For more information on Chang’s teaching method, called the “Nicholson Journal”, see www.ucl.ac.uk/sts/staff/fellows/jackson/nicholsons.

¹⁹ Some commentators have argued that this constraint could be avoided by offering philosophy of science courses designed for specific disciplines, e.g. physics, economics or dentistry (as for example offered at Aarhus University). However, as I argued in section 2, an interdisciplinary perspective is also an important teaching tool: studying the methods of other disciplines helps to overcome a conventional methodology. Thus I am hesitant to call for a specialisation of PoS courses at the cost of an interdisciplinary perspective, although this perspective is of course a challenge for teaching.

designed in more engaging ways. These lectures then should be enriched with a few purpose-built seminars, which require mandatory preparations from the students and ask them to employ the analytical tools, learned in the lectures, to specific texts or problems. Finally, these lectures can be further enriched by the projects discussed above, including research paper presentations, lab log methodology, simulation exercises, or methodological essays on Masters or PhD theses.

Of course, this section could only raise a few problems that arise from how to teach such a PoS course for science students and only offer a few strategies how to deal with them. The aspects discussed are far from complete. Nevertheless, I believe that it is a useful starting point for a big and exciting project: providing the best PoS instruction suited to the needs of science students.

5. Conclusions

The ability to identify, analyse and evaluate justifications of scientific methods is needed by and highly useful for scientists. Philosophers of science are uniquely qualified to teach such skills, and therefore should provide specifically-designed courses of PoS to scientists on a much broader base than currently practices.

Furthermore, teaching philosophy of science to science students offers great potential benefits to philosophers. To name but a few: (i) we get to spread the result of our research to a large audience, (ii) we can make a difference in the way science will be done, (iii) we learn from interacting with scientists, (iv) we can show that philosophers have something important to contribute to the scientific community in times of general university scale-backs, (v) we open new sources of revenue for philosophy departments, ensuring the continuation of our work, (vi) we open new employment options for those who we train to become philosophers of science.

Yet in order to reap these benefits, we have to convince administration, scientists and students. To do so, we need to revise our standard philosophy curriculum to provide optimal instruction for budding scientists, and we have to adopt new teaching strategies for such a course.

In this paper, I have argued that we should convince administration, scientists and students by showing that the genuinely philosophical work of analysing procedures of evaluating and justifying scientific knowledge production contributes to making better scientists. These arguments imply that we change what we teach in PoS courses, and how we teach it. We should revise the standard PoS curriculum in such a way that it focuses on this reconstruction work and tools necessary for the analysis, comparison and evaluation of scientific practices; and we should expanded it by ethical and decision-theoretic analyses of scientific practices. Finally, we should teach such a course in a way that picks up science students from the problems and perspectives they likely experience at that point of their training, and be mindful of their possible lack of philosophical sophistication, lack of motivation and lack of experience with actual scientific research.

Acknowledgements

For valuable discussions and insightful comments I thank Hanne Andersen, Mieke Boon, Marion Godman, Inkeri Koskinen, Jaakko Kuorikoski, Caterina Marchionni, Samuli Pöyhönen, Kristina Rolin, and the audiences at Helsinki University and at the 2013 Meeting of the German Society for Philosophy of Science in Hannover. Remaining mistakes are my own. Financial support from the Swedish Research Council (Vetenskapsrådet) and the Finnish Centre of Excellence in the Philosophy of the Social Sciences (TINT) is gratefully acknowledged.

REFERENCES

- Allchin, D. 2001. Error Types. *Perspectives on Science* 9: 38-59.
- Blachowicz, J. 2009. How science textbooks treat scientific method: A philosopher's perspective. *The British Journal for the Philosophy of Science*, 60(2): 303-344.
- Biggs, J., & Tang, C. 2011. *Teaching for quality learning at university*. Maidenhead: Open University Press.
- Bird, A. 1998. *Philosophy of Science*: London: Routledge.
- Bortolotti, L. 2008. *An Introduction to the Philosophy of Science*. Cambridge: Polity.
- Cable, M. 2005. *Calibration: A Technician's Guide*. ISA.
- Chalmers, A.F. 1999. *What is this thing called science?* Brisbane: University of Queensland Press (3rd ed.).
- COPE. 2011. Code of Conduct and Best Practice Guidelines for Journal Editors. http://publicationethics.org/files/Code_of_conduct_for_journal_editors_Mar11.pdf. Accessed June 12th, 2013.
- Curd, M. and J. A. Cover. 2012. *Philosophy of science: The central issues*. New York: WW Norton (2nd ed).
- Doran, P. 2006. Cold, Hard Facts. *The New York Times*. July 27, 2006.
- Ennis, R. 1979. Research in Philosophy of Science Bearing on Science Education. In *Current Research in Philosophy of Science*, eds. P. Asquith and H. Kyburg, 138–70. East Lansing, MI: Philosophy of Science Association.
- European Commission. 2008. *The European Qualifications Framework for Lifelong Learning (EQF)*. Luxembourg: Office for Official Publications of the European Communities. http://ec.europa.eu/education/pub/pdf/general/eqf/broch_en.pdf. Accessed May 15th, 2013.
- European Commission. 2012. *FP7 Grant agreement, Annex II*. ftp://ftp.cordis.europa.eu/pub/fp7/docs/fp7-ga-annex2-v7_en.pdf. Accessed May 15th, 2013.
- Grayson, D. J. 2006. Rethinking the content of physics courses. *Physics today* 59(2): 31.
- Godfrey-Smith, P. 2003. *Theory and reality: An introduction to the philosophy of science*. Chicago: University of Chicago Press.

- Hansson, S.O. 2012. Risk. *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Edward N. Zalta (ed.). <<http://plato.stanford.edu/archives/win2012/entries/risk/>>. Accessed June 15th, 2013.
- Hertwig, R. and A. Ortmann. 2001. Experimental practices in economics: A methodological challenge for psychologists? *Behavioral and Brain Sciences* 24(3): 383-403.
- Hillerbrand, R. and M. Ghil. 2008. Anthropogenic climate change: Scientific uncertainties and moral dilemmas. *Physica D: Nonlinear Phenomena* 237(14): 2132-2138.
- Hodson, D. 1991 Philosophy of Science and Science Education. In *History, Philosophy, and Science Teaching*, ed. M. Matthews, 19–32. Toronto, ON: OISE Press.
- Hoover, K. D. and S.J. Perez. 2000. Three attitudes towards data mining. *Journal of Economic Methodology* 7(2): 195-210.
- IPCC. 2007. Summary for Policymakers. In: *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, eds. Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- Israel, M. and I. Hay. 2006. *Research ethics for social scientists*. Thousand Oaks, CA: SAGE Publications Limited.
- Kitcher. P. 2003. *Science, Truth and Democracy*. Oxford: Oxford University Press.
- Ladyman, J. 2002. *Understanding Philosophy of Science*. London: Routledge.
- Lehtinen, A. and J. Kuorikoski. 2007. Computing the perfect model: Why do economists shun simulation? *Philosophy of Science* 74: 304-329.
- Macrina, F. L. 2005. *Scientific integrity*. Washington, DC: ASM Press.
- Martin, M. 1976. The Relevance of Philosophy of Science for Science Education. In *PSA 1974*, R. Cohen, C.A. Hooker, A. C. Michalos, J. van Evra, 293–300. Dordrecht: D. Reidel.
- Matthews, M. 1994. *Science Teaching: The Role of History and Philosophy of Science*. New York: Routledge.
- Mazur, E. 1997. *Peer instruction: A user's manual*. Upper Saddle River, NJ: Prentice Hall.
- Medawar, P. B. 1963. Is the scientific paper a fraud. *The Listener* 70(12): 377-378.

- Medawar, P. B. 1969. *Induction and intuition in scientific thought*. London: Routledge.
- Okasha, S. 2002. *Philosophy of science: A very short introduction*. Oxford: Oxford University Press.
- Ontario Council of Academic Vice Presidents (OCAV). 2008. *Guidelines for University Undergraduate Degree Level Expectations*.
http://www.uwo.ca/univsec/handbook/general/OCAV_Guidelines_2005.pdf. Last accessed May 16th, 2013.
- Prather, C. M., D. M. Choate, M. J. Michel and T. A. Crowl. 2009. Putting the “Ph” back into “PhD”: framing graduate research in a theoretical context. *Frontiers in Ecology and the Environment* 7(7): 389-390.
- Rosenberg, A. 2005. *Philosophy of Science A Contemporary Introduction*. London: Routledge (3rd ed.).
- Salmon, M.H., J. Earman, C. Glymour, J. Lennox, K.F. Schaffner, W.C. Salmon, J.D. Norton, J.E. McGuire, P. Machamer, J.G. Lennox. 1999. *Introduction to the Philosophy of Science: A Text By Members of the Department of the History and Philosophy of Science of the University of Pittsburgh*. Indianapolis, IN: Hackett Publishing Company.
- Schickore, J. 2008. Doing Science, Writing Science. *Philosophy of Science* 75(3): 323-343.
- Shamoo, A. E. and D.B. Resnik. 2009. *Responsible conduct of research*. Oxford: Oxford University Press.
- Siegel, H. 1989. The rationality of science, critical thinking, and science education. *Synthese* 80(1): 9-41.
- Stigler, S. 2008. Fisher and the 5% level. *Chance* 21(4): 12-12.
- Turner, S. 2001. What is the Problem with Experts? *Social Studies of Science* 31(1): 123-149.
- Varian, H. R. 1997. How to build an economic model in your spare time. *The American Economist* 41: 3-10.