Philosophia
https://doi.org/10.1007/s11406-024-00721-9

Philip Kitcher's Purge of Philosophy

Timothy Williamson[1]

Received: 26 January 2024 / Revised: 26 January 2024 / Accepted: 30 January 2024

According to Aristoxenus, when Plato gave his lecture 'On the Good', he lost his audience. As he went on about numbers, geometrical figures, and astronomy, finally concluding 'The Good is One', they realized that they were not going to hear anything useful about means to health, wealth, or happiness. Most of them left disgruntled.

According to Philip Kitcher in his latest book, *What's the Use of Philosophy?*, much of contemporary analytic philosophy has lost its audience, because it offers them nothing useful, though he makes no analogy to Plato (43, 58; all page references to the book). Kitcher's sympathies are with the fleeing audience. He wants philosophy to change, to win them back.

Surprisingly, Kitcher treats his call for reform as novel. He complains that the question of the worth of analytic philosophy '*never even surfaces*' and seeks 'to *initiate* a process of self-interrogation' (119, 145; his italics). This is unfair to a long-standing genre: critiques of analytic philosophy as pedantic, pointless, boring, shallow, trivial, irrelevant, etc. As an undergraduate at Oxford in the mid-1970s, I was familiar with such complaints; one of my tutors was Alan Montefiore, who had started in analytic philosophy, and still had to teach it, but had switched his own allegiance to philosophy in the style of Jacques Derrida. In his hopes for the future of philosophy, Kitcher goes further back into the past for a model, to John Dewey. In 1959, Ernest Gellner's crude attack on linguistic philosophy, *Words and Things*, had caused quite a stir: Bertrand Russell's denunciation of Gilbert Ryle's refusal to have it reviewed for *Mind* initiated a month-long controversy in the correspondence columns of *The Times*. Well before that, some of the same accusations were made against logical positivism, by both the extreme Right and the extreme Left. Throughout my life in academic philosophy, a recurrent event has been reading, with mild curiosity, the latest contemptuous dismissal of analytic philosophy. Many other philosophers have had similar experiences.

✉ Timothy Williamson
timothy.williamson@philosophy.ox.ac.uk

[1] University of Oxford, Oxford, UK

Published online: 02 March 2024

I do not mean to suggest that all critiques of analytic philosophy are interchangeable. Kitcher's variation on the theme is more nuanced than most. He issues more exemptions, especially for his friends in the philosophy of science. He is also more optimistic about the prospects for reform: analytic philosophy is not beyond hope, provided it gets rid of most of its metaphysics, epistemology, philosophy of language, philosophy of mind, and metaethics (24, 117, 147). Instead, it should spend more time facilitating the moral progress of humanity, and doing stuff to attract the interest of non-philosophers.

One area where valuable work does get done, in Kitcher's view, is the philosophy of the special sciences, including physics, biology, psychology, economics, and linguistics (23). It therefore seems fair to judge the book as a contribution to the philosophy of philosophy, by the standards of the philosophy of the special sciences. Of course, the philosophy of a special science does not usually take the form of a scathing critique of much contemporary work in that special science, but we may assume for the sake of argument that it is not in principle illegitimate for it to take that form.

By general agreement, work in the philosophy of a special science should be based on detailed engagement with work in the special science itself, so we may fairly apply that test to Kitcher's work in the philosophy of contemporary analytic philosophy.

Alarm bells ring when we read a passage like this (60):

> Contemporary "analytic" philosophy, dominant in the English-speaking world, retains part of the positivist program. Philosophy's central task is seen as one of providing analyses of concepts, analyses exact enough to make the concept completely clear.

That description is several decades out of date. Although conceptual analysis still has some defenders, there is no longer a consensus that it is philosophy's 'central task'. There are signs of a growing consensus that it is *not* philosophy's central task. One would not be allowed to get away with such a mischaracterization in the philosophy of a special science.

Kitcher's comments on analytic epistemology are slightly more up to date. He sighs (62):

> The zeal for analyzing "S knows that *p*" abates, but controversies about the merits of externalism and internalism and contextualism thrash on. And on.

His yawn at the epistemological debate between externalism and internalism is striking, for a recent influential contributor on the externalist side is Srinivasan (2020), whom he later presents as a brilliant model for the kind of philosopher we most need (172-3, 178). Her arguments connect the epistemological question with political issues about the role of knowledge which people possess while unable to vindicate it discursively. Kitcher could also have reflected that internalism makes the justification of belief purely a matter of internal coherence: the most abhorrent beliefs of a consistent Nazi may count as fully justified by internalist standards (Williamson, 2023). Yet, from a wider intellectual perspective, internalism is no mere philosopher's caprice. The most popular version of probabilistic epistemology, both within philosophy and

far beyond, is subjective Bayesianism, an internalist approach on which the only epistemic norms are purely formal coherence constraints on the agent's credences (degrees of belief): synchronically, they must satisfy the axioms of mathematical probability theory; diachronically, they must be updated only by Bayesian conditionalization. A more traditional route to internalism is quasi-Cartesian: the justification of belief should depend only on factors accessible from the conscious subject's perspective. Although such motivations for internalism are arguably wrong-headed, they speak to issues central to our self-understanding. A community of epistemologists who did not engage with those questions would be missing something deep.

When philosophers of science discuss epistemological issues, they often seem to take internalist assumptions uncritically for granted, in ways they could not get away with in 'mainstream epistemology'. Kitcher praises philosophers of science for not wasting their time engaging with 'mainstream epistemology' (42). One wonders whether such dogmatism on the internalism-externalism issue is part of what he finds admirable.

Kitcher's mention of contextualism in the quoted passage also backfires. Far from being a niche topic confined to a sub-community of epistemologists, it is one where epistemology has interacted fruitfully with semantics as part of both linguistics and philosophy of language; some contextualist hypotheses are open to experimental test.

In a different connection, Kitcher notes 'the (healthy) trend among some philosophers of language to engage more closely with work in linguistics' (83). This presents the flow of ideas as being from linguistics into the philosophy of language. On closer inspection, the picture is more complex. As one might expect, the areas of greatest interaction are semantics and pragmatics. In both, ideas from analytic philosophy of language have been massively influential in shaping the theoretical frameworks within which linguists work. One can easily document the influence in linguistics of J. L. Austin on illocutionary acts, Paul Grice on conversational and conventional implicature, Richard Montague on grammar, Peter Strawson and Robert Stalnaker on presupposition, Saul Kripke and David Lewis on intensional semantics, David Kaplan on content, character, and monsters, Donald Davidson on adverbs, Hans Kamp on discourse representation, Jon Barwise and John Perry on situations, and so on. Linguists such as Barbara Hall Partee, Angelika Kratzer, James Higginbotham, and many more took up and applied those theoretical approaches, and developed them further. In these areas, interaction between linguists and philosophers is comparatively easy, because they have so much in common methodologically and theoretically, with much of the shared heritage the outcome of the flow of ideas from analytic philosophy of language into formal semantics and pragmatics. Research communities on topics such as conditionals and modal verbs seamlessly span the divide between departments of philosophy and departments of linguistics. There are slight differences in style and skill-sets—those trained in linguistics tend to bring in more considerations from syntax and examples from languages other than English, while those trained in philosophy and logic tend to prove more formal theorems—but each side sees the relevance of the other side's results, and communication rarely breaks down. There is no more reason to diagnose pathologies in analytic philosophy of language than there is to diagnose them in semantics and pragmatics as branches of linguistics.

The role of possible worlds semantics is especially notable, because critics of analytic philosophy are so liable to treat the invocation of possible worlds as an irresponsible metaphysical extravaganza, a grotesque symptom of detachment from any serious interest in reality, a paradigm of what is wrong with contemporary analytic philosophy. But linguists cannot reasonably be charged with overweening metaphysical ambition: they quickly saw the value of the possible worlds framework because it is exactly what one needs to explain perspicuously how the meanings of mundane auxiliary verbs such as 'can' and 'might' combine with the meanings of other expressions to generate the truth-conditions of sentences in which they occur. It is a salutary warning not to judge the cognitive potential of ideas on first appearances.

A third area of contemporary analytic philosophy which could do with purging, in Kitcher's vew, is metaethics. Quite how much purging is unclear, since Kitcher does some quick metaethics of his own in the book, and refers us to another book where he does more. For example, he tells us that he takes 'morality and ethics to be human inventions' (17) but 'it is still possible to find a kind of objectivity in our moral codes' (18). If Kitcher's metaethical reflections are not to be purged, why should the same dispensation not apply to everyone else's metaethical reflections too? He writes (105):

> once the tale I have briefly rehearsed is clearly appreciated, there's an obvious way to reform meta-ethics: give up the intricate exchanges about meaning and truth; focus instead on improving whatever methods we have for making moral progress.

In this bright vision of the future, we all accept Kitcher's metaethical views (by 'clearly appreciating' his 'tale'), so there are no more 'intricate exchanges about meaning and truth' because there is no more metaethical disagreement; we can all cooperate in preparing the way for moral progress on Kitcher's terms. However, he may not find it so easy to install himself as metaethical dictator. Some awkward characters may insist on making objections to his 'tale', his arguments with them may get complicated (unless complexity in metaethics is outlawed); there may even be a recurrence of 'intricate exchanges about meaning and truth'. After all, as already seen, analytic philosophy of language remains in good standing, with its theories about meaning and truth.

Kitcher treats metaethical uncertainty like a professional deformation of philosophers. But that is insensitive to the conditions of modern life in a complex society. Almost everyone has experienced interactions with others whose norms and values evidently differ from one's own, in practically relevant ways. People follow different religions, or none; they support different political parties; they respect different dietary restrictions. Even without prompting by a philosopher, such encounters can provoke a reflective person into at least inchoate metaethical thought about whether norms and values are in some sense absolute or relative, objective or subjective, real or merely apparent. In a modern society, the terms for asking such questions are readily available, however vague. When the questions are asked, it is far from obvious how to answer them. If philosophy did *not* take those metaethical questions seriously, and at least try to give them reasoned answers, it would arguably be failing to fulfil

its responsibility to society. Even if philosophers feel ill-equipped to address those questions, who is *better* equipped than they are to address them? When people turn to philosophy, they do not always expect or even want clear, unequivocal answers to their questions. They may just be seeking help in taking their thoughts a few steps further on.

In any case, as long as people are worried that norms or values may be in some sense relative, or subjective, or unreal, they are liable to find Kitcher's call to help prepare the way for moral progress question-begging. They may wonder: progress by what standard? In short, metaethics is much harder to sidestep with integrity than he suggests.

Kitcher asks what would change if, tomorrow, the philosophical community were to achieve consensus in favour of moral realism, but without consensus in favour of any specific moral code (104). He suggests that very little else would change. That may be true. But he does not consider the opposite thought experiment, in which, tomorrow, the philosophical community achieves consensus in favour of some hard-line form of moral *anti*-realism, with no substitute *ersatz*-morality to fall back on: moral norms and values are seen as just a myth, as unreal as fairies. Even then, not much else might change at first. After all, even advocates of the crudest moral expressivism did not rush into the street to commit adultery. But how psychologically and socially stable is the combination of conventional adherence to a moral code with contempt for it as a fairy story in the long run? Some people are better than others at disjoining theory and practice. The collapse of moral ideology would not be a trivial event. The interrelations between ethics and metaethics run deeper than Kitcher suggests.

The nearest the book gets to a case study of analytic philosophy is a discussion of its use of hypothetical cases to test a philosophical theory (70–80). He says 'analysis' rather than 'theory' (70), but his examples—such as trolley cases in moral philosophy—involve testing theories, not analyses. Kitcher reports that, when he presents such scenarios in class, his undergraduates are very reluctant to make any judgment as to what to do, and desperately seek loopholes in the small print of the case descriptions. He claims that the cases outrun our imaginative capacities: for example, the agent in the scenario is supposed to be absolutely certain that the relevant conditions obtain, but we cannot fully imagine what such absolute certainty what would be like.

I suspect that Kitcher is doing his students' imaginations an injustice. In my experience, most students have little difficulty in imaginatively pushing various epistemic options off the table, just as in science classes most have little difficulty in engaging imaginatively with perfect spheres, frictionless planes, and perfectly closed systems. Stipulation plays a major role in imagination: we need it in assessing a conditional by imagining its antecedent holding and then assessing its consequent under that supposition (Williamson, 2020). I conjecture that Kitcher's students could imagine the hypothetical scenario quite well, understood that it put the agent in something like a moral dilemma, and were reluctant to give a verdict because they knew that, whichever horn they picked, Kitcher would immediately make its downside vivid, putting them on the spot. Desperately seeking a loophole is a natural response to being presented with such a dilemma. Contrary to what Kitcher claims, the scenario's unfamiliarity is not the problem: even in bizarre science fiction, sometimes there is an

obviously right choice and an obviously wrong one. If we could not engage imaginatively with bizarre science fiction, it would not be so popular.

Misleadingly, Kitcher frames much of the discussion around the question: 'What would *you* do in those circumstances?' That may be a useful pedagogical technique for bringing the problem home, but it distracts from the question supposedly at issue, which concerns not what you *would* do, but what you *should* do. Most people can imagine situations in which pride, fear, embarrassment, hatred, greed, or lust *would* lead them to do something they know they *should* not do. Many thought experiments need to be assessed from a third-person perspective, because they turn on factors unknown to the protagonists: for example, lucky coincidences in Gettier cases, and the chemical constitutions of liquids in Twin Earth cases. Kitcher's emphasis on the first-person perspective in assessing hypothetical cases is a distraction.

As all these examples indicate, as a critique of contemporary analytic philosophy, the book does not meet the standards of the philosophy of the special sciences. It is too impatient, too biased, too unsympathetic, too contemptuous of its target, to bother looking at it with care, too busy seeing through it to see it properly. That is a pity. Contemporary analytic philosophy is not without its methodological problems, as I discuss elsewhere (Williamson, 2007, forthcoming). Some of Kitcher's shots are on target. For example, there is genuine insight in his critical comments on recent work on the metaphysics of grounding (69):

> It is entirely reasonable to scrutinize claims that some aspects of the world are "grounded" in others. Equally reasonable to appreciate a kinship with the concept of explanation: it seems, at least initially, that *x* is grounded in *y* just in case we would explain *x* by appealing to *y*. If that equivalence is even in the ballpark of the truth, it might be advisable to reflect on the reasons for which, in a number of alternative traditions, the idea of a complete formal theory of explanation was abandoned. That has implications for those who seek "the logic of *ground*."

One can start to substantiate Kitcher's suspicions by investigating how explanatory helpfulness is used as a heuristic for assessing statements of the form 'P because Q', in the relevant constitutive sense of 'because' (Williamson, 2021, forthcoming). But, to be effective, a methodological critique of current analytic philosophy needs to be far more discriminating than Kitcher's, far more sensitive to what its target is trying to achieve, far more willing to engage with it in detail.

In the final part of this essay, I will discuss some of the more general assumptions at work in the book.

Kitcher aligns himself with the pragmatist tradition of James and Dewey. He has little time for pure intellectual curiosity or love of knowledge as a basis for philosophical inquiry. He asks rhetorically (100):

> Should societies grant blanket permission to people who, in at least some instances, are privileged in the level of support they enjoy, to pursue any venture that arouses their intellectual curiosity, without any responsibility to account for the benefits they take it to deliver?

He goes on to sneer at 'philosophical rhapsodies about the unworried pursuit of projects whose probability of adding to the broader human good is infinitesimal' (ibid.). The book contains many memorable potential quotes that a populist politician could use to justify shutting down philosophy departments, even though Kitcher intends no such consequence (151).

Of course, similar questions can be raised about the utility of most research in mathematics, natural science, and the humanities. Kitcher writes triumphantly of the British mathematician G. H. Hardy (8):

> Hardy, distressed by the outbreak of World War II, gloried in the uselessness of number theory—and did not foresee how cryptography would later apply his field.

Kitcher seems to regard this as a point on his side. But what the case illustrates is that someone can do practically useful research without any intention to do so. If Hardy's research funding had depended on foreseeable practical utility, his research would not have been funded. Indeed, if he had been forced to do research with foreseeable practical benefits, he would very likely have worked on easier problems in applied mathematics, which would not have had those benefits for cryptography. The point is familiar: in the long run, 'blue skies research', done out of intellectual curiosity, with no eye to applications, can have the greatest practical benefits, because the questions it tackles go deeper and wider than more practically-motivated questions.

Kitcher might concede that much, but insist that it has no realistic chance of applying to most work in contemporary analytic philosophy. However, that response ignores two factors. First, very few pieces of research anywhere but in the most applied science have practical applications all by themselves. For almost all research, the practical applications (if any) are mediated in complex ways. Second, when philosophy has practical applications, they were often unpredictable in advance.

One example is the modern study of epistemic logic. It goes back to Jaakko Hintikka's book *Knowledge and Belief* (1962), which provided a formal framework for exploring issues in traditional epistemology, such as whether knowing implies knowing that one knows. Nevertheless, multi-agent epistemic logic has found important applications in both theoretical economics and computer science (Fagin et al., 1995). Consequently, other work in analytic epistemology which bears on principles of epistemic logic may also turn out to be practically relevant.

Another example is mereology, the branch of analytic metaphysics which studies the abstract structure of relations between parts and wholes. Peter Simons, the author of a standard monograph on mereology (Simons, 1987), was also employed part-time by engineering companies such as Lockheed Martin and Rolls Royce as a consultant on ontology (McCarthy, 2006). His work in mereology turned out to be relevant to the classification of engineering components, which for practical reasons needs to be as perspicuous as possible. Consequently, other work in analytic metaphysics which bears on principles of mereology may also turn out to be practically relevant.

As such examples show, predicting which pieces of research will eventually have practical utility is much harder than it looks. It is fortunate that none of Kitcher's

pragmatist predecessors managed to prevent Hintikka from doing his work in analytic epistemology, or Simons from doing his in analytic metaphysics.

One may concede Kitcher this much: something is wrong when a discipline operates in isolation from all others. But contemporary analytic philosophy is clearly not such a discipline. In addition to all the examples already given, one should mention contemporary analytic philosophy of mind—another item on Kitcher's hitlist—much of which is in continuous interaction with various branches of psychology. Of course, when analytic philosophers are talking to each other, the conversation often gets intricately technical, to the bafflement of non-philosophers, as Kitcher repeatedly complains (62, 101, 118, 152). Obviously, the same happens when mathematicians are talking to each other. The difference is that non-mathematicians *expect* not to understand mathematicians when they are talking to each other about mathematics, whereas many non-philosophers still seem to feel entitled to understand philosophers when they are talking to each other about philosophy. But no self-respecting discipline should feel bound by such a constraint. Although it needs an outward-facing side, it also has the right to get as intricately technical as its questions demand.

Contemporary analytic philosophy has plenty to offer other disciplines, as its track record of inter-disciplinary interaction shows. It will have even more to offer other disciplines if it gets better at answering its own questions. There is plenty of room for methodological improvement. Some of what Kitcher says is salutary in that respect, for instance: learn more about modern evolutionary theory before you try to apply it. But most of his recommendations would make philosophy *worse*: more anxious to imitate the neighbours and please the general public, and so with less to contribute of its own. That is not exactly his intention, but as a pragmatist he should stand ready to be judged by his polemic's consequences in practice.

References

Fagin, R., Joseph Halpern, Yoram Moses, and Moshe Vardi. (1995) *Reasoning about knowledge*. Cambridge, Mass.: MIT Press.

Gellner, E. (1959). *Words and things: A critical account of linguistic philosophy and a study in ideology*. London: Gollancz.

Hintikka, J. (1962). *Knowledge and belief: An introduction to the logic of two notions*. Ithaca, N.Y.: Cornell University Press.

McCarthy, N. (2006). 'Philosophy in the making'. *The Royal Academy of Engineering*, Issue 26; https://www.ingenia.org.uk/ingenia/issue-26/philosophy-in-the-making.

Simons, P. (1987). *Parts: A study in ontology*. Oxford: Oxford University Press.

Srinivasan, A. (2020). Radical externalism. *Philosophical Review*, *129*, 395–431.

Williamson, T. (2007). *The philosophy of philosophy*. Wiley-Blackwell. Enlarged (2021). edition.
Williamson, T. (2020). *Suppose and tell: The semantics and heuristics of conditionals*. Oxford: Oxford University Press.
Williamson, T. (2021). 'Degrees of freedom: Is good philosophy bad science? *Disputatio*, *13*(61): 73–94.
Williamson, T. (2023). 'Boghossian, Müller-Lyer, the parrot, and the Nazi', in Luis Oliveira (Ed.), *Externalism about knowledge*. Oxford: Oxford University Press, 377.
Williamson, T. Forthcoming. *Overfitting and heuristics in philosophy*. New York: Oxford University Press.