Book of Abstracts

"New Directions in Metaphilosophy" Conference

13 – 14 May 2021

Department of Philosophy and Centre for Reasoning

University of Kent, UK

Sponsored by the Aristotelian Society and the Analysis Trust

Contents

T-Philosophy (Christopher Daly)

Criticism of metaphilosophy may take one of two forms, general or special. By the general case I have in mind a root and branch rejection of the enterprise of metaphilosophy. The enterprise is rejected as misconceived, pointless or inimical to genuine intellectual inquiry. The general case is usually made as part of a still more general claim: that philosophy, not just the philosophy of philosophy, is itself misconceived, pointless or inimical to good thinking. Some of D. H. Mellor's arguments for the general case are assessed. A special case against metaphilosophy targets a particular conception of metaphilosophy, a particular way of thinking of philosophy. The later Wittgenstein sought to displace what he saw as the prevailing conception of philosophy as a theoretical fact-uncovering enterprise in favour of his conception of philosophy as a form of therapy that frees us from the perplexities induced by a muddled understanding of language. Wittgenstein's case has recently been taken up and refurbished by Paul Horwich. Horwich's arguments are evaluated.

The Epistemic Aim of Philosophy (Michael Hannon & James Nguyen)

The discipline of philosophy is puzzling in a number of ways. Unlike many other disciplines, philosophy circles back to the same questions, even over millennia. This indicates that philosophy is not a reliable way to arrive at philosophical truths, and that philosophers are not reliable sources of knowledge about which philosophical answers are the right ones. Moreover, academic appointments in philosophy often ignore "the advantage of being right" (Lewis 1999). When deciding whom to appoint, a job search committee will typically behave as if the truth or falsity of a candidate's philosophical views is not a legitimate consideration. We also don't seem to care whether our students provide true answers to philosophical questions. After all, we uphold the idea that two classmates can provide utterly opposed answers to a philosophical question and yet each may get an A+ (or a First) for their work. Likewise, we are untroubled by the fact that our colleagues often impart (what we take to be) false doctrines on our students.

These puzzling aspects of philosophy raise a number of additional questions, such as:

- Is perennial disagreement in philosophy consistent with success or progress?

- How can we recognize philosophical expertise in the face of systematic peer disagreement?

- Why do philosophers dedicate their lives to philosophy even when they know that it circles back to the same questions that preoccupied Plato and Aristotle?

- Why is it often inappropriate to form philosophical beliefs on the basis of testimony?

- Why do some philosophers argue for conclusions that almost everybody already accepts?

In this paper, our goal is to provide a unified answer to these questions (and several more). We demonstrate that new light can be cast on all these issues by first answering a more

fundamental metaphilosophical question, namely: What is the aim of philosophy as an intellectual activity? The answer, we believe, is: *understanding*.

Our argument proceeds as follows. We will first outline the various core features of understanding, drawing on recent literature in the epistemology of understanding. Second, we argue that the primary intellectual aim of philosophy is understanding. This hypothesis has far more explanatory power then its main rival: the view that the primary goal of philosophy is the *attainment of truth*. The idea that philosophy aim's at understanding is used to argue for the following claims:

- Philosophy makes considerable progress;

- There are philosophical experts despite widespread peer disagreement;

- Skepticism about the reliability of philosophy is misguided;

- It is often unacceptable to form philosophical beliefs on the basis of deference;

Overall, we demonstrate that that many aspects of philosophical practice become intelligible on the assumption that philosophy aims at understanding (rather than truth or knowledge). In this way, our hypothesis has more explanatory power than its competitors.

Philosophy in Interdisciplinary Collaboration (Stacie Friend)

A standard question in discussions of philosophical methodology is whether philosophers should get out of their 'armchairs' and engage in empirical research of a more scientific sort. Advocates of the latter often contribute to the burgeoning field of experimental philosophy, creating their own empirical studies to test (and challenge) standard philosophical methods or to shed new light on philosophical problems. Defenders of more traditional approaches often point out that philosophers possess training and expertise that participants in these experiments typically lack. In this paper I suggest that interdisciplinary collaboration between philosophers and empirical scientists offers a more promising approach to at least some philosophical questions. At the same time, I discuss reasons why the promise is difficult to fulfil.

Extending Philosophical Thought through Poetry (Karen Simecek)

What makes new philosophical thought possible? What role might literature, and in particular, poetry play in the development of philosophical thought? The idea that poetry might have something to offer the philosopher is not new; there is a long history of experimenting with philosophy through poetry, take for instance, Lucretius' *De Rerum Natura* (The Nature of Things) or Pope's *An Essay on Man*. However, in the reception of such works there has been a tendency to treat such works as either poetry or philosophy, rather than taking seriously philosophy through poetry as distinct from (prose)philosophy. In

this paper, I will argue that engaging with poetry – its particular use of language and the perspectival nature of its structure – allows one to experience extension of thought and possibilities of language, thus laying the foundation for new philosophical thought. Reading poetry encourages the reader to experience the limits of their own perspectives, which makes new insights possible. My argument centres on the claim that some philosophical thought is perspectival, and to entertain such thought one must engage with it actively and first-personally.

Hypothesis-Driven Digital Philosophy of Science (Charles Pence)

Recent trends toward both experimental philosophy and digital philosophy seem to have begun collapsing the distinction between scientific and philosophical methodology. In at least some cases, philosophers are turning toward empirical methods, in the process having to consider how those methods should be put in dialogue with traditional philosophical questions and approaches. In this talk, I want to pick up on an element of this methodological discussion in the context of digital philosophy of science.

Digital humanities (and, by extension, digital philosophy) is, at its heart, a "big data" discipline. It is not uncommon to perform analyses on dozens or hundreds of books, or thousands or tens of thousands of journal articles, in an effort to extract generalizations about scientific practice that can, in turn, be used to ground philosophical claims (Lean, Pence, and Rivelli forthcoming). This poses a variety of interesting problems, already familiar to any well-versed data scientist. Analyses performed on any dataset of such a size will offer significant problems of noise. A number of spurious correlations are certain to be present (for a humorous visualization of the problem, see https://tylervigen.com/spurious-correlations). Temptations to simply "read off" philosophical conclusions from the data themselves should be resisted; any philosophical conclusion of merit will require significant human interpretation (boyd and Crawford 2012).

One common proactive way to approach these problems and others like them is to turn toward hypothesis-driven research. Specific, testable, empirical hypotheses about the content of the literature can be evaluated with the aid of digital tools, if we clearly know what it would mean for them to succeed or fail. What is less clear, I think, are the answers two further questions. First, how does the evidence of the success or failure of such a hypothesis lead to confirmation or disconfirmation of generalizations about scientific practice, which are in turn the ingredients we hope to use to build novel claims in the philosophy of science? And second, what are the disadvantages of such a method? How can we balance the utility of these digital approaches as tools for discovery of unexpected trends in the sciences with the risk that we will simply see the connections that we want to see?

For all that these are widely acknowledged questions, the ways in which they might bear on digital philosophy of science are less clear. I hope in this talk to take some first steps toward an analysis of these issues, with the goal of offering insight for philosophers of science that is both theoretically grounded (in an understanding of the utility of empirical generalizations about scientific practice for the philosophy of science) and practically useful (in that it can give philosophers advice about how to formulate and use such hypotheses).

Imagination and the Epistemology of Interventionism (Amanda Bryant)

Defenders of causal interventionism (Cartwright 2003; Meek and Glymour 1994; Pearl 2009; Woodward 2000, 2003) face foundational problems in the articulation and defence of their view, which have been well-explored in metaphysics and philosophy of science, including the challenge of defining interventionism in a non-circular manner that does not invoke causal notions, as well as the implausibility of understanding an objective metaphysical relation in terms of human agency. While the basic coherence and plausibility of the view have been thoroughly discussed, metaphilosophical matters concerning its accompanying methodology and epistemology have so far remained largely unexamined. In this paper, I will argue that the methodology and epistemology that naturally accompany causal interventionism further detract from its appeal as a philosophical framework.

The issue takes on additional import and interest given recent connections drawn between causation and metaphysical grounding (Schaffer 2016, A. Wilson 2018). Some metaphysicians have even taken up interventionist language in their discussion of grounding, such as Jonathan Schaffer, who claims that "there is a straightforward and informative parallel working test of token grounding to be had, in terms of counterfactual covariation: wiggle the ground, and the grounded wiggles" (2016, §3.2). Here we have an epistemological analog of causal interventionism. If causation and grounding are strongly metaphysically analogous, their respective epistemologies will likewise dovetail, and exploring the epistemology that falls out of various causal theories can potentially illuminate, and perhaps spell trouble for, the epistemology and methodology of the metaphysics of grounding.

If true, causal interventionism suggests that in order to know about causal relationships, we must somehow discover the relevant counterfactuals — for example, that if we were to bring about A via some intervention, B would occur. That is, we must discern what would happen were we to 'wiggle' certain putative causes. We must ask ourselves on what basis philosophers form beliefs about such counterfactuals and, importantly, whether it constitutes an epistemically optimal — or even merely permissible — basis for belief.

Imagination is one likely source of our claims about the relevant counterfactuals. Timothy Williamson has argued that we have a developmentally and evolutionarily un-mysterious cognitive capacity for handling counterfactual conditionals, which often employs the imagination, "radically informed and disciplined" by the empirical background of beliefs and an accompanying folk physics (2007, 143). For instance, imaginative simulation allows us to discern that "If the bush had not been there, the rock would have ended in the lake" (2007, 142).

If Williamson's epistemology of counterfactuals is correct, so much the worse for causal interventionism and its philosophical cousins. Empirical experience and folk physics prepare us well to assess ordinary counterfactuals regarding things like rock trajectories — i.e. in situations where folk physics is adequate for our predictive purposes. Yet those are highly limited adequacy conditions. If interventionists believe we can have knowledge of counterfactuals pertaining to, for instance, unusual physical circumstances, non-empirical matters, interventions beyond the realm of practical possibility or beyond our capacity to

scientifically model, then they must either defend the epistemic permissibility of stretching our imaginative resources beyond their adequacy conditions or point to some alternative mode of epistemic access. I believe that in cases where empirical experience doesn't enable imaginative simulations, imagination calls on antecedent intuitions that lack adequate epistemic footing and evidential weight. If so, then causal interventionism and its methodological analogs invoke an epistemically deficient form of modal rationalism — or so I will argue.

Impossible Worlds and the Safety of Philosophical Beliefs (Zack Garrett and Zach Wrublewski)

Epistemological accounts that make use of a safety condition on knowledge, historically, face serious problems regarding beliefs about necessary truths. This is because necessary truths are true in all possible worlds, so such beliefs are trivially safe. The existence of trivially safe beliefs would undermine a major motivation for the condition itself: the ability to evaluate how well a belief tracks the truth. We call this problem the "Triviality Problem." In a recent paper, Guido Melchior suggests that this problem may be ameliorated by allowing the close worlds considered in the definition of safety to include impossible worlds. Since beliefs that are necessarily true in the actual world can be false in impossible worlds, these beliefs would not be trivially safe. But, for this to work, *many* impossible worlds would have to be closer to the actual world than *many* possible worlds. Melchior argues that this seems implausible, and, thus, we should accept a similar version of sensitivity rather than safety (as sensitivity doesn't require closeness, whereas safety does).

In this paper, we'll show that according to various accounts of the similarity of worlds, many impossible worlds are incredibly similar to the actual world. So, *contra* Melchior, we contend that there is good reason to accept that many impossible worlds are closer to the actual world than many possible worlds. As such, we need not fall back to sensitivity to deal with the Triviality Problem. Further, we'll argue that including impossible worlds in this way would mean that we can properly evaluate the safety of many philosophical beliefs— in particular, philosophical beliefs that, if true, would be necessarily true.

For many kinds of philosophical beliefs, *e.g.* , beliefs about the correct logic, the correct ethical theory, the impossibility of phenomenal zombies, *etc.* , we'll argue that there are sufficiently many close impossible worlds to render these beliefs unsafe. So, we contend that many kinds of philosophical beliefs are unsafe. As such, we will be agreeing with philosophical skeptics, like Helen Beebee, that many of our philosophical beliefs do not amount to knowledge.

But, we'll resist the conclusion that philosophy cannot make progress toward knowledge. We'll argue that the inclusion of impossible worlds in this discussion can help to explain what would need to change about specific philosophical methodologies in order to obtain safe beliefs; and, even if we cannot get safe beliefs through these methodologies, we can make progress toward *safer* beliefs.

Humanistic Philosophy as Interpretative (Alan Millar)

The title derives from Bernard Williams' essay, 'Philosophy as a Humanistic Discipline', in which Williams speaks of philosophy as 'part of a more general attempt to make the best sense of our life, and so of our intellectual activities, in the situation in which we find ourselves'. He contrasts this conception with a scientism that assimilates philosophy to the natural sciences. The present discussion defends the idea that an important motivation to philosophising is a concern with human nature, and in particular, with rational thought and action. It works towards the idea that humanistic philosophy is in large measure interpretative. In its attempts to get a grip on topics such as knowledge, inquiry, belief, intention, desire, emotion, abilities, reasons, responsibility, personal identity, values, rights, and much else besides it should have due regard to 'ordinary' thought as well as existing philosophical thought about these things, and so must engage with conceptual issues. It should take into account what we already know, not just what strikes us as intuitive. If we are investigating intention, for instance, we should try to shed light on what intention must be given what is true in what people say about intention. A preoccupation with concepts need not confine itself to clarifying, far less analysing, concepts, as opposed to illuminating what they can be used to describe or express. In pursuing these themes some aspects of the history of philosophy are briefly considered.

Conceptual Engineering as Lexical and Conceptual Abandonment: "Democracy" as a Case Study (Herman Cappelen)

Some of our terminology shouldn't be improved, but instead abandoned. Abandonment theory is the study of the conditions under which lexical and conceptual abandonment is appropriate. The first part of this talk is an introduction to abandonment theory and its relationship to amelioration, replacement and elimination. The second part applies abandonment theory to a core concept in political philosophy: 'democracy'. I argue that 'democracy' is an ideal candidate for abandonment.

Attentional Progress by Conceptual Engineering (Eve Kitsik)

Conceptual engineering has been getting a good deal of attention recently: many voices call for evaluating and improving our concepts rather than just analysing them. But some feel that the attention is largely undeserved: the important philosophical questions are about the world, not about what our concepts are or should be. Such critics of the conceptual engineering bandwagon apparently accept that it is important to distribute attention properly; for example, conceptual engineering should not get more than its fair share. But conceptual engineering, I argue, is precisely an instrument for distributing attention well. As such, conceptual engineering can serve attentional progress in philosophy and in other areas of life and inquiry. I focus on attentional progress by conceptual engineering in philosophy.

According to the relevant notion of attention, to attend to X is to persistently select X for action, including mental action. (This draws on Alan Allport's and Wayne Wu's accounts of attention as selection for action.) Proper distribution of attention means persistently selecting the right objects for the right actions. Conceptual engineering can help achieve such proper distribution of attention in philosophy. At a very general level, this can happen when philosophers purposefully shape labels like "metaphysics" and "epistemology", or even "philosophy" itself, to legitimize non-mainstream and under-appreciated areas of inquiry. But there are also more local ways of engineering concepts to engineer attention. For example, we can facilitate selecting specific phenomena for discussion by introducing new terms (like "analyticity" or "supervenience"). Further, we can reassign terms to facilitate selecting a more significant phenomenon in the vicinity of what was originally associated with the term. For example, perhaps epistemologists should think more about knowledge* rather than knowledge; and a way of facilitating that is reassigning the term "knowledge". (There is, of course, room for disagreement on whether reassigning terms is the best way to achieve the desired shift of attention, in any given case.)

Recognizing attentional progress in philosophy stands in some contrast with the common idea (associated, among others, with David Chalmers) that philosophical progress consists in establishing the answers to the big philosophical questions. When philosophers make attentional progress, whether by conceptual engineering or otherwise, they often replace their questions with better ones, as opposed to answering the questions. Granted, such replacement cannot plausibly be all the progress there is in philosophy. After all, improving the goals is worth little if one does not at all approach achieving the goals. But it is nevertheless an important form of progress.

The bottom line is: if we care about what gets attention in philosophy, then we should care about conceptual engineering as a method that facilitates progress on that front.

One Person's Modus Ponens is Another's Modus Tollens: Pantomemes and Nisowir (Jon Williamson)

That one person's Modus Ponens is another's Modus Tollens is the bane of philosophy, I argue, because it strips many philosophical arguments of their persuasive force. I show that appeals to intuition, evidence or truth fail to alleviate the problem. However, I develop two strategies that do help in certain circumstances: an appeal to *normal informal standards of what is reasonable*, and *argument by interpretation*. The method of explication features prominently in both strategies. I illustrate the problem and the two strategies with examples of arguments in formal epistemology, and suggest that one of the strategies can help to defend against philosophical scepticism by shifting the burden of proof to the sceptic.

In Praise of (Logical) Practice (Benjamin Martin)

When it comes to understanding the epistemology of logic, current debates in the literature are in as bad of a condition as discussions of the methodology of the sciences were in the 1950s. Rather than demonstrating a recognition of the realities of research in the field, accounts of logic's methodology are more concerned with formulating a proposal which fits with the advocate's wider metaphysical and epistemological commitments, while respecting some of the traditional properties of logic. What has resulted from these theory-first, "topdown" approaches, which start from assumptions about knowledge and logic, are accounts of logic's methodology which struggle to make any serious sense of the actual debates logicians engage in (Martin 2020; Martin & Hjortland 2020).

A case in point is the traditional picture of logical knowledge resulting from analyticity, which still has some advocates (Warren 2020). According to the conventionalist tradition, made famous by the logical positivists (Ayer 1936; Carnap 1937), as the validity of logical inferences is wholly a matter of linguistic convention, in order to become justified in believing that an inference is valid all we need to do is appropriately understand its content. The perceived philosophical strength of the conventionalist's account of logical knowledge is that it manages to simultaneously respect certain traditional philosophical assumptions about logic, such as its evidence being wholly *apriori* and its laws being necessarily true, while retaining a commitment to metaphysical naturalism (Warren 2020). This was, as we know, one of the significant attractions of the account for the positivists (Carnap 1963).

The problem with the proposal, however, is it significantly distorts the disputes that logicians get into, and the actual business of theory-choice and justification in logic. If the traditional picture of logical justification in terms of epistemic analyticity were correct then we would expect logicians, when engaged in debates over the validity of a rule of inference, to directly appeal to the meaning of the terms involved in the inference. Yet, this is far from what we find. Instead, a far more varied range of evidence are appealed to when arguing for particular

logics, including their ability to explain the validity of steps within mathematical proofs and solve problems blighting competitors, such as the logico-semantic paradoxes (Martin 2020; Martin & Hjortland 2020). This weakness of current accounts of logic's methodology holds just as true for other prominent proposals, including forms of rationalism (BonJour 1998) and the currently favoured abductivism (Priest 2014; Williamson 2017).

This talk will argue for the need to embrace a new *practice-based approach* to the epistemology of logic if we are to rectify the failures of past accounts. According to this *approach*, we should begin by looking in detail at the actual practice of logicians and then extract methodological principles from this practice, gradually building up a detailed account of logic's epistemology. The working assumption underpinning such an approach being that experts in a field are generally very reliable at what constitutes relevant evidence for or against a theory in that field. Drawing on both the failures of past accounts and lessons learnt from the philosophy of sciences, we'll argue for the advantages of the approach, the new avenues of research it opens up, and why traditional concerns over such an approach are mistaken.

Philosophy Doesn't Need the Concept of Progress (Yafeng Shan)

Philosophical progress is one of the most controversial topics in metaphilosophy. It has been widely debated on whether philosophy makes any progress in history. In this paper I revisit the concept of philosophical progress. First, I identify two criteria of an ideal concept of philosophical progress. Then I argue that our accounts of philosophical progress fail to provide such an ideal concept. Finally, I argue that not only do we not have a good concept of philosophical progress, but also we do not need a concept of philosophical progress in order to make a good understanding of the history of philosophy.

What is Naturalized Metaphysics and What Could It Be? (Jack Ritchie)

Metaphysical theorizing is usually thought of as an attempt to say how things are. Science too, at least if we are realists, is an attempt to describe the world. What is the right way to understand the connection between these two descriptive activities? Are they in competition? Most contemporary philosophers think not. Science and metaphysics ought to and do, when done well, complement one another: metaphysical theorizing is continuous with scientific investigation. Sometimes this view is called naturalized metaphysics and there are, I think, three broad (possibly overlapping) ways in which contemporary metaphysicians think of their discipline as continuous with the natural sciences. According to some, although metaphysicians are concerned with more general questions than the sciences -- what are things, properties, laws, fundamental structure, etc. -- the methods which are used to answer metaphysical questions are the same as those used by the natural sciences. The best theory is the one which fits all the data and displays various theoretical virtues like simplicity, fecundity and others. A second kind of naturalized metaphysics thinks of the metaphysical project as an attempt to synthesize the diverse theories and ideas of the natural sciences. One way to do this is to claim that there is a general metaphysical picture that is best supported by our scientific knowledge. It is the philosopher's job both to articulate this general theory and show how (if they do) non-fundamental aspects of reality fit into this picture. Physicalism is a version of this form of naturalized metaphysics and so too is Ladyman and Ross's version of ontic structural realism. Finally, in the foundations of physics community there is the idea that the project of interpreting our best fundamental theories is a way, perhaps the only honestly naturalistic way, to pursue metaphysics. I argue that a reflective naturalist ought to reject all three ways of doing naturalized metaphysics. The methods used by metaphysicians

debating the nature of properties are not the same as scientific methods; over-arching views like physicalism or ontic structural realism either lack clear content or when formulated with more precision have no empirical support from current science; debates in the foundations of physics show us either, as in the case of non-relativistic quantum mechanics, there are too may reasonable interpretations and no good empirical reasons to choose among them or, as in the case of quantum field theory, that there is no consistent interpretation. The lesson for the naturalist is that we ought to think of metaphysics as a different sort of game to science. The aim is or ought not to be to articulate truths about how the world is but to be, for lack of a better word, interesting. I suggest three ways of being interesting connected to the three kinds of naturalized metaphysics mentions above: metaphysics as the pursuit of multiple interpretations with aim of better understanding our theories and their applications, metaphysics as model building without a target and metaphysics as a species of metaphorical talk.

On Philosophical Model-Building (Miguel Egler)

What should we make of the fact that there is widespread disagreement among philosophers on virtually every philosophical issue? For the most part, answers to this question paint a rather bleak picture of the philosophical discipline. For instance, some have argued that empirical findings showing a great level of disparity in people's intuitions undermine our well-established approaches to philosophical inquiry. Moreover, others have worried that the widespread disagreement we find in the current philosophical landscape invites a pernicious kind of meta-philosophical scepticism that makes it imprudent to hold any philosophical view whatsoever . And to make matters worse, these concerns are not mutually exclusive. When taken together they motivate the conclusion that all philosophical research should come to a screeching halt. Call this the 'problem of philosophical paralysis'.

Not all philosophers share this grim outlook on the philosophical discipline. One notable example are proponents of the view that philosophical inquiry can (at least sometimes) involve practices of *model-building* . Call this view `Phil-mod'. Defenders of Phil-mod argue that this view gives us two reasons to be more optimistic about the future of the discipline. First, they claim that by acknowledging the role of model-building in philosophy, we can recognise that the widespread disagreement we find in philosophy is not as disconcerting as some would have us believe. And second, they suggest that Phil-mod offers valuable resources with which to better measure philosophical progress: namely, in terms of the development of successful models in philosophy.

Despite these promises, current formulations of Phil-mod remain cursory at best. So far, these accounts have failed to explain how we should deal with conflicts among philosophical models, or why any such conflicts fail to motivate the problem of philosophical paralysis. Moreover, these accounts are scarce in details regarding what *epistemic* achievement is involved in the development of successful philosophical models. In this paper, I seek to remedy these shortcomings.

For the sake of argument, I will assume Phil-mod in this paper. My central purpose will be to develop this view in order to show how it can help push back against the problem of

philosophical paralysis. I first rehearse the motivations for the most prominent formulations of Phil-mod and highlight how they fail to suppress concerns arising from the fact of widespread disagreement in philosophy. I then explain how conflicts among models in biology fail to motivate any epistemological concerns about research in this field of study. I then trace parallels between such examples and purported cases of model-building in philosophy to show how similar conclusions apply there. Lastly, I flesh out this picture in terms of the claim that, like scientific models, philosophical models can be said to afford *understanding* of phenomena.

Disagreement in Metaphysics (Timothy Williamson)

Disagreement in metaphysics is often seen as pathological, a symptom of something deeply wrong with the whole field. Sometimes it is treated as 'merely verbal', between theories which disagree over nothing 'substantive' and may even be mutually equivalent. On one currently fashionable view, much disagreement in metaphysics is 'metalinguistic negotiation' in disguise. Even more negatively, disagreement in metaphysics is diagnosed as a product of 'conceptual confusion'. In opposition to such ideas, I will argue that disagreement in metaphysics is typically just what it seems, theoretical disagreement on very basic, general matters about the nature and structure of reality. Given the kind of hard questions metaphysicians ask, general agreement on their answers would be more disturbing than disagreement.

Some deflationists about metaphysics invoke principles of charity in interpretation. However, they misapply those principles. In effect, they treat dialogues between metaphysicians as several monologues, to be interpreted independently of each other; the most charitable interpretation of the dialogue need not be the sum of the most charitable interpretations of the monologues, taken in isolation from each other. Inappropriate hermeneutic devices can trivialize speakers' contributions or make them conversationally irrelevant. No sound principle of charity works like that. Interpretations in terms of metalinguistic negotiation are often problematic in similar ways.

Critiques of metaphysics which appeal to the prevalence of disagreement risk self-defeat. For they are contributions to metametaphysics, where there is just as much disagreement as in metaphysics itself.

To a striking extent, critiques of metaphysics still rely on some kind of logical empiricism, with an implicit or explicit taxonomy of conceptual analysis and empirical inquiry as the only two possibilities for respectable cognition. Such a view is both epistemologically and psychologically naïve. To oppose it in more detail, I will consider the role of disagreement in *logic*, as in disputes between classical and alternative logics, since logic is generally agreed to be a form of respectable cognition of use in both metaphysics and metametaphysics. Although attempts are often made to deflate disputes in logic, usually by interpreting them as merely verbal, such attempts tend to be hermeneutically inadequate. For example, they ignore (i) technical results about the impossibility in some cases of 'peaceful co-existence' between two logics with different logical constants in the same language; (ii) the application of social externalism about meaning to logical constants, and (iii) the extension of many logical disputes to the metalogic. In practice, the standard methodology for such logical disputes is abductive, with each side appealing to the alleged theoretical virtues of its preferred logic.

Even in principle, no better methodology for these disputes is available. Such argumentation fits neither logical empiricists' stereotype of conceptual analysis nor their stereotype of empirical inquiry. Although one could Procrusteanly assimilate it to one side of the dichotomy or the other, or to a combination of both, the same could then be done for paradigmatic metaphysical argumentation. Indeed, given the metaphysical motivations in play, many disagreements in logic are themselves good exemplars of disagreement in metaphysics as normal theoretical disagreement.