


Concluding Meditation

University Press Scholarship Online

Oxford Scholarship Online



Being, Freedom, and Method: Themes from the Philosophy of Peter van Inwagen

John A. Keller

Print publication date: 2017

Print ISBN-13: 9780198715702

Published to Oxford Scholarship Online: January 2017

DOI: 10.1093/acprof:oso/9780198715702.001.0001

Concluding Meditation

Peter van Inwagen

DOI:10.1093/acprof:oso/9780198715702.003.0019

Abstract and Keywords

In this final chapter, Peter van Inwagen responds to the essays of Louise Antony, David Chalmers, John Keller, Thomas Kelly and Sarah McGrath, Michael Loux, Laurie Paul, and Alex Rosenberg. These responses clarify (and in some cases modify) van Inwagen's views, and give a nice indication of where the next rounds of debate will be conducted on the problem of evil, metaphilosophy, constituent ontology, and the compatibility of theism and evolution. Van Inwagen's responses also provide helpful methodological insight into his approach to philosophy in general, and to ontology, the philosophy of religion, and philosophical success and failure in particular.

Keywords: problem of evil, metaphilosophy, constituent ontology, theism and evolution, method

Concluding Meditation

Introduction

There can be no greater honor for a scholar than to be the recipient of a Festschrift. My gratitude to John Keller and the (other) contributors is so great that, in attempting to write this opening paragraph, I discover that I am able to find no words that are adequate to express it. I will only say, thank you all. Thank you all so very much.

It is perhaps unusual for a Festschrift to contain responses by the recipient to the essays it contains. (A Festschrift, after all, is not a “Schilpp volume.”) Apparently, however, the editors at Oxford thought that such replies would be advisable. In any case, the purpose of this preface is to say a few words about my replies.

I have not responded to all the essays. The reason is twofold yet simple: there was not enough time and there would not have been enough space. I am rarely able to write a short, simple reply to or comment on a philosophical essay. I may be able to do this if I agree almost entirely with the author’s arguments and need only make a few minor points. This is the case with my reply to Michael Loux’s essay. Or I may be able to do this if the essay employs a system of philosophical concepts that is radically different from my own—in which case I can often find something brief but helpful to say about the nature of the gulf that separates the concepts the author employs from the concepts I employ. This is the case with my reply to Laurie Paul’s essay. It even occasionally happens that I agree entirely with an author—that the author’s essay is one that, matters of style apart, I might have written myself. This is the case with Eric Olson’s essay. In such a case, no reply is needed beyond what Alvin Plantinga once called “a brief celebratory ceremony.” I will accordingly enact a suitable brief celebratory ceremony: I hereby declare that Eric’s essay is a fine piece of work.

But cases like these are the exceptions. If I am commenting on or replying to an essay that is addressed either to my own work or to some problem or topic that deeply interests me, I almost always have a very hard time keeping my responses

Concluding Meditation

and comments within reasonable bounds. However much I say, there always seems to be more.

For this reason, if I were to attempt to reply to all the essays this volume contains, it would almost certainly be published only after my death—and with some of the responses still unwritten. And it would probably not be a volume. It would probably comprise at least two volumes—assuming that anyone was willing to publish such a monster.

(p.344) When I began to write replies to the essays in this book, I started with the ones with whose authors I am in most fundamental disagreement—but authors whose ideology (in Quine’s sense) and ontology are sufficiently similar to my own that a useful and extended discussion of the matters on which we disagree is possible. Accordingly, the first replies I wrote were to the essays of Louise Antony and Alex Rosenberg. No one, I think, would dispute the statement that their essays represent points of view that are fundamentally opposed to my own. But Louise’s and Alex’s points of view are not “fundamentally opposed” to my own point of view in the way Laurie Paul’s is. It’s evident that Louise and Alex’s opinions concerning God, Freedom, and Immortality (or to use Douglas Adams’s tripartite list, Life, the Universe, and Everything) are radically different from mine. Nevertheless, these disagreements, radical though they are, can be framed using a common vocabulary and a common system of metaphysical concepts—and this is not so in the case of my radical disagreements with Laurie.

My work on the replies to Louise and Alex’s essays illustrated the basic soundness of Hofstadter’s Law: Any project always takes longer than you expect, even when you take Hofstadter’s Law into account. I next wrote the replies to the essays of Michael Loux and Laurie Paul, which went more quickly (although, of course, less quickly than I expected), and then went on to the essays on “Method”—a topic that is of deep and abiding interest to me. (The essays, that is, of David Chalmers, John Keller, and Thomas Kelly and Sarah McGrath.) When I had finally finished these three replies, I had written replies to only seven of the seventeen essays this volume contains (counting the “Symposium on the Fixity of the Past” as two

Concluding Meditation

essays), and had written almost 30,000 words. At that point, I was simply out of time and out of space.

I am deeply saddened by the fact that I have not written replies to many essays that are by old students and old friends of mine—in several cases, both. (Sara Bernstein is too young to be an old friend, but she is a very good friend.) I am deeply saddened by the fact that I have not written replies to many excellent essays—ten of them, to be exact, for every essay in the volume is an excellent essay. In only one case—the case of Eric Olson’s essay, for the reasons I have given—was the failure to reply to an essay a matter of deliberate choice. Setting that case aside, that I wrote replies to the seven essays I did write replies to and not to the other equally deserving nine essays is to a very large degree a matter of chance: an unintended consequence of the point at which I chose to begin, my inability to write short replies, and the amount of time and space available to me. In most of the possible worlds closest to actuality, I wrote replies to the members of some other proper subset of the seventeen essays. I will close by thanking those contributors to whose essays I did not write a reply—Eric Olson, Sara Bernstein, Mark Heller, Alicia Finch, Neal Tognazzini, John Martin Fischer, Wes Holliday, Eleonore Stump, Frances Howard-Snyder, Daniel Howard-Snyder, and Lynne Rudder Baker—for their contributions, which I read with pleasure and gratitude. All but two of these eleven philosophers are close friends of mine, and five of them are former students. I wish I had been able to do more. I hope my friends and students will forgive me.

17.1 Reply to Michael J. Loux

I turn now to Michael Loux’s rich and rewarding paper. I can think of no paper that provides a better overview of what is going in ancient, medieval, recent, and present-day (p.345) discussions of the problem of universals—or, better, discussions in which the concept “universal” plays a central role.¹ In these brief remarks, I will touch on only two things that Loux says. Both these things are statements about my own work, but, given the nature of this book, that is perhaps excusable. The first of these statements concerns a very minor matter. The second, however, is of central metaphysical

Concluding Meditation

importance. Both statements occur in footnotes. The first is in note 1. In this note, Loux expresses concern about my

...use of the framework of sets or classes in the definition of a category. Since sets or classes have their memberships essentially, I do not see how van Inwagen's definition can avoid the conclusion that a category in one possible world is different from a category in another world merely in virtue of their having different extensions in those worlds. The two could be composed exclusively of ontologically indistinguishable objects. Since kinds (like properties and relations) do not have their extensions essentially, I would have employed the notion of a kind in defining the concept of a category.

On this matter, I must refer the reader to note 6 to "What is an Ontological Category?",² which reads in part:

The "classes" that figure in this essay are—or are if they really exist—much more like biological taxa than they are like sets....Like taxa, and unlike sets, they can change their membership with the passage of time and the membership of a class in one possible world may not even overlap its membership in another. Like taxa, and unlike sets, moreover, they may have "borderline members"....I am not, however, seriously asserting that there really are things that have the properties I have ascribed to classes. I issue this promissory note: I could—the result would be rather awkward, I concede—eliminate the apparent reference to and quantification over classes in the sequel by paraphrase.

Here are two related points. First, my official ontology does not contain sets or classes. "In the ontology room," I eliminate (or at least I have committed myself to being able to eliminate) by paraphrase all reference to and quantification over sets in favor of quantification over properties. Secondly, my ontology contains "kinds" only insofar as kinds can be identified with properties—the properties that, as it were, correspond to them. Thus, for example, I would identify the kind "horse" (supposing horses to constitute a kind) with the property equinity or horsehood or "being a horse." And if, in

Concluding Meditation

the course of a philosophical discussion, I am told that properties have the wrong, well, properties to be kinds, I will simply make the philosopher who (p.346) has advanced that thesis a present of the word 'kind': I will stop using it and frame my theses by using words like 'property' and 'attribute'.

The second point pertains to Loux's note 16. It is a very long note indeed, and I will quote only a part of it. (The reader should study the whole note carefully before proceeding.)

Van Inwagen...denies that it counts as a substantive explanation of the fact that a concrete particular is, say, green to claim that it exemplifies the color green. Whether he is right or not, most of those involved in the debate would take such a claim to involve a genuine explanation; and despite his denial here, van Inwagen himself takes character to be grounded in what he calls properties, and his properties include what have traditionally been called universals. Thus, van Inwagen... tells us that, in our prephilosophical moments, we all believe that there are anatomical features that insects and spiders share, and he argues that there is (likely) no way of understanding how that belief could be true without conceding the existence of shared properties. But surely this is to claim that the universals in question ground one or more facts about shared character...For van Inwagen...properties are...what he calls unsaturated assertibles, and, on his view it is because they share the relevant assertibles that spiders and insects agree anatomically.

I certainly agree that "most of those involved in the debate would take such a claim to involve a genuine explanation." They are, I say, *wrong, wrong, wrong*.³ At any rate, they are wrong if they take properties or universals to be anything like what I take them to be: unsaturated assertibles. And I'm not going to argue about whether unsaturated assertibles *should* be called properties or universals. Once more I find myself in a generous mood: I'm happy to make a present of the words 'property' and 'universal' to anyone who has strong feelings about their proper use and who thinks that I am using them

Concluding Meditation

improperly. (“But who sees not that all the dispute is about a word?”)

Let’s look at what an unsaturated assertible is. As a saturated assertible (a proposition) is something one can say—and something such that if one says it what one thereby says is true or false⁴—an unsaturated assertible is something one can say *of* something. (And if one says an unsaturated assertible of something, it will either be true of that thing or false of that thing.⁵) Suppose, for example, that I ask you what color this apple is, and you reply, “It’s green.” And suppose Alice asks Jerry what color his copy of *A Theory of Justice* is, and he replies, “It’s green.” Then there is one thing, one unsaturated assertible, such that you said it of the apple and Jerry said it of his copy of *A Theory of Justice*. Call it “that it is green.” (“That it is green” is supposed to be a proper noun—or noun-phrase. If you want words that sound more proper-noun-phraseish, call it ‘what one says of something when one says of it that it is green’. Or if you want a perfect nominal, call it ‘the thing one says of something when one says that it is green’.) Now it happens that what one says of Jerry’s book (p.347) when one says that it is green is true of that book. But would anyone assent to the following statement?

It counts as a substantive explanation of the fact that Jerry’s copy of *A Theory of Justice* is green to claim that the unsaturated assertible “that it is green” is true of it.

Or this one?

It is because the unsaturated assertibles “that it is bilaterally symmetrical” and “that it has a segmented body” and “that it has an exoskeleton” are true both of every spider and of every insect that spiders and insects agree anatomically.

I certainly hope not. If anything, the order of explanation is the other way round:

That Jerry’s copy of *A Theory of Justice* is green counts as a substantive explanation of the fact that the unsaturated assertible “that it is green” is true of it.

Concluding Meditation

It is because spiders and insects agree anatomically (in that every member of the order Araneae and every member of the class Insecta is bilaterally symmetrical, has a segmented body, and has an exoskeleton) that the unsaturated assertibles “that it is bilaterally symmetrical” and “that it has a segmented body” and “that it has an exoskeleton” are true both of every spider and of every insect.

I refer anyone who disagrees to The Philosopher: It is because you are pale that we who say that you are pale have the truth (*Metaphysics* IX 10).

We should remember that unsaturated assertibles, like numbers, are abstract objects and can figure in explanations only in ways analogous to the ways in which numbers can figure in explanations. (If every member of the US Supreme Court votes on an issue, the vote cannot result in a tie because the Court has an odd number of members; if you throw a pair of dice, you have one chance in 18 of throwing 11 because there are 36 ways the dice can fall and only two of them add up to 11 and $2 \text{ divided by } 36 \text{ is equal to } 1 \text{ divided by } 18$.) Anyone who accepts the existence of numbers will accept the following biconditional: ‘A planet has two moons only if and only if the number 2 is the number of its moons’. But who would say that the fact that the number 2 is the number of the moons of Mars counts as a substantive explanation (or any sort of explanation at all) of the fact that Mars has two moons? Or, again, anyone who accepts the existence of (singular) propositions will accept the biconditional ‘A planet has two moons if and only if the proposition that it has two moons is true’. But who would say that the fact that the proposition that Mars has two moons is true counts as a substantive explanation of the fact that Mars has two moons?

In my view, the idea of a *metaphysical* explanation of the fact that (say) this copy of *A Theory of Justice* is green is meaningless—that is to say, there is no such idea, there are only the words ‘metaphysical explanation of the fact that this copy of *A Theory of Justice* is green’, words that darken metaphysical counsel. There are, of course, efficient-causal explanations of the fact that the book is green (explanations

Concluding Meditation

involving the manufacturer's use of green ink), and there are formal-causal explanations of fact that the book is green (its surface absorbs photons and then, the book being more or less in thermal equilibrium with its surroundings, emits photons; the photons it emits have wavelengths predominantly in the 520–570 nm range). And that is all the explanations of the book's being green that there are. No meaningful sequence of (p.348) words in any possible language counts as a metaphysical explanation of the book's being green.

17.2 Reply to L. A. Paul

Laurie Paul and I stand in a relation that I stand in with almost every other metaphysician whom I admire: near total disagreement. I find my root-and-branch disagreement with her immensely profitable, however, because I learn so much from it. I will give an example of the way I have learned from our disagreements that I think is particularly important. She and I once arranged to have a lunch during which we would try to get clear about this "property" business. I knew that she advocated a one-category ontology, an ontology according to which everything is a property. The Pauline ontology, as I shall call it, tells us that there are nothing but properties and sums or fusions of properties—and that a fusion of properties is itself a property. (According to the Pauline ontology, its author, Laurie Paul herself, is a property—as are we all.)⁶

Now Paul and I both say we believe in objects we call properties, but (I supposed) we couldn't mean anything like the same thing by the word 'property'. Given what I mean by 'property' to say, "I am a property" is nonsense. Not nonsense like 'Gubble buggle guggle', of course, and not (exactly) nonsense like '*Das Nichts nichet*' or 'The world is a progressively realized community of interpretation', but nonsense like 'The shadow of Caldwell Hall itself casts a shadow' or 'If you want to extract a cube root, you will need to use a number 18L dental forceps' or 'Sara Bernstein is drinking coffee from a two-dimensional cup'.

At some point during our luncheon conversation, I asked Paul why she believed that the things she called 'properties' existed—why she supposed that there *were* things that had the

Concluding Meditation

properties (in my sense of 'property') that she ascribed to properties (in her sense of 'property'). Her answer transfixed me: she told me that she thought that one of the strongest arguments for the existence of properties_{LP}⁷ was that one could see some of them.

Just look at this Granny Smith apple, for example. You can see the property greenness right there before you; it's *in* the apple. Some would use some such conveniently obscure phrase as 'it inheres in the apple', but the admirably forthright Pauline ontology has no use for convenient obscurity: according to the Pauline ontology, greenness is a constituent, and, in fact, a *part*—a part in the strict and mereological sense—of the apple.⁸ The apple is a fusion of many properties, greenness among them, and that's why (p.349) greenness is before you when the apple is before you. ("But Professor Paul, is it the particular greenness of this apple that is before you when you look at the apple or is it the *universal* greenness?" "Both. The universal greenness is nothing other than the fusion of all particular greennesses. Or put the matter this way: a particular greenness is a connected part, or perhaps a maximally connected part, of the universal greenness. If we use the term 'trope' in the sense the word has in current analytical metaphysics, a trope—like the greenness of this apple—is a part of the object of which it is a trope. For an object like an apple to instantiate a universal is for it to overlap that universal, to share a part, a trope, with that universal. The objects that are wrongly called particulars by some metaphysicians—apples, all 'moderately sized specimens of dry goods,' stars—are fusions, maximal fusions, of tropes. But when I was speaking of the property greenness, I meant the universal." "Ah, but then you never see the universal, you only see parts of it." "Yes, and, unless you're an astronaut, you've never seen the Atlantic Ocean—you've only seen parts of it. The analogy is exact: you can see the color greenness, the universal, when you look at the apple in the same sense of 'see the' as the sense in which you can see the Atlantic Ocean when you're standing on the Outer Banks facing east.")

This was a revelation to me. Suddenly I realized that "platonism" concerning universals (my sort of platonism, at

Concluding Meditation

any rate) and “Aristotelianism” concerning universals are not sisters-in-arms, allies in the battle against nominalism who disagree only on a few small, technical points. (Of course Paul is not an Aristotelian, since she denies that there are “substances,” things that are not properties or universals, things of which other things are predicated but which are not themselves predicated of things. But her picture of the intrinsic nature of universals is very like that of an Aristotelian who contends that universals exist by inhering in substances. She, as it were, “subtracts” the substances from Aristotelianism, and identifies all the objects the Aristotelian typically gives as examples of substances—the Granny Smith apple, a bronze ball, Socrates—with fusions of parts of universals.) Rather, they differ from each other quite as much as either differs from nominalism, even from austere nominalism. It is, indeed, a defensible position that (my sort of) platonism is more like austere nominalism than it is like Aristotelianism in that each gives the same account of concrete particulars: particulars are what David Armstrong has called “blobs”; the only kinds of “structure” a concrete particular has are spatial and mereological structure—and the only proper parts a concrete particular has are other, smaller concrete particulars.

For my part, I do not believe in properties_{PVI} because I can see them. I believe in them because I think that many sentences that express true propositions—‘Spiders and insects share many important anatomical characteristics’, for example—would not express true propositions if ‘There are properties_{PVI}’ did not also express a true proposition.⁹ Indeed I believe that no one can see them—just as no one can extract a cube root with a forceps.¹⁰ Properties_{PVI} are very much the same sort of thing as (p.350) propositions—but where propositions are true or false *simpliciter*, properties are true or false *of* things. Properties_{PVI} stand to declarative sentences in which one variable is free much as propositions stand to closed declarative sentences: ‘The apple that Newton saw fall was green’ expresses the proposition that the apple that Newton saw fall was green; ‘x is green’ expresses the property greenness.

Concluding Meditation

While Paul and I mean different things by ‘property’—at least when we are doing metaphysics—, there is nevertheless a certain *role* that English words like ‘property’, ‘attribute’, ‘quality’, ‘feature’, and ‘characteristic’ play in our discourse, a role defined by the tacit rules embodied in and governing that discourse. Paul thinks that properties_{LP} (assuming that there are such things) are better fitted to play this role than properties_{PV} (assuming that there are such things). That is not why she believes in their existence, but that is why she calls them ‘properties’. And I, of course, think just the opposite. How could we have arrived at these two radically opposed positions?

I believe that a large part of the explanation—perhaps the whole of it—is methodological. Paul believes that the proper method of metaphysics is to construct explanatory theories: there are certain metaphysical data (for example: the book and the apple are both green; we can establish that they are both green by visual examination) and our job as metaphysicians is to order and explain those data.¹¹ It turns out that the best explanation of those data involves postulating properties—properties_{LP}. My method is to examine those truths that we bring to metaphysics from everyday life or from the sciences and other non-philosophical disciplines (and perhaps even from other parts of philosophy, such as ethics) and to attempt to discover their metaphysical consequences—to discover which of their logical consequences (if any) are metaphysical propositions. I have tried to explain (in “Relational vs. Constituent Ontologies”¹²) why I think that Paul’s method (and she is hardly the only analytical metaphysician to employ—and to endorse—this method) and my method result in radically opposed ontologies.¹³ What is really at stake in our methodological disagreement is this: is there such a thing as explaining what it is for something to be, e.g., green (other than a physical explanation: it emits photons with wavelengths predominantly in the 520–570 nm range). As I said in “Relational vs. Constituent Ontologies,”

(p.351) ...no set of statements among all possible sets of statements counts as [a metaphysical or ontological] explanation of what it is for a particular to have a

Concluding Meditation

property or for two distinct particulars to have the same property.

In my view, Paul's methodology is a methodology for solving a problem that does not exist (a problem that cannot be coherently posed), and thus it is not surprising that it leads her to postulate entities that make no sense. (I should perhaps mention that I would say the same thing about *On the Plurality of Worlds*,¹⁴ which I regard as a wholly admirable book, one of the greatest philosophical achievements of the twentieth century. As I've said many times, meaninglessness is what we *risk* in doing metaphysics; what one risks in metaphysics isn't being wrong—except in the sense in which someone who believes that a cube root can be extracted with a forceps is “wrong.” What we risk is *not even being wrong*.)

I am going to say just two more things about the issues that separate Paul and me. The first is a comment on the following definition.

Substance theory takes objects to be, fundamentally, primitively unanalyzable or irreducible substances of different sorts, and holds that substances have properties by standing in some sort of relation to universals or other entities. (34)

My ontology certainly includes substances. (There are at least two importantly different definitions of substance. No matter: the things I call substances satisfy them both. I believe, moreover, that the terms ‘substance’, ‘concrete object’, ‘particular’, ‘individual’, and ‘object that can be an agent or a patient’, although not equivalent in meaning, all have the same extension.) And, as I have said, I agree with the austere nominalist (who may or may not be happy with the term ‘substance’) about the nature of such things. But I reject, I most emphatically reject, the thesis that “substances have properties by standing in some sort of relation to universals or other entities.” I give my reasons for this in my reply to Michael Loux's paper in this volume. (The reader of that reply will see how emphatic my rejection of this thesis is: it would not be entirely unfair to describe what I say there as a rant.)

Concluding Meditation

Secondly, it seems to me that the Pauline ontology must recognize the existence of properties_{PVI}. Consider our Granny Smith apple. Granny (so to name the apple) is green. Granny has (exemplifies, instantiates) the universal greenness. That is to say (according to the Pauline ontology) Granny overlaps the universal greenness: *its* greenness is a common part of the apple and the universal. Well and good. But now consider Great-granny, a slightly larger but very similar apple in the vicinity of Granny. One of the things one can say (and say truly, according to the Pauline ontology; in the sequel, I'll suppose that the Pauline ontology gets the ontological structure of apples right) about Granny is that it overlaps the universal greenness. (I myself said that very thing about Granny a few sentences back.) And one can say the same thing about Great-granny (also truly). Here, I'll *do* it: Great-granny overlaps the universal greenness. So: one of the things that is true of both Granny and Great-granny is "that it is green" (or "the thing that one says of something when one says of (p.352) it that it is green"). And, of course, that's not the only thing that is true of them both. Surely the following statement is true?

There are many things that are true of both Granny and Great-granny.

And that statement certainly *appears* to imply the existence of things that can be true of things; that is, the existence of "unsaturated assertibles"; that is, the existence of properties_{PVI}. I do not see any way in which to paraphrase this statement in such a way that the paraphrase implies the existence of nothing but properties_{LIP}. For—surely?—Paul does not want to say that there is such a property_{LIP} as "overlapping the universal greenness" or that overlapping the universal greenness is a part of both Granny and Great-granny.

17.3 Reply to Louise Antony

Louise Antony's paper is an admirable combination of clarity of argument and moral passion. Nevertheless, either I have misunderstood its central line of argument or it turns on a misunderstanding of my project in *The Problem of Evil* (van Inwagen 2006).¹⁵ I'm going to quote a paragraph from the

Concluding Meditation

paper that illustrates the nature of the misunderstanding on which I believe the paper rests.¹⁶ (The paragraph nicely *illustrates* the misunderstanding, but it is not the only place at which that misunderstanding is manifested.) At one point in the book I said, “It is at least very plausible to suppose that it is morally permissible for God to allow human beings to suffer if the inevitable result of suppressing the suffering would be to deprive them of a very great good, one that far outweighs the suffering.” The following paragraph is a response to this statement.

But this is not good enough. It’s not good enough for a Defense to say that it is “at least very plausible to suppose” that there’s *something* that makes it all right for God to behave in a way that would be morally wrong if a human being did it. The whole point of the Defense is to *show* us what that difference might be. Just as the fundamentalist wants an example of how a part of an eye or a proto-wing might be adaptive, Atheist wants to know what kind of difference between God’s situation and the situation of the human doctor would make it morally permissible for God to do what it would not be permissible for the doctor to do. (183)

To explain why I say that this paragraph illustrates a misunderstanding of my project—of what I was up to, if you like—, I will remind you of the “dialectical context” in which the sentence ‘It is at least very plausible to suppose that it is morally permissible for God to allow human beings to suffer...’ occurs.

Theist and Atheist are conducting a debate before an audience of neutral agnostics.¹⁷ At a certain point in this debate, Atheist is trying to convince the audience of agnostics that, in light of the (global) argument from evil, they should cease to be (p.353) agnostics and become atheists like herself.¹⁸ Theist is not trying to convince the agnostics to become theists (not in this part of the debate, anyway); he is simply trying to convince the agnostics that their reaction to the argument Atheist has presented should be this: “For all we know, at least one of the premises of that argument is false.”¹⁹

Concluding Meditation

In this dialectical context, Atheist bears a much heavier burden of proof than Theist. *She* has to convince the agnostics that all the premises of the argument are true. (That is what she has to do if her intention is to convince them that God does not exist. But if, failing that, she would at least like to convince them that it's significantly more probable that God does not exist than that God does exist—see note 18—then she must at least convince them that it's significantly more probable that all the premises of the argument are true than it is that at least one of them is false.) All *he* has to convince the agnostics of is that they shouldn't be convinced by her arguments. The relations between Atheist and Theist and their audience are very like the relations between the counsel for the prosecution, the counsel for the defense, and the jury in a criminal trial. (Think of matters this way: the proposition that God exists is on trial; the charge is falsity.) In a common-law criminal trial, the burden of proof (the burden of having to prove things) falls on the prosecution. (Or—a picky philosopher's qualification—a very *light* burden of proof falls on the defense. The defense will sometimes tell the jury things like this: that the prosecution has *failed to show* that the defendant was in O'Malley's Bar on the night of the 11th. The defense is then required to prove—that is, is required to convince the jury—that the prosecution has indeed failed to show this.)

As I have described the debate, Theist attempts to convince the “jury” of agnostics that, for all they know, at least one of the premises of the (global) argument from evil is false by telling a story. The story is to have this feature:

(a) It logically implies that at least one of the premises of the argument from evil is false. (If the story implies both that God exists and that the world contains evils of the same sorts and in the same amounts and distributed in the same way as the evils of the actual world, then the story will imply that at least one of the premises of any logically valid version of the global argument from evil is false. And I represent Theist as telling a story that has just those two properties.)

And Theist *hopes* it will have *this* feature:

Concluding Meditation

(b) In the end, when all is said and done, when the debate has run its course, the members of audience of neutral agnostics will react to it, the story, by saying—imagine that they are sequestered and are discussing the plausibility of Theist’s story among themselves—things along the lines of, “For all I know, that story is true” or “The story he has told seems to be a real possibility, one we should take seriously” or “Gee, if God exists, the rest of that story *could* be true, too. I mean for all we know.”

(p.354) I had thought that the ideas appealed to in my statement of (b) were reasonably clear. But I may have been wrong. Antony sometimes uses language that suggests to me that she thinks I use variants on the phrase ‘for all one knows’ with a meaning weaker than the meaning I intended.²⁰ I will, therefore, try to make my intended meaning clearer than I perhaps have.²¹

Let us imagine a murder trial. Winifred has died of arsenic poisoning, the coroner’s verdict was “unlawful homicide,” and Charles has been arrested for the murder and brought to trial. The prosecution’s case against Charles is essentially that he was the only person who both had a reason to desire Winifred’s death (he stood to inherit her vast wealth upon her death) and had the means and opportunity to introduce arsenic into her final meal. If the counsel for the defense suggests to the jury—by an adroit series of questions put to a witness she has called—that a certain person other than Charles also had “motive, means, and opportunity,” and if the prosecution is unable to rebut this suggestion, it is reasonable to suppose that the members of the jury will conclude that “for all they know” Charles is not the murderer, that that other person’s having committed the murder is a real or serious possibility. If, in contrast, the defense counsel calls the lead investigating officer as a witness, and says to him (out of clear blue sky, as it were), “I put it to you, Detective, that the evidence you have presented to this court does nothing to rule out the possibility that Winifred was a spy in the service of a hostile power, and that she was poisoned by a C.I.A. assassination squad—the C.I.A. of course having known that everyone would suppose that she had been poisoned by the accused,” she will get a very different reaction from the

Concluding Meditation

members of the jury. None of them will say, “Gee, for all I know that’s how things happened.” They, or the comedians among them, will rather say things along the lines of, “Oh, sure—and for all I know, my Aunt Harriet’s turnip pancakes are a cure for cancer” and “Yeah, that *might* be what happened. And flaming monkeys *might* fly out of my...um, ears.”²² (The trial judge and the Bar Association will no doubt have something to say about the matter as well.)

A story that has the feature (a)—and which is presented by someone who hopes it will have feature (b)—I call a defense:²³ Theist attempts to convince the agnostics that, for all they know, at least one of the premises of the argument from evil is false by presenting them with a defense. Antony has no objection to this strategy, but insists that the two defenses I have presented²⁴ are defective because they include (p.355) *mysteries*, things that—if they exist, and there is no reason to think that they do—are beyond our understanding. It is as if, she says (well, she doesn’t say it; the example is mine, not hers, but I think she will not object to it), I tried to show that a time machine was (for all we know) possible by presenting a blueprint for the construction of a time machine that included a neat box labeled “flux capacitor,” and explained that a flux capacitor is a mechanism such that, if you constructed a device according to my blueprint, and placed the mechanism in the device at the place indicated in the blueprint, the result would be a working time machine. (After all, “for all we know” it’s possible for there to be a mechanism that has exactly that property.)

I reply that it’s perfectly legitimate for Theist to include “mysteries” in a story that is to serve as a defense—if it is reasonable for him to hope that the agnostics will (despite Atheist’s best attempts to block this reaction) react to that story by saying, “For all we know, that story is true.”²⁵ Or, to put my point in terms of propositions and alethic modality, it’s perfectly legitimate for Theist to include a proposition *p* in a defense even if he has no argument for the metaphysical possibility of *p*. The metaphysical possibility of the propositions his defense comprises is not something Theist hopes to establish or to convince the agnostics of. He hopes,

Concluding Meditation

rather, that the agnostics will—in the end, when Atheist has said everything about the story that she has to say—react to each proposition contained in his defense by saying, “For all we know, that proposition is true”—a statement that is, of course, consistent with, “For all we know that proposition is true, and, for all we know, that proposition is metaphysically impossible.”²⁶ (For all *I* know, a sequence of thirty 7s occurs somewhere in the first 10^{30} digits of the decimal part of π ; and, for all *I* know, it is metaphysically impossible for there to be such a sequence. For all *I* know, signals are sometimes transmitted faster than light—perhaps by advanced extraterrestrial beings who would regard our physics as quaint; and, for all *I* know, it’s metaphysically impossible for signals to be transmitted faster than light.)

I will give an example of an obviously legitimate defense that employs a mystery.

Many people who would like to believe that there are many technologically advanced civilizations inhabiting the planets of other stars are troubled by the fact that we have—as of this date—detected no radio signals whose origin was an extra-solar civilization. Let us call this phenomenon “cosmic radio silence”—‘silence’ for short. And let us use the term “alienist” to describe someone who believes that there is at least one technologically advanced civilization other than our own within 1000 (p.356) light-years of the earth. (There’s no *other* good use for the word.) If theists face “the problem of evil,” alienists face “the problem of silence.”

It would be easy to construct an argument for the falsity of alienism that is similar to the (global) argument from evil in its structure, and in which silence plays the role that evil plays in the argument from evil—the “argument from silence.” It would be so easy that I won’t bother to do it. Now imagine a debate before an audience of neutral agnostics (*sc.* about whether there are other technologically advanced civilizations within 1000 light-years of us). The debate is between Alienist (male) and Humanist (female). (Humanists think that *Homo sapiens* is the only technologically advanced species within 1000 light-years of the earth—“within the Sphere” I’ll say from now on.) Humanist has presented the argument from silence, and

Concluding Meditation

Alienist must respond. His response need not be in the form of a *refutation* of the argument from silence—that is, he need not show that at least one of its premises is false.²⁷ It would, of course, be *nice* (from his point of view) if he could do that, but he need not do it. To respond effectively to Humanist’s argument, he need only convince the agnostics that, for all they know, at least one of its premises is false. What will he say? Well, there is a great deal of relevant and easily available material he can draw upon, for many scientifically literate people have proposed reasons why we might detect no radio transmissions from extra-solar civilizations—reasons other than there being no extra-solar civilizations near enough to us for their radio transmissions to be detectable. Imagine that Alienist offers the following defense, a defense based on one strand of this voluminous material:

Many advanced civilizations inhabit planets of some of the approximately fifteen million stars within the Sphere. Many of them have discovered that other such civilizations exist, and communicate with those neighboring civilizations across intersidereal distances—or with the colonies they have themselves established on the planets of other stars. But only briefly in the “career” of a technological civilization does that civilization use radio waves to communicate over large distances. Technological civilizations typically exist for many millions of years, and typically use radio waves for communication for only a few hundred years. There is at least one other method—perhaps there are many—of sending signals between the stars that is vastly more efficient than radio waves. We have not yet discovered this method—or any of these methods—for it, or they, depend on a level of understanding of the physical world that we have not yet achieved. When a civilization achieves a certain level of understanding of the physical world, it abandons radio waves, just as we abandoned smoke signals and hilltop beacons when we achieved a sufficient understanding of the physical world to invent telegraphy. We can no more detect the signals they send by this method, or these methods, than the Elizabethans could have detected radio waves.

Concluding Meditation

My story, of course, is pure speculation, for obviously we do not know anything about what could be accomplished by physics we have not yet discovered. But it is *plausible* speculation, as a glance at the history of human signaling shows. Advances in physics have frequently led to the discovery of new and better ways of sending signals, signals it would have been impossible for anyone to have detected before those advances were made. And we early-twenty-first-century human beings still have much to learn about physics. We know this because we know that our physics is radically incomplete. We have no inkling of an understanding of why the (p.357) expansion of the universe is speeding up, for example. And our two finest physical theories, the general theory of relativity (our theory of gravity) and the “standard theory” of elementary particles (our theory of everything else) are *logically inconsistent* with each other.

Here, then, is a defense—a defense that Alienist has presented in an effort to counter Humanist’s attempt to use the argument from silence to convert the audience of neutral agnostics to humanists like herself.

But will the effort be successful? Will Alienist’s defense lead the agnostics to conclude that—for all they know—at least one of Humanist’s premises is false? (Presumably the defense would lead them to that conclusion with respect to at most one of Humanist’s premises—‘All or almost all technologically advanced species that send signals to other stars use radio waves for that purpose’, or some premise very much like it.) I can say only that it would lead *me* to that conclusion. (In the absence of some cogent argument from my creature Humanist that I should not take the story seriously. And I can’t think of one to put into her mouth.) And I am no neutral agnostic. I think alienism is almost certainly false—for reasons unrelated to “cosmic radio silence.” But, of course, an argument whose conclusion is true can have false premises, and *a fortiori*, premises that are false for all anyone knows.

Alienist’s defense convinces me—and I believe should convince anyone—that for all we (all we early-twenty-first-century human beings) know, the proposition

Concluding Meditation

All or almost all technologically advanced species that send signals to other stars use radio waves for that purpose.

is false. (That is, it convinces me that the falsity of this proposition is a *realpossibility*, a possibility that is to be taken seriously.²⁸) It doesn't, of course, convince me (nor should it convince anyone) that this proposition *is* false. For all I know, for all any early-twenty-first-century human being knows, there are many technologically advanced species (very distant ones, ones too far away for us to "hear" their signals) who communicate with one another by radio—and none of them uses any other method because radio is pretty much *it*: the regions of physics as yet unexplored by human beings contain nothing that would permit intersidereal signaling that was in any way preferable to good, old-fashioned electromagnetic radiation. But that's beside the point, for—for all we know—the "unexplored regions" *do* contain something of exactly that description, and many "nearby" extra-solar species *do* make use of it to communicate.

Now Alienist's defense, you will notice, contains a mystery: it entails the existence of something that is beyond—far beyond—our present understanding. For all we know, Alienist's defense appeals to something that is simply not permitted by the laws of physics. He certainly—by his own forthright admission—has no proof that, no argument of any sort for the conclusion that, a means of intersidereal signaling more efficient than radio waves is physically possible. But physical possibility is not what Alienist's defense appeals to: it appeals to *epistemic* possibility.

(p.358) Are we to make anything of the fact that my example of a defense that contains a mystery contains a mere "physical mystery," contains something that is, for all we know, physically impossible, while Theist's defense (and likewise the defense that I advance *in propria persona* to account for "the sufferings of beasts") appeals to things that for all we know are *metaphysically* impossible—and asserts that certain actions are morally permissible for God, when, for all we know, those actions would be morally impermissible for any

Concluding Meditation

rational being? No, for what counts is only this: that it is *epistemically* possible (for us, for anyone in our present epistemic condition) that these things are actual—or that it is epistemically possible that a morally perfect being should act as God is represented as acting in the two defenses.

All Theist has to do, all I have to do, is lay these defenses before the agnostics and contend that they are true for all they, the agnostics, know—and then wait for Atheist’s rejoinder. For, of course, we must not forget that the defenses are being presented to the agnostics in the presence of Atheist: that’s an essential part of my “set-up.”²⁹ Atheist is perfectly free to present arguments to the agnostics intended to show that they should *not* regard Theist’s defenses as epistemically possible. That is, Atheist is perfectly free to offer arguments for the falsity of those defenses. And she is perfectly free to present arguments for the conclusion that, for one reason or another, the defenses are simply unbelievable—and thus do not represent real possibilities.³⁰ (Of course, she must do this in the presence of Theist, part of whose dialectical function is to do everything—provided it is rational and intellectually honest—that can be done to undermine the plausibility of such arguments as Atheist may present to the agnostics in her attempts to convince them that they should not regard his defenses as “true for all they know.”) She is perfectly free to offer arguments for their metaphysical impossibility or their “moral impossibility,” or to contend that they are inconsistent with some known facts—the known facts of human evolution, for example—and thus *not* “true for all anyone knows.” That’s *her* dialectical function.

I do not, therefore, think I should have included more in my stories, my defenses, than I did. But I could have. I included as much as I did and did not include more for a simple, practical reason. The story, when it included everything I regarded as essential, had got pretty long—to the point of bordering on unwieldiness. And I did not want to push it over that border. But I will give an example of something more I *could* have included. Toward the end of her paper, Antony asks,

Concluding Meditation

Why think that, once they are removed from the world of horrors, restored to safe and blissful union with God, that these human beings won't just do the same damn thing again?

(That is, why think that they won't rebel against God a second time?) I could have included something in the story I put in Theist's mouth to explain why they won't. In [\(p.359\)](#) fact, in another presentation of more or less the same defense (there called a 'theodicy'—but the word was being used in a carefully defined but non-standard sense) I included just such a "something".³¹ The core idea of that "something" is well expressed by the proverb, "Once burned, twice shy." My present point is not to present that "something" but to point out that *Theist is not required to include any such thing in his defense*. All he is *required* to do is tell stories, and if necessary, to reply to Atheist's critiques of his stories, her arguments for the conclusion that the stories—if they are possible at all—do not represent possibilities that the agnostics should take seriously. Now it happens that one of the stories he tells includes the following statement (it is made in different words):

Once they are removed from the world of horrors, restored to safe and blissful union with God, these human beings *won't* just do the same damned thing again.

Of course, while he's telling the story that includes this statement, Atheist is right there listening to him tell it. And she's free to try to turn the idea contained in Antony's pointed question into an argument for the falsity or "simple unbelievability" of that statement. And, if she does try that, Theist will be right there listening to that attempt, and will be ready to present the "Once burned, twice shy" rejoinder or any other rejoinder that may occur to him. Antony writes as if the "debate" aspect of the interaction between Theist, Atheist, and the neutral agnostics did not exist, as if that all Theist had to do was first to present a defense to the agnostics—the Extended Free Will Defense, let us say—and then, having presented it, to sit back in the hope that they will react to it by saying, "Golly, that story is true for all we know." But that is

Concluding Meditation

not the point at which he hopes to hear those words, or not the only point. He hopes to hear the agnostics say, “Golly, that story is true for all we know,” *after* Atheist has done her worst—after she has brought all the dialectical pressure to bear on the Extended Free Will Defense that an *ideal* representative of atheism can.

In the end, I have to say that I am puzzled by Antony’s demand that a proper defense would contain more than my proposed defenses do—more details about how various things that figure in the defenses *work*, arguments in support of the thesis that this, that, or the other element in the defenses is metaphysically possible, reasons for supposing that an action that would be morally impermissible for human beings is, nevertheless, morally permissible for God.... I *think* the explanation has to be this: she thinks I was doing something other than what I was doing. I can only suppose that she thinks I was trying to do at least one of these two things:

To show that the vast amounts of ill-distributed suffering the world contains are consistent (and not merely logically consistent) with the existence of God—that is, that there are *metaphysically* possible worlds in God and “actual evil” co-exist.

To show that it is or at least could be reasonable for someone who was aware of the vast amounts of ill-distributed suffering the world contains to believe in the existence of God.

Well, I was not trying to show that the existence of God was consistent with existence of “actual evil.” I wouldn’t know how to do that—although of course I *believe* that (p.360) they’re consistent. In fact, I wouldn’t know how to show that *any* two things were consistent (other than by showing that they both existed or were both true, depending on the nature of the “things”). I know what it is to rebut arguments for the conclusion that two things are inconsistent, but I’m not at all sure what would count as showing that two things were consistent. (Formally or logically consistent, yes, *that’s* certainly possible—for what it’s worth, and that’s not much. Consistent relative to something else, yes: it’s certainly

Concluding Meditation

possible to show that Riemannian plane geometry is consistent *if* Euclidean solid geometry is consistent. But to show that two things are consistent *simpliciter*—what could that mean?)

And I was not trying to show that it is or could be reasonable to believe in the existence of God given our knowledge of “actual suffering.” Again, I wouldn’t know how to do that. I wouldn’t even know how to prove that belief in the existence of God was or could be reasonable in a world that contained no suffering at all. (Of course, I believe that belief in the existence of God is reasonable—and reasonable even in this sorry world—, but I wouldn’t know how to *show* that it was reasonable or even that it is reasonable for some beings in some metaphysically possible world containing suffering comparable to that of the actual world.)

I had thought I had made it clear that I had only this goal: to defend the thesis that the argument from evil has premises that are, for all anyone knows, false. That’s a much weaker thesis than either ‘It is metaphysically possible for God and “actual evil” to co-exist’ or ‘It could be reasonable for someone who was aware of the existence of “actual evil” to believe in the existence of God’. Nevertheless, that conclusion is not without interest. The argument from evil is an important philosophical argument—one of the *most* important philosophical arguments. And if, for all anyone knows, at least one of its premises is false, then no one should regard it as a cogent argument for its conclusion.³²

I would also remind Antony that my intention was not to convince someone like *her*—an atheist; if I may so express myself, a hardened atheist—that at least one of the premises of the argument from evil is false for all *she* knows. I’m not up to that task. But it might be easier for one to convince neutral agnostics that a premise of the argument was false for all *they* knew (and to convince them of this in the presence of someone like her, someone whose dialectical function it was to call their attention to every debatable point in one’s arguments) than it would be for one to convince Antony that that premise was false for all *she* knew.³³

Concluding Meditation

(p.361) I would remind her, finally, that in my view, whatever else it is, the argument from evil is a philosophical argument—and, in one important respect, at least, a *typical* philosophical argument. There is, in my view, *no* philosophical argument for any substantive, positive³⁴ conclusion that would fare well in a debate before an audience of neutral agnostics (neutral agnostics with respect to its conclusion). I might, for example, debate with an able compatibilist (let's suppose it was David Lewis) about the compatibility of free will and determinism before an audience of neutral agnostics (agnostics about whether free will and determinism are compatible). At one point in the debate, I might propose to convert the agnostics to incompatibilism by presenting the Consequence Argument. But I am very close to being certain that the result of an exchange in which Lewis and I argued about the premises of the Consequence Argument till the cows came home would be this: in the end, he would be able to convince the agnostics that the argument had at least one premise that was, for all they knew, false. He wouldn't be able to convince *me*—a hardened incompatibilist—of this (even David Lewis wasn't up to that task) but he would be able to convince *them*.³⁵

17.4 Reply to Alex Rosenberg

There are several points in Alex Rosenberg's description of the content of "Weak Darwinism"³⁶ that I'd like to dispute. But nothing a scholar can write is more boring than a detailed, paragraph-by-paragraph documentation of the thesis that that scholar has been misrepresented by a critic. I choose instead to discuss two substantive arguments in Rosenberg's paper, arguments that do not depend on the way Rosenberg has represented my views—arguments that, if they are cogent, refute what I would concede are the central theses of "Weak Darwinism." I will first discuss Rosenberg's argument for the conclusion that "Allism" is supported by the second law of thermodynamics. I will then turn to his argument for the conclusion that one of *my* arguments—the argument for the compatibility of Darwinism and theism—is flawed.

Concluding Meditation

17.4.1 Allism and the Second Law of Thermodynamics

Allism, Rosenberg contends, is supported not so much by the data of biology and paleontology (although data from those sciences certainly support a thesis I explicitly affirmed, namely that natural selection must play an important and essential role in any explanation of the features of life on the earth, either as it is at present or as it has been at any point in its long history) as by physical theory: by the second law of thermodynamics.

Rosenberg's argument—if I understand it—proceeds from two premises: the thesis of a “zero adaptation starting point,” and the thesis that every episode in which some (p.362) type of organism becomes better adapted to its environment must (at least to a very high probability—a qualification I will ignore in the sequel) take place in accordance with the second law.

I will discuss the second premise first. I would not dream of disputing it. I agree that every aspect of the biological history of the earth—whether it involves adaptation or not—has happened in accordance with the second law. The surface of the earth, its atmosphere, and its oceans contain high local concentrations of order (Rosenberg and I are two of them and a bacterium is another), but that order was not created *ex nihilo*: it was “paid for” by a vast reduction in the order inherent in the physical content of other places, such as the solar core. When the earth was lifeless—when there were none of these organisms, these high local concentrations of order, on its surface or in its atmosphere and seas—, the earth and sun taken together contained vastly more order than they do now (most of it distributed pretty uniformly throughout the solar core, among as-yet-unfused nuclei under immense pressure). Or consider an automated factory, which builds, let us say, widgets out of materials available in its immediate environment. Each individual widget may contain more order that can be found in any widget-sized region of space within the factory or within the local deposits of raw materials it draws upon; nonetheless, the total amount of order in the factory, its sources of energy, the remaining deposits of raw materials, and the stock of produced widgets, not only does not increase when a new widget is produced but must decrease every time one is produced. (The “lost” order will be

Concluding Meditation

dispersed over an eternally expanding volume of space in the form of thermal radiation—heat produced, for example, when one of the factory’s automated lathes breaks molecular bonds in a chunk of raw material, heat that the factory is unable to utilize to power its operations.)

But why is it that any explanation of the high local concentrations of order in the biological world (whether in terms of adaptation or not) must appeal to no order-concentrating mechanism but natural selection? Natural selection can certainly concentrate order, but there are other mechanisms by which order that is dispersed over some region can be concentrated or “pumped into” a smaller region, and these other mechanisms do not violate the second law. We human beings, for example, often produce high local concentrations of order by *designing and building* things, by applying *foresight* and *knowledge* and *intelligence*, and we are not miracle workers: the high local concentration of order in a stack of firewood or in a twenty-dollar bill or in an ampule of Botox has been paid for by the loss of a greater amount of order somewhere in the world. I do not say this to suggest that “intelligent design” has been at work in the history of life on the earth. I am rather pointing out that *any* natural process that produces high local concentrations of order—even intelligent design—will have to work within the constraints of the second law.

Let me tell you a story (the genre is science fiction). It is 100,000 years in the future, and human beings have become like gods (albeit they are only what the sci-fi writer Charles Stross has called “*weakly godlike*”: they cannot violate the laws of physics). Our remote descendants of those days change the trajectories of stars at a whim and design and implement entire biospheres for hitherto lifeless planets. And the biospheres they design, while natural selection inevitably plays a significant role in them, employ other mechanisms of—well, let us not say *evolution* but *biological* (p.363) *development*. Inherent in each biosphere there is a “biopilot,” a sort of computer. (The physical substrate of this “computer” is entirely biological, the biosphere as a whole—its operations are “distributed” throughout the totality of living things. The

Concluding Meditation

biopilot exists as an *aspect* of the biosphere: it is a virtual and not a “real” entity.) Its function is to search out ecological niches that are unoccupied or only marginally or inefficiently occupied and to *design* taxa to occupy them. And, being a part of the physical world—if only a virtual one—the biopilot is able to bring its designs to reality. It “observes,” for example, that the niche “apex predator” in a certain rainforest is occupied by an organism that is not exploiting the opportunities that niche provides very efficiently, designs an optimum apex predator, works its magical genetic engineering—“Any sufficiently advanced technology is indistinguishable from magic”—, and within a few score generations, the old apex predator either conforms to the biopilot’s design or has become extinct and has been replaced by a new species that does. (Unaided natural selection does things like that, too, of course—but much, much slower. That is precisely why our godlike descendants build biopilots into their designer biospheres.) It seems to me that a species that can control the behavior of stars might well be able to design and implement biospheres with such virtual pilots. “Allism” would of course be false as a thesis about such biospheres—and yet there would be nowhere within them any violation of the second law: every local decrease of entropy brought about by the actions of the biopilot would, must, be paid for, and more than paid for, by an increase in the entropy of other parts of the universe (mainly in the core of the primary of the planet “inhabited” by the biopilot and the organisms in its charge). The biopilot is a designer, to be sure, but not an *intelligent* designer, for it is no more intelligent than Watson (or, more exactly, than the DeepQA computer architecture).

“But, the biopilot is a product of intelligent design—as was the DeepQA architecture. So its productions *would* be a result of intelligent design—at one remove.”

True—although that doesn’t touch the point that the biopilot operates within the constraints of the second law—, but who’s to say that *natural selection* couldn’t at some point produce something like the biopilot? (This would be something like the “Gaia hypothesis,” although, I concede, considerably stronger.) After all—I know that you, the interlocutor, will not

Concluding Meditation

dispute this statement—natural selection produced *us*, produced intelligence. In my sci-fi story, therefore, natural selection produced the biopilot at, to use your phrase, one remove. Are you sure that natural selection couldn't, as it were, eliminate the middleman and produce something like the biopilot without "going through" intelligence? If natural selection can indeed produce intelligence, I'd be very cautious about saying what biological objects or phenomena (consistent with the laws of physics) natural selection *can't* produce. In "Weak Darwinism," I suggested that, *for all anyone knows*, other biological mechanisms than natural selection *may* have played a central role in the story of how life on the earth got into its present state. In saying this I explicitly left open the possibility that at one time natural selection was the only such mechanism and that "additional" mechanisms as there may be came to be solely as a result of the operations of natural selection. But if that has happened, Allism is *false*, for Allism is not simply the thesis that natural selection is *source* of all other mechanisms that have operated to produce such things as taxonomic diversification, biological complexity, (p.364) and apparent teleology in nature: it is the thesis that natural selection is the only such mechanism there has ever been.

I suppose I should explicitly say that I am not putting forward the hypothesis that something like a biopilot has actually played a role in the history of terrestrial life as a real possibility, one we should assign a significant credence to. The biopilot-produced-by-natural-selection is nothing more than a philosopher's counterexample to another philosopher's contention that a certain proposition is a necessary truth: if there *were* a natural-selection-produced biopilot, Allism *would be false* and both Rosenberg's premises *would be true*.) The intent of my story has been to demonstrate that it is not evident that there is any inherent impossibility in the idea of other mechanisms than natural selection that both figure in the explanation of taxonomic diversification (etc.) and are consistent with the second law of thermodynamics. If there were such mechanisms, I could hardly be expected to produce a plausible story about what they might be—any more than Lord Kelvin's contemporaries could have been expected to

Concluding Meditation

produce a plausible story about what mechanisms might have enabled the sun to shine for many hundreds of millions of years.

I contend, therefore, that Rosenberg's conclusion does not follow from his premises.

17.4.2 The Consistency of Darwinism and Theism

Rosenberg's critique of my argument for the consistency of Darwinism and theism makes extensive use of the concept of objective chance. I find his use of this concept extremely puzzling.

Here is one of the questions that puzzles me: is the existence of "objective chance" consistent with strict causal determinism? If it is not, then Darwinism entails causal indeterminism, which is a very implausible idea. Adherence to Laplacian determinism was widespread in the nineteenth century, and no one in those days saw Laplacian determinism as an objection to Darwin's theory. Computer programs that mimic and illustrate the operations of natural selection are common enough, and these programs are entirely deterministic: the only source of "objective chance" in a particular application of one of these programs is the filling in of open parameters in the program by the human beings making the application—the fixing of initial conditions, as it were. In any case, it is not the business of a biological theory to pronounce on the question of determinism. Whether the world is deterministic may be a question for physicists or it may be a question for metaphysicians. One could argue about that. But it is certainly not a question to be answered by biologists.

Suppose then that the existence of objective chance (in whatever sense Darwinism requires "objective chance") is consistent with determinism. And suppose that Allism is true and that the universe *is* deterministic. (Let us say that it is deterministic in the sense imagined by Laplace, just to have a version of determinism to play with.) Then an omnipotent God can easily bring it about that natural selection produces *any given* organisms that it is possible for natural selection to produce. And, of course, the way he can do this is by an

Concluding Meditation

exquisitely precise choice of initial conditions. Suppose, for example, that God wanted natural selection eventually to produce human beings. Well—so the Darwinists say—, natural selection *did* produce human beings. If they are right, therefore, all God had to do to achieve his end was to create an array of (p.365) particles with certain precisely specified relative positions and momenta—these being the positions and momenta those particles *actually* had at some point in time before there were any living things. Being omnipotent, he would be able to do that. Being omniscient, he would know that *if* he created such an array, the physical universe would after such and such an interval evolve (in the physicists', not the biologists' sense of 'evolve') into a state that included the existence of human beings (and, for that matter, *Darwinilus sedarisi* and the Taj Mahal).

I have not been able to see the force of Rosenberg's objection to my little argument for the compatibility of theism and Darwinism.³⁷ I will therefore present a slightly more elaborate version of the argument and ask him what his objection to this version of the argument would be:

God is often said to be achronic or timeless or "outside time." (This was the position of Augustine and Anselm and Thomas Aquinas and many other great Christian philosopher-theologians.) Let us suppose that this doctrine is correct. Consider the timeless God contemplating (timelessly) the creation of a cosmos. Since he is omniscient, all possible distributions of matter and radiation in space-time are present to his mind. Consider the one among them that is displayed by the actual cosmos. God chooses that one and says, "Let it be!" or "Let there be a cosmos that consists of matter and radiation distributed in space-time in *that* way!" And, by that act, a complete four-dimensional whole—from the Big Bang to, well, whatever—timelessly *is*.

Now if Darwinism is both a true theory and a scientific theory, its truth must have been established by this timeless act of creation. The truth of a scientific theory must "supervene on" (must be "settled by") the

Concluding Meditation

distribution of matter and radiation in space-time. A theory whose truth or falsity does not supervene on the distribution of matter and radiation in space-time is not called a “scientific” theory. Such a theory is called a “metaphysical” theory.

I have told a story in which an omnipotent being has created a Darwinian world. At any rate, those who accept the Darwinian theory should agree that I have told a story with that feature, since they believe that Darwinism is both a true theory and a scientific theory, and every (actually) true scientific theory must be true in any cosmos in which matter and radiation are arranged as they are in the cosmos we inhabit.

(A few years ago, I presented this argument in a sort of symposium on science and religion in Rome. One of my fellow symposiasts was the mathematical biologist Martin A. Nowak, who is Professor of Biology and Mathematics at Harvard, and Director of the Program for Evolutionary Dynamics. Professor Nowak and I had not seen each other’s papers before they were presented, and we were amused—and I, at least, was gratified—when it turned out that our two papers presented essentially the same argument for the compatibility of Darwinism and theism.) I should say that I do not think the above “four-dimensionalist” story is actually true. (For one thing, it seems to me to be inconsistent with human free will.) But I do not put it forward as true. I put it forward as a logically consistent story in which the physical universe has (p.366) been created by God and which incorporates the truth of any scientific theory that is true in reality. That the story is logically consistent does not prove that it is metaphysically possible, of course, but the story is metaphysically possible if it is metaphysically possible for there to be a timeless omnipotent being *and* metaphysically possible for Allism to be both a scientific theory and true. And if the story is possible, its possibility implies that God can create a world—it need not be governed by deterministic laws—in which human beings are the product of natural selection.

And does “objective chance” exist in this world? Well, it does if the existence of objective chance is necessitated by the way

Concluding Meditation

matter and radiation are distributed in the space-time of the actual world. If the existence of objective chance is necessitated by the way matter and radiation are actually distributed in space-time, then God can create a world that contains all the objective chance any scientific theory that applies to the actual world needs. If it is not so necessitated, then either (a) objective chance does not exist and ‘Objective chance exists’ is not a component of any *true* theory, or (b) ‘Objective chance exists’ is a metaphysical proposition—owing to the fact that its truth-value does not supervene on the distribution of matter and radiation in space-time—and should not be a component of any *scientific* theory.

17.5 Reply to David J. Chalmers

I am largely in agreement with David Chalmers’s fine paper. But I wish to say a few things about some of his statements about my own views (or, in one case, Colin McGinn’s views)—to set the record straight, as it were. I shall quote some passages from the paper and comment on them. None of the things I say is of very great relevance to any of the main theses of the paper. Each quoted passage from Chalmers’s paper (after the first) marks a transition from a discussion of the previously discussed passage to a discussion of the “new” passage.

17.5.1 Disagreement and Knowledge

I begin with:

A strong version of this view, suggested by van Inwagen’s discussion, is that where there is sufficient disagreement among experts, no individuals can be said to know the truth. Even if some individuals have hit on good arguments for true conclusions, how can they have justified confidence that these are good arguments, when so many of their peers disagree? I am not so sure: I think that at least in some cases, a good argument can ground an individual’s knowledge of a conclusion even when peers reject it. For example, I think that the presence of any number of peers who deny the existence of consciousness would not undermine my knowledge

Concluding Meditation

that I am conscious. Likewise, it would not undermine arguments that take this claim as a premise. (p. 285)

Whether that view is suggested by my discussion or not, I will point out that I have presented explicit arguments for a position that might be loosely expressed in these words: ‘Where there is sufficient disagreement among *philosophers*, no individual philosopher can be said to know the truth’.³⁸ But that position applied only to (p.367) disagreement about philosophical theses or beliefs that satisfied the following condition: the theses must not be theses are accepted

...by almost all human beings. I shall not be concerned with philosophical theses that have been accepted by all sane non-philosophers and have been denied only by a few philosophers—generally practitioners of “revisionary metaphysics.” I assume that there are such philosophical beliefs because I assume that the denial of a philosophical belief is itself a philosophical belief, and many philosophers have believed things (in, as it were, their professional capacity) that almost everyone—even most philosophers—would deny. Or so it seems at least plausible to maintain. Plausible examples of things that fall into this category would be: “Change and motion are not real features of the world”; “One has no reason to suppose that there are minds other than one’s own”; “There are no material objects.”...Thus, philosophical beliefs like “Change and motion are real features of the world,” “One does have reason to suppose that there are minds other than one’s own,” and “There are material objects” do not satisfy my second condition.³⁹

Let us call philosophical propositions that are not “accepted by almost all human beings” “controversial philosophical propositions.”⁴⁰ The thesis I defended was (it was not stated in these words; I have rewritten it to make it as verbally similar to ‘Where there is sufficient disagreement among experts, no individuals can be said to know the truth’ as possible):

If there is extensive disagreement among philosophers⁴¹ about whether a certain controversial philosophical

Concluding Meditation

proposition p is true (assuming that all philosophers who accept either p or its denial are aware of and have a perfect understanding of all the arguments that any of them has given that are relevant to the question the truth or falsity of p ⁴²), then it is irrational for any philosopher to accept p .

Call this thesis the Philosophical Disagreement Thesis or PDT. (If PDT is true, then I believe many, many things it is irrational for me to believe—and I expect you do, too. My only defense is of the “ought implies can” variety: I can’t help it; I can’t help believing things that seem to me to be obviously true.)

Now let us ask, is ‘Consciousness exists’ a counterexample to PDT? Before we can address this question, we must answer a prior question. Suppose a philosopher uttered the sentence ‘Consciousness does not exist’ and intended thereby to be stating a philosophical thesis. What would that thesis be—what would it *imply*? The only answer to this question that comes to my mind is that it would imply that *all* the following statements (and many millions of statements of the same sort) were false:

(p.368) Lots of people suffer from migraine—and migraines are extremely painful.

Sometimes people have to make difficult and unpleasant decisions—but they do often somehow manage to make them.

There are now [i.e., in the summer of 2015] people who believe that Hillary Clinton will almost certainly be the next President of the United States.

Some couples who want to adopt a child are unable to.

At any rate, *I* am unable to see why a thesis would deserve to be called ‘the thesis that consciousness does not exist’ if it didn’t imply that those statements (and perhaps 90 percent of the statements that human beings make in the course of conducting the business of their lives) were false—and, of course, if it didn’t imply that their denials were true.

Concluding Meditation

Is 'Consciousness exists', so understood, a controversial philosophical proposition? I'm willing to stipulate (as the lawyers say) that it's a philosophical proposition. But it's certainly not "controversial": it's certainly held by almost all human beings—at least tacitly or implicitly. That is to say, almost all human beings accept, if not precisely the four statements given above, millions upon millions of statements of much the same kind. So even if a given "man on the Clapham omnibus" has never considered the proposition 'Consciousness exists' he can certainly be said to accept it in the same sense as that in which he can be said to accept 'Fish do not wear shoes' or 'Butter is not a high explosive'. And, I expect, what goes for the passengers on the Clapham omnibus goes for most philosophers. I very much doubt whether very many philosophers accept any proposition that implies that the above four statements are false. And I am convinced of the truth of a rather stronger statement: while there may well be possible worlds in which a sizable proportion of philosophers accept such a thesis, such worlds are very *distant* from the actual world. (And I am convinced that such worlds are not only distant but *rare*—that is, that the objective probability of philosophy's being in such a state, while greater than 0, is very low.)

Now let's suppose that the stronger statement is true. Let w be such a distant (and improbable⁴³) world: although the medical and political (and so on) features of w are much like those of the actual world, the condition of philosophy in w is very different: w is a world in which a sizable proportion of philosophers accept a thesis that they express as 'Consciousness does not exist'—words that, in their mouths, have a meaning that entails the falsity of our four statements. I am certainly willing to concede that if David Chalmers exists in w (and if his "doxastic state" in w is as similar to his actual doxastic state as the bizarre condition of philosophy in w allows⁴⁴) he knows (in w) that "consciousness exists." (And his belief that "consciousness exists" is rational—in case that doesn't follow from his knowing that "consciousness exists.") Nevertheless, 'Consciousness exists' is not a counterexample to PDT (in w), owing to the fact that, in w as in the actual

Concluding Meditation

world, that proposition is accepted by almost all human beings.

(p.369) Are there possible worlds in which 'Consciousness exists' is *not* accepted by almost all human beings? I very much doubt whether there's a world that has that feature *and* contains "extensive disagreement among philosophers" about whether consciousness exists. (A "Zombie" world would not satisfy both these requirements—for, although in a Zombie world *no* proposition is accepted by *any* human being,⁴⁵ by that very token there will be no such thing as extensive disagreement among philosophers in a Zombie world.) But even if I'm wrong about that and there is an "oddfolk-madphil-saneDave" world—a world in which 'Consciousness does not exist' is *consistent* with all the beliefs of the Clapham-omnibus passengers, in which a significant proportion of philosophers *accept* that proposition, and in which Chalmers (nevertheless) *knows* that consciousness exists—, an oddfolk world (whether or not it was also a "madphil" and "saneDave" world) would be *vastly, radically different* from the actual world. And it would be *vastly more* different from the actual world than a "mere madphil" world—a world in which, although the folk, almost to a person, tacitly accept 'Consciousness exists', a goodly proportion of philosophers deny it. If an oddfolk-madphil-saneDave world exists, its existence implies that PDT is not a necessary truth. PDT is, nevertheless, false only in worlds that are *vastly, radically different* from the actual world. In that respect it could perhaps be compared with 'If human beings exhibit the kinds of complex linguistic and non-linguistic behavior they exhibit in actuality, they are conscious': even if it is not a necessary truth, it is true and its truth is firmly grounded in the Nature of Things as They Are.

17.5.2 Establishing Lack of Success

Does this mean that all philosophical arguments for positive theses are unsuccessful, as van Inwagen...has suggested?...This depends on what one means by success. If one defines success in sociological terms, so that success requires convincing almost everyone in a community, then we have seen that at best very few

Concluding Meditation

philosophical arguments for positive theses have been successful in our community. Van Inwagen defines success in idealized epistemological terms: a successful argument for a proposition p is one that would convince an audience of ideal reasoners who are initially agnostic concerning p , in the presence of an ideal opponent of p . I do not think that the sociological observations above (or the sociological observations that van Inwagen appeals to) come close to establishing that no philosophical arguments are successful in that sense. (p. 289)

Well, of course my arguments don't *establish* that conclusion. Any argument whose conclusion is a philosophical thesis is a philosophical argument. (At any rate, I don't see any other way to understand 'philosophical argument'.) Say that an argument for a proposition p that "would convince an audience of ideal reasoners who are initially agnostic concerning p , in the presence of an ideal opponent of p " is 'successful_{PVI}'. The thesis

Neg No philosophical argument is successful_{PVI}

is a philosophical thesis. If some argument did establish the conclusion that no philosophical arguments are successful_{PVI}—did establish the truth of *Neg*—, then *Neg* would be false!

(p.370)

And while I didn't explicitly say that no philosophical argument "comes close" to establishing its conclusion, I believe that, too. (And, anyway, suppose that some philosophical argument "came close" to establishing its conclusion—that human beings have free will, let's say. It's at least a defensible position that that argument would *establish* the conclusion that human beings *probably* have free will—and isn't 'Human beings probably have free will' a philosophical thesis?) I will say that I regard my arguments for *Neg* as pretty good ones—as *philosophical arguments go*. They are, I would say, no worse than the arguments of Kripkenstein or Frank Jackson's "Mary the color scientist" argument or one of the more sophisticated versions of the cosmological argument. But, despite being pretty good as philosophical arguments go, those three arguments don't establish their conclusions—nor,

Concluding Meditation

sadly, do they come close to establishing their conclusions. (“But Jackson’s argument not only comes close to establishing its conclusion, it *does* establish its conclusion.” Really? If so, why does Dan Dennett think it’s a tissue of fallacies? If it establishes its conclusion, what are we going to say about poor Dan? What is he? Stupid? Philosophically inept? Intellectually dishonest?)

17.5.3 Explaining Lack of Success

Evolutionary explanations: It is sometimes suggested (e.g. by McGinn 1993) that there is a Darwinian explanation for the lack of progress in philosophy. The rough idea is that we did not evolve to be good at philosophy, since in the evolutionary environment there were no selection pressures that favored philosophical ability or anything that strongly correlates with it. Perhaps there is something to this, though it would take some work to explain why the same does not apply to the ability to do abstract mathematics or highly theoretical science. (p. 295)

McGinn’s position⁴⁶ is that that the same *does* apply to our ability to do abstract mathematics and highly theoretical science. That we (as a species, if not in every case as individuals) happen to be good at these things and bad at philosophy has no evolutionary explanation. It’s just one of those things that have *happened* in the course of biological history—like the pentadactyl limb. (A species can’t be good at everything. As Chomsky says somewhere, to be equally good at everything would be to be very bad at everything.) We may one day discover among the stars a species that, for no reason having to do with selection pressure or reproductive advantage, is good at philosophy.⁴⁷ And we may also imagine that this species is bad at physics. (At any rate, the philosophers of this alien species all agree about the answers to philosophical questions, or come as close to agreeing as human scientists do to agreeing about the answers to scientific questions. The alien philosophers are—they tell us—puzzled and confounded by the inability of human philosophers simply to *see* that Zombies are metaphysically impossible and that free will is compatible (p.371) with

Concluding Meditation

determinism. And when the alien philosophers present what they regard as decisive arguments for such conclusions to the human philosophers, the humans can't follow them: every time an alien philosopher presents human philosophers with a careful, step-by-step presentation of one of the aliens' supposedly decisive arguments, at some crucial point in presentation of the argument, the human philosophers confess that—to borrow a useful term from the Lexicon⁴⁸—the reasoning they are being asked to follow goes all wilfrid in their minds. And the aliens, for their part, very soon become confused and unable to understand what we're saying when we try to explain rainbows or planetary orbits to them.)

[Van Inwagen has] advocated unknowability: humans are just not smart enough to answer the big questions. The idea is that there is some level of intelligence or aptitude that would suffice to answer these questions, but that humans fall below that level.

Van Inwagen argues for this conclusion as follows. He suggests that it is implausible that we are much above that level, given the lack of progress to date, and that it is antecedently improbable that we should be just barely at that level. So it is much more likely that the level lies above us. I am not so sure about this argument. I think we already know that for a vast range of questions, humans are just barely at the level for doing them well: scientific and mathematical questions, for example. Because of this, it is arguable that we lie at a special intelligence threshold at which an extraordinarily wide range of questions come to be within our grasp over time. It is not obvious whether or not philosophical questions fall within that range, but it is not obviously more likely that they do not than that they do. (p. 297)

This misrepresents the conclusion of my argument. The conclusion of the argument is not that we are “not smart enough to answer the big questions.” (The argument pertained only to metaphysics—it occurred in the “Concluding Meditation” of an introductory textbook of metaphysics⁴⁹—, but I am willing to say that a parallel argument would apply to all questions of the form ‘Is the case that p ?’, where p is one of

Concluding Meditation

those “controversial philosophical propositions” I spoke of earlier.) The conclusion of the argument was closer to this statement.

If it turns out that we human beings are not smart enough to answer the big questions (*sc.* of metaphysics), we should not find that surprising.

It was in fact *this* statement:

...if metaphysics does indeed present us with mysteries we are incapable of penetrating, that fact is not itself mysterious. (p. 289)

(Suppose that Jane is about to draw a card at random from a standard deck. Obviously the two statements ‘The card Jane will draw is the three of clubs’ and ‘If it turns out that the card Jane draws is the three of clubs, we should not find that surprising’ are not equivalent.⁵⁰) And I would say the same thing with respect to (p.372) physics: if physics presents us with mysteries we are incapable of penetrating, we should not find that mysterious—not, at any rate, if we know that we are *just* able to get as far in physics as we have. Suppose, for example, that the nature of “dark energy” is a mystery—in the sense that we human beings simply do not, *as a species*, have within us cognitive resources that would enable us to discover its nature. (If physicists and cosmologists actually manage to refer to something when they use the phrase ‘dark energy’, it must be that “something” *has* a nature.) Suppose we one day encounter another intelligent species, the Spicans, and discover that they are demonstrably less good at physics than ourselves, although only *a bit* less good. The Spicans had a pretty good analogue of Newtonian mechanics before we first visited them, for example, and a sort of gesture in the direction of special relativity, but they did not have anything resembling general relativity. When we teach general relativity to their physicists, they can *learn* it, but we’re convinced, and they themselves sadly admit, that they’d never have figured it out for themselves: no member of their species (we believe and they concede) would ever have been able to do what Einstein did (and what Hilbert would almost certainly have done in 1915 or 1916 if Einstein hadn’t done it first).

Concluding Meditation

Suppose we're all agreed—on the basis of these facts about the Spicans' scientific abilities—that the Spicans (now that we've informed them of the cosmological phenomenon that dark energy has been postulated to account for) are unable, biologically unable, to discover the nature of dark energy. Then, I contend, it would be reasonable for us to conclude that if the nature of dark energy is a mystery (to us), there's nothing mysterious about its *being* a mystery. If the Spicans, who are only a bit less good at physics than ourselves, are unable to discover its nature, then, if we are able to discover its nature, we're only *just barely* able to discover its nature. (As some of us are just barely able to run a mile in four minutes—unlike horses and Houyhnhnms, who are able to run a mile in four minutes and are *not* just barely able to run a mile in four minutes.) And that seems rather improbable. Why should we be “just over the line”? So there's a reasonable probability that we're not able to discover the nature of dark energy—a probability high enough that we shouldn't be surprised if we're not able to discover its nature.⁵¹ If moreover, we are able to penetrate all the mysteries—or what had at some point in the history of science *seemed* to be mysteries—of physics, we should say (on the basis of our knowledge of the Spicans)

(p.373) If physics *had* presented us with mysteries we were incapable of penetrating, that fact *would not have been* mysterious.

As I said in *Metaphysics*, commenting on an argument similar to the argument I have just presented,

What the argument...suggests is that if we *were* unable to explain [rainbows and planetary orbits and such], this would not be mysterious. The success of physical science does not suggest that the inability of metaphysics to achieve the same sort of success is a mystery. If anything in this area is a mystery, it is the fact that science is a success. (p. 290, n 3)

I have said that I did not claim to have shown that we are unable to solve the problems of metaphysics. I will add that I did not claim to have shown that we are *probably* unable to

Concluding Meditation

solve the problems of metaphysics. My conclusion was rather this: that the considerations I appealed to showed that the probability that we are unable to solve the problems of metaphysics is high enough that, if we somehow come to learn that we are unable to solve them, we should not find that particularly surprising.

17.6 Reply to John A. Keller

I do not have the space to discuss everything John Keller says in his very rich and closely argued paper. I am going to do only two things: I will defend the criterion of philosophical success and failure I have proposed in various places (but primarily in Lecture 3 of *The Problem of Evil*⁵²), and I will say something about the alternative criterion that Keller proposes. I will proceed mainly by quoting some passages from his paper and commenting on them.

17.6.1 Analyzing Success

I begin with these two passages:

...van Inwagen's is not the most liberal possible criterion. **Philosophical individualism**—roughly, the view that an argument *A* for conclusion *c* is successful for individual *i* if and only if *A* is convincing to *i* (regardless of *i*'s previous attitude towards *c*)—is more liberal, in that it allows some arguments for substantive philosophical conclusions to be successes. (p. 305)

What makes philosophical individualism “individualistic” is that it relativizes success to individuals, holding that ‘*x* is a successful argument’, like ‘*x* is in motion’, has implicitly relativistic truth-conditions, such that an argument might be a success for you without being a success for me. Philosophical individualism relativizes the success of an argument to the person evaluating it. (p. 306)

I don't think ‘philosophical success’ has enough pre-analytic content for it to be possible to say that ascriptions of philosophical success do or do not have “truth-conditions that relativize success to individuals.” I think that there is room for

Concluding Meditation

criteria of success that relativize success to individuals and for criteria that relativize success to very large subsets of the set of all human beings who have lived or will live. (My (p.374) own criterion, as we shall see, relativizes success to historical period and culture.) Now it may be that it is impossible for a philosophical argument (for a substantive, positive conclusion) to be a success relative to almost all persons of a certain time and culture (or to any other such comprehensive class of people) and that it would therefore be advisable for us to turn our attention to the possibility of criteria of success that are relativized to individuals. I don't dispute that. My thesis about philosophical success, after all, is just exactly that there *are* no philosophical arguments (for etc.) that are successes relative to almost all persons of *our* time and culture (in the most liberal possible sense of 'success relative to almost all persons of our time and culture'). It seems to me that even if it is *impossible* for there to be an argument that is a success according to my definition, *that* such an argument is impossible is an interesting and important thesis. And, I believe, it would remain an interesting and important thesis if it could be shown that there were other standards of success, reasonably plausible and intuitive standards, according to which some philosophical arguments (for etc.) *were* successes. In any case, Keller has not presented a definition or analysis of "success relative to persons of our time and culture" that is more liberal than mine because that is not the concept he has presented a definition or analysis of. His criterion and mine are not in competition. One could consistently employ them both, although not, of course, for the same purposes. It does not follow, however, that he has not shown that my time-and-culture-relative criterion of philosophical success is, well, not a success.

17.6.2 Three Difficulties for Success_{PVI}

Keller has found three principal difficulties with my criterion of success—Success_{PVI}, let us call it. The remainder of this reply will largely be an examination of these difficulties. (In most of what follows, I will suppress reference to 'persons of our time and culture'.)

Concluding Meditation

Difficulty 1: Success_{PVI} cannot serve as a criterion of success for first-person arguments—arguments such that each individual member of any audience to whom the argument is addressed is expected to go through the argument “for himself” or “for herself.”

‘*Cogito, ergo sum*’ is an example of a first-person argument—an example that is, I think, sufficient to establish the philosophical importance of first-person arguments.

Here, I think, Keller is right: Success_{PVI} cannot, as it stands, be applied to first-person arguments.⁵³ I am, however, fairly confident that it could be revised so as to apply to them, and in such a way that the revised criterion would be very much in the spirit of the original. If one presents a first-person argument to an audience, one generally does so by going through the argument for one’s own case—that is, by using one’s own case as a model or paradigm or template—and inviting the members of one’s audience to follow one’s example. Here is an imaginary application of this dialectical method. René is lecturing to Jean-Paul and Simone. He intones, “*Je pense, (p.375) donc je suis,*” and looks expectantly at the two of them. Jean-Paul, following René’s lead, says, perhaps subvocally, “*Je pense, donc je suis*” and Simone also says, “*Je pense, donc je suis.*” Each of them, or so we may imagine, is going through a certain argument. Following David Kaplan, we may say that Jean-Paul’s argument and Simone’s argument are the same in *character* but differ in *content*. (The proposition Jean-Paul expresses by saying “*Je suis*” is true in exactly those possible worlds in which *he* exists; the proposition Simone expresses by saying “*Je suis*” is true in exactly those possible worlds in which *she* exists. Since those are distinct sets of worlds, the conclusion of Jean-Paul’s argument and the conclusion of Simone’s argument are distinct propositions.) We may say that our story involves one “character argument” and three “content arguments.” (If a content argument can be identified with a sequence of propositions, a character argument can be identified with a function from sets of indices—of some agreed-upon kind: person, moment of time, and place, it may be—to content arguments.) I propose to mend the flaw (the very real flaw)

Concluding Meditation

that Keller has found in Success_{PVI} by modifying Success_{PVI} in such a way that it applies to character arguments as well as content arguments. I will not undertake this task here, however. In the sequel, I will defend only the thesis that Success_{PVI} is, as it stands, a useful and interesting criterion of success for “third personal” philosophical arguments—that is, arguments such that every member of the audience to whom the argument is addressed is considering the same content argument.

Difficulty 2: Success_{PVI} involves the concept of an audience of whose members are *neutral* agnostics with respect to the conclusion of an argument. But *strict* neutrality with respect to a proposition is a condition that is much more demanding than I have supposed it to be; it is in fact so difficult to achieve that a criterion of success that appeals to it is bound to be unworkable.

Keller says:

There are, however, a number of problems with thinking that successful arguments must convince neutral agnostics. First, neutral agnosticism is infectious. If neutral agnostics about p only had to be agnostic about p itself, some such agnostics might firmly believe that almost all arguments for or against p are sophistical. Such agnostics would be unlikely to be convinced by any argument for or against p , successful or not.

But once we see that neutral agnostics about p must be agnostic about more than p itself, it is hard to know where to stop. Would neutral agnostics about the existence of free will have to be compatibilists? Moral realists? Consequentialists? Physicalists? Theists? Scientific realists? Even if our agnostics were strictly neutral about the existence of free will, consider how differently “primed” towards free will someone would be who was an incompatibilist, moral nihilist, atheist, physicalist, and scientific realist compared with someone who was a compatibilist, moral realist, theist, dualist, and scientific anti-realist. Of course, one can believe in

Concluding Meditation

free will as an incompatibilist, moral nihilist, atheist, physicalist, and scientific realist without manifest irrationality: Mark Balaguer accepts all of those except moral anti-realism, and there is no reason to think that adding moral anti-realism to the mix would make his position inconsistent. But convincing Balaguer to reject free will seems much easier than convincing someone who is agnostic about free will, but who is a compatibilist, moral realist, theist, dualist, and scientific anti-realist. Explicit opposition to a claim doesn't guarantee that the totality of one's views doesn't "mesh" better with that claim than does another totality that is neutral about the claim. Of course, to be truly neutral (p.376) with regard to free will one would also have to be agnostic about determinism, consequentialism, reductionism, and much else. If the members of a jury of "ideal neutral agnostics" were not agnostic about these theses, they might come to different verdicts on the basis of their differing opinions. But an audience that was agnostic about all of these would believe so little that it is hard to imagine convincing them of anything at all. (p. 308-9)

I will first consider the case of those agnostics who 'firmly believe that almost all arguments for or against p are sophisticated'. When laying out the features of the members of an ideal "audience of agnostics" in *The Problem of Evil*, I said

The audience is composed of what we may call agnostics. That is, they are agnostic as regards the subject-matter of the debate. If the debate is about nominalism and realism..., each member of the audience will have no initial opinion about whether there are universals, and no predilection, emotional or otherwise, for nominalism or for realism.⁵⁴ As regards a tendency to accept one answer or the other, they will stand to the question whether there are universals as you, no doubt, stand to the question whether the number of Douglas firs in Canada is odd or even. But that is not the whole story; for you, no doubt, have no desire to have the question whether that number is odd or even settled. My

Concluding Meditation

imaginary agnostics are not like that in respect of the question of the existence of universals. They would very much like to come to some sort of reasoned opinion about the existence of universals—in fact, to achieve knowledge on that matter if it were possible. They don't care *which* position, nominalism or realism, they end up accepting, but they very much want to end up accepting one or the other. (p. 44)

Perhaps 'firmly believe that almost all arguments for or against nominalism are sophistical' is consistent with the letter of this description. It doesn't exactly seem to be consistent with the spirit of 'They would very much like to come to some sort of reasoned opinion about the existence of universals—in fact, to achieve knowledge on that matter if it were possible'. I should have no objection to explicitly adding 'regard it as an open question whether there can be cogent arguments for *p*' to my list of the features of 'neutral agnostics with respect to *p*'. I would regard that as a friendly amendment. But suppose we *don't* add that. That will no doubt make the task assigned to the proponent of the argument more difficult—but it will not, by itself, make that task hopeless. Suppose the debaters are Atheist and Theist, that the proposition in question is 'God does not exist', and that the argument for that proposition that is being tested the argument from evil. I had been supposing that in the course of the debate before the audience of agnostics only Atheist and Theist would speak. But we could allow questions and remarks from the audience as well. (In fact, that's probably a good idea quite independently of the question whether it is possible for there to be an audience of "strictly neutral" agnostics.) Suppose a member of the audience says, or says something that implies, that any argument (p.377) either for the existence or the non-existence of God is almost certain to be sophistical. It's certainly open to Atheist to introduce a line of reasoning that opens with a remark along the lines of, "All right, but do you think that *this* argument, the argument from evil as I have presented it, is sophistical? This is a particular, concrete argument. If it is sophistical, wherein does its sophistry lie?" In other words, it's open to Atheist to try to convince the members of the audience that, whatever may be

Concluding Meditation

the case in general with arguments for the non-existence of God, *this one*, the one *she* has presented, is not sophistical.

“Yes, but in such a case, at least some members of the audience are not strictly neutral.” The Interlocutor’s question brings us to Keller’s more general argument as regards the impossibility or near-impossibility of there being an audience of “strictly neutral” agnostics. I can say only that I did not intend the audience to be “strictly neutral with respect to p ” in Keller’s very demanding sense but only in *my* sense. That is, the sense I intended is the sense I intended: what I have written, I have written. And that sense is: when the debate begins, each of them refuses to assign a higher subjective probability to the conclusion of the argument than to its denial—and vice versa. Suppose that the argument under consideration (call it ‘FW’) is an argument for the conclusion ‘Human beings have free will’. And suppose the audience comprises people who are neutral agnostics in my sense (“neutral_{PVI} agnostics”) on the question whether human beings have free will. Suppose, finally, that some of them are dualists and some are materialists. Ophelia, the opponent of the argument is able to convince the members of the audience that premises (2) and (4) of FW jointly imply dualism. Then, unless Protagoras, the proponent of the argument, is able somehow to convert the materialists in the audience to dualism or is able to convince them that (2) and (4) do not, after all imply dualism, he will find that Ophelia has effectively blocked his attempt to convince the audience (*all* the members of the audience) of the existence of free will by presenting the argument FW. But I don’t see why this case should be supposed to imply that the members of the audience who are materialists should not be described as (having initially been) neutral agnostics with respect to the existence of free will. What the case does do is provide an additional reason to suppose that it is no easy task to find an argument for a positive, substantive philosophical conclusion p that will convince an audience whose membership comprises neutral_{PVI} agnostics with respect to p to accept p . And that reason is: any given argument A whose conclusion is p may well be such that various members of that audience are convinced adherents of various philosophical theses that are inconsistent with various

Concluding Meditation

premises of *A*. I am happy to allow the following addition to my description of the audience: it is large and its members hold all possible combinations of philosophical positions that are consistent with their being neutral agnostics—neutral_{PVI} agnostics—with respect to the conclusion of the argument under consideration. (By ‘*x* holds a combination of philosophical positions consistent with *x*’s being a neutral_{PVI} agnostic with respect to *p*’ I mean this: Let ϕ be the combination of philosophical positions that *x* holds; the statement that *x* holds the combination of philosophical positions ϕ and is a neutral_{PVI} agnostic with respect to *p* is not a self-contradictory statement. I am quite willing to allow that many of the combinations of philosophical positions held by individual members of the audience are inconsistent with neutral agnosticism with respect to *p*—although, of course, the contradictions (p.378) can’t be very obvious ones, owing to the agnostics’ enjoying “the highest degree of logical acumen.”⁵⁵)

Note that it is not clear to what extent bias accounts for disagreement in philosophy, since there is persistent apparently rational disagreement about philosophical questions about which no one (or hardly anyone) has a preconceived opinion: Newcomb’s Paradox, the nature of mental representation, quantifier variance, whether properties are parts of their instances, and so on. (p. 309, n 24)

I can’t agree. Even in the most abstract areas of philosophy, we tend to accept the propositions (ideas, points of view...) that were advanced and defended by the most persuasive (charismatic, forceful, overbearing, dialectically able...) of those who taught us in graduate school. If one studied metaphysics at Notre Dame when John Keller did, one would, ten years on, be much more likely to take a Quinean meta-ontology for granted than one would be if one had been an MIT or a Pittsburgh “product”; and one would be much more likely to reject a constituent ontology than one would be if one’s doctoral studies had been at UNC or Durham University.

Consider this passage from Mark Lilla’s *The Reckless Mind: Intellectuals in Politics*:⁵⁶

Concluding Meditation

[At Marburg] Heidegger attracted students who traveled from the four corners of Europe to study with him. Although he had not yet published any works, he quickly became known throughout Germany as a genius and radical thinker simply by the force of his teaching.

It would be hard to deny that almost all the adoring students who hung on every word of “the little magician from Messkirch” would have been strongly inclined to dispute Hegel’s characterization of *Sein* as “the most barren and abstract of all categories”—and would have been so inclined for the rest of their lives. This is simply an extreme example of a very common phenomenon. (It is true that in graduate school we academics tend to be stripped of the biases we acquired as undergraduates. Nevertheless, even our undergraduate education has almost certainly had a lasting influence on our philosophical views, for the opinions of our undergraduate teachers influenced our choice of where to pursue graduate study—and the philosophical “schools” to which our letter-of-recommendation writers belonged had a lot to do with where our candidacy for admission had even a chance of success. A Marquette undergraduate is much more likely to get a PhD in philosophy from Penn State than from Princeton.)

An additional problem is that agnostics are not neutral in debates about agnosticism itself, religious or otherwise. Pyrrhonian skeptics aim only to get us to suspend belief about whether there is an external world, so someone who is agnostic about whether there is an external world already accepts the skeptic’s conclusion. When it comes to arguments that are skeptical in this sense, there is no neutral vantage point. But if anti-skeptical arguments do not need to appeal (p.379) only to premises acceptable to the skeptic, and they do not need to appeal only to premises acceptable to someone agnostic about skepticism, what is left? The only plausible answer is that they can appeal to premises that non-skeptics themselves accept. But that supports philosophical individualism over views that claim that successful arguments must utilize premises acceptable to everyone, a neutral audience, etc. (p. 309)

Concluding Meditation

Let us consider the case of the Skeptics. (I will not call them Pyrrhonian skeptics for historical reasons that I will not go into.) Let us in fact consider two particular skeptics, Pyrrho and Carneades. Each is standing before an audience comprising persons of our time and culture⁵⁷ and each wishes to bring it about that the members of his audience accept neither the proposition that there is an external world nor the proposition that there is no external world. But how do they attempt to do this? There would seem to be two dialectical procedures such a skeptic might employ. Pyrrho chooses one and Carneades the other.

Dialectical procedure 1: Pyrrho will proceed by presenting the members of his audience (*AP*) with irrefutable arguments for the existence of an external world *and* irrefutable arguments for the non-existence of an external world.⁵⁸ (Let us call these two propositions the Positions.) And he will, as it were, allow these arguments speak for themselves. He will present no final, comprehensive “the lesson of these arguments is” argument that refers to them (as Kant did with respect to his Antinomies of Pure Reason) and whose conclusion is a proposition along the lines of ‘Both Positions are without epistemic warrant’: He will not go on to say, “Since, as we have seen, there are irrefutable arguments for both Positions, both Positions are therefore without epistemic warrant.”

A philosopher who, like Pyrrho, employs dialectical procedure 1 presents no problems for Success_{PVI} , owing to the fact that that procedure involves no argument that is intended to convince those to whom it is addressed of the truth of its conclusion. It is, in fact, essential to Pyrrho’s strategy that none of the arguments the skeptic employs should, in the end, convince any member of *AP* that its conclusion is true—owing to the fact that, for each argument he presents, he will have presented at least one equally convincing argument for the denial of its conclusion.

Dialectical procedure 2: Carneades will proceed by presenting the members of his audience (*AC*) with an argument whose conclusion is a certain epistemological proposition. (Or more than one argument. But let us make the case as simple as possible and suppose he presents only one argument.) Let us

Concluding Meditation

call that argument *SKA* and let us suppose that the conclusion of *SKA* is

Skepticism Both Positions are without epistemic warrant.

(p.380) Note that, unlike the arguments Pyrrho presented, which were not intended to convince the members of *AP* that their conclusions were true, this argument *is* intended to convince the members of *AC* that its conclusion is true. But *should* *SKA* convince them that its conclusion is true? Suppose I say that that *SKA* should convince the members of *AC* that everyone should suspend judgment about whether there is an external world just in the case that *SKA* is a *Success_{PVI}*. Does that mean that *AC* must be entirely composed of persons who *neutral_{PVI}* agnostics about the truth or falsity of *Skepticism*? No, not at all. What it implies is this:

SKA should convince the members of *AC* that *Skepticism* is true just in the case that:

If presented in IDEAL CIRCUMSTANCES⁵⁹ *SKA* would convince the members of an audience of initially *neutral_{PVI}* agnostics (about the truth or falsity of *Skepticism*) that *Skepticism* is true.

Here is a strictly parallel statement (let '*MAT*' be any argument for materialism, understood as the thesis that human beings are wholly material):

MAT should convince Robert M. Adams and Richard Swinburne and Alvin Plantinga that human beings are wholly material just in the case that:

If presented in IDEAL CIRCUMSTANCES *MAT* would convince the members of an audience of initially *neutral_{PVI}* agnostics about the truth or falsity of materialism that human beings are wholly material.

Each of these two biconditionals makes reference to *two* audiences: the former to *AC* and an audience of agnostics about *Skepticism*, and the latter to Adams, Swinburne, and

Concluding Meditation

Plantinga (on the one hand) and an audience of agnostics about materialism (on the other).

Let us suppose that *INA* is a representative or typical audience of “initially neutral_{P_{VI}} agnostics about the truth or falsity of *Skepticism*.” It may well be that each of the members of *INA* must—in order to exhibit the appropriate species of agnosticism about *Skepticism*—accept neither Position. (For it seems plausible to suppose that there is some sort of pragmatic contradiction involved in saying anything of the form, ‘*p*, and my belief that *p* is without epistemic warrant’.) Let us suppose that that is so. Has it any consequences that are relevant to the question whether Success_{P_{VI}} is a useful or interesting criterion of philosophical success? It certainly does not follow that, since *AC* is presumably composed mostly of people who believe that there is an external world, Carneades should not present *SKA* to the members of *AC* with the intent of demonstrating to them that both their Position and its denial are without epistemic warrant. And it would not follow that ‘*SKA* is a Success_{P_{VI}}’ is not sufficient for ‘*SKA* should convince the members of *AC* that both Positions are without epistemic warrant’. There would, of course, be no point in Carneades’ presenting *SKA* to an audience whose members were neutral_{P_{VI}} agnostics about the existence of an external world in order simply to convince them neither to accept nor to reject the proposition that there is an external world, since they “already” neither accept nor reject that proposition. That fact, however, does not (p.381) prevent us from “testing” *SKA* by asking whether it could be used to cause an audience of initially neutral_{P_{VI}} agnostics about the truth of the proposition *Skepticism* to accept that proposition. As I have said, it is plausible to suppose that an audience that was agnostic about *Skepticism* would also have to be agnostic about whether there was an external world. But an audience that was agnostic about whether there was an external world could certainly *also* be agnostic about whether both ‘There is an external world’ and ‘There is no external world’ are without epistemic warrant. There is no contradiction, pragmatic or otherwise, in my saying something of the form:

Concluding Meditation

I accept neither p nor its denial; and I accept neither the proposition that p is without epistemic warrant nor the proposition that the denial of p is without epistemic warrant.

Consider, for example, the proposition that the power of the continuum is either aleph-one or aleph-two. I accept neither that proposition nor its denial. It has been reported that Gödel believed that the power of the continuum was either aleph-one or aleph-two. Let us suppose that that report is correct. I do not believe (do not have the belief that) Gödel's belief was without epistemic warrant. (For all I know, his belief was based on some profound mathematical insight.) Nor do I have the belief that the proposition that the power of the continuum is neither aleph-one nor aleph-two is without epistemic warrant. (For all I know, Gödel's "profound insight" was non-veridical—even Homer nods—, and the proposition that the power of the continuum is neither aleph-one nor aleph-two is true. *And*, for all I know, some other mathematician has had a veridical mathematical insight that revealed its truth to him or her.) Therefore:

I accept neither the proposition that the power of the continuum is either aleph-one or aleph-two nor its denial; and I accept neither the proposition that (the proposition that the power of the continuum is either aleph-one or aleph-two is without epistemic warrant) nor the proposition that (the proposition that the power of the continuum is neither aleph-one nor aleph-two is without epistemic warrant).

Difficulty 3: Success_{PVI} is incompatible with (subjective) Bayesianism.

Keller says,

Bayesians hold that what is rational to conclude on the basis of new arguments or evidence depends on one's priors. But people have different priors, and according to (subjective) Bayesian orthodoxy no set of coherent priors is more rational than another. Hence, two people exposed to the same evidence and arguments may reach

Concluding Meditation

different conclusions. Indeed, if they are rational, they *must* reach different conclusions, each in accordance with her priors. But then there can be no objective fact about whether an argument makes acceptance of its conclusion rational, and hence whether the argument is successful. For this will depend on one's priors, and priors are individualistic. Hence, so is argumentative success. [The following footnote is appended to this passage: 'Yes, priors can often be "swamped" by the evidence....But they are not always swamped, and so as long as we hold that any consistent set of priors is rational—or even a wide range of consistent priors—exposure to the same arguments will not necessarily lead to agreement, even between people who respond to evidence perfectly.'] (pp. 311-12)

(p.382) If this argument is cogent, it tells against the possibility of there being any argument in any field of enquiry that is a non-individual-relative success. This one, for example: 'Only a ball casts a circular shadow from all angles; hence, if the boundary of the shadow the earth casts on the moon during a partial lunar eclipse is always an arc of a circle, then the earth must be a ball.' Or consider currently available arguments for the conclusion that smoking cigarettes causes lung cancer or for the conclusion that the continents are in motion. (I don't mean the arguments that have convinced epidemiologists and geologists that these propositions are true; I am thinking of arguments addressed to the general public that rest, at least in part, on appeals to the testimony of such experts.) These arguments are successes for (almost) all of us, for all us twenty-first-century adult human beings who are members of a European or "European-descended" culture—or almost all. (Remember that in my account of the audience of agnostics, I said, "I mean the agnostics to be drawn from our time and our culture."⁶⁰) Almost all adult persons of our time and our culture share certain priors—and a very large and rich class of them. Almost everyone "of our sort" assigns a subjective probability of 1 to 'Some people live in houses' and 'Lions are carnivores' and 'It is colder at the poles than at the equator'.⁶¹ Almost everyone assigns a subjective probability of 0 to 'Butter is a high explosive' and 'Tomatoes are imported

Concluding Meditation

from the moon' and 'Bacteria are capable of learning calculus'. Almost everyone assigns a subjective probability of 0.5 to 'When Vladimir Putin dies, he will have had an odd number of birthdays'. If these rather fanciful examples are not examples of *true* priors—owing to the fact that most people who accept them accept them on the basis of the probabilities they assign to other propositions—, their near-universality implies that there must be *some* probability assignments that constitute the priors of almost every adult of our time and culture.⁶² I think we may speak of the “core priors” of the (adult) members of that class of people, the priors that they (almost) all share. (Naturally, there are bound to be borderline or disputable cases of membership in this class: perhaps there is no determinate answer to the question whether 'If a fair coin has fallen heads ten times in a row, the probability that it will fall heads on the next toss is 0.5' is a core prior.) I am happy to relativize my account of success to “success with respect to those whose set of priors (whose *consistent* set of priors) includes the core priors of our time and culture.” My point, after all, was that no philosophical argument for a positive, substantive conclusion is an “impersonal” success in the way 'Only a ball casts a circular shadow from all angles; hence, if the boundary of the shadow the earth casts on the moon during a partial lunar eclipse is always an arc of a circle, then the earth must be a ball' is an impersonal success. It is therefore essential to my purposes that I should appeal to a criterion of success that allows some arguments to be impersonal successes in that way.

(p.383) 17.6.3 A Difficulty for Success_{JK}

This concludes my discussion of the three difficulties that, in Keller's view, confront my position on dialectical success and failure in philosophy. I will close this reply by mentioning a difficulty I see in Keller's own individual-relative account of success, the “Knowledge Account”:

An argument is successful for an individual just if she *knows* it is sound (and non-fallacious).

I do not think this will do. I know that Genuine Modal Realism is not an acceptable modal ontology, and I know that Alvin

Concluding Meditation

Plantinga has said that Genuine Modal Realism is not an acceptable modal ontology. I know that *modus ponens* is a valid argument-form. I therefore know that the following argument is sound:

Alvin Plantinga has said that Genuine Modal Realism is not an acceptable modal ontology

If Alvin Plantinga has said that Genuine Modal Realism is not an acceptable modal ontology, then Genuine Modal Realism is not an acceptable modal ontology⁶³

hence,

Genuine Modal Realism is not an acceptable modal ontology.

But I can see no sense in saying that this argument is a success “for me.”⁶⁴ If our *definiendum* is

A is a success for *i*,

I propose that the following clause, or something very much like it, appear somewhere in the *definiens*:

i knows the premises of *A* to be true, and would know those premises to be true whether she knew its conclusion to be true or not.⁶⁵ (p.384)

17.7 Reply to Thomas Kelly and Sarah McGrath

Thomas Kelly and Sarah McGrath’s paper is stimulating, thoughtful, clear, and closely reasoned. Because it is stimulating and thoughtful, I am moved to reply to it at some length. Because it is clear and closely reasoned, it is easy for me to see on what points I disagree with them about success and failure in philosophical argument.

In *The Problem of Evil*, I gave a definition of what it is for a philosophical argument to be a failure. And I presented an argument for the conclusion that all the members of a very broad class of philosophical arguments were, in the sense provided by that definition, failures. (If that argument is sound, its soundness implies that it is itself a failure: its

Concluding Meditation

conclusion is a substantive metaphilosophical thesis, and substantive metaphilosophical theses are substantive philosophical theses. I don't mind that. There are many *good* philosophical arguments in the "broad class" to which I have alluded.⁶⁶ On pain of pragmatic inconsistency, I can't hope to produce a *compelling* philosophical argument—so to call a philosophical argument that is not a failure⁶⁷—for that conclusion, but I think my argument *is* a good argument.)

The definition was framed in terms of an ideal debate between a defender of the conclusion of an argument and an opponent of its conclusion (Desmond and Ophelia, let us call them) before an audience of ideal, neutral agnostics (agnostics with respect to the conclusion of the argument)—an audience whose members are, in Kelly and McGrath's happy phrase, *disinterested but not uninterested* with respect to the conclusion of the argument. Very roughly speaking, the argument in question will be a failure just in the case that Desmond will ultimately be unable to use that argument to turn the agnostics into believers like himself—will be unable to use the argument to convince them to accept its conclusion. I say 'ultimately' because he must do this in the presence of Ophelia, his opponent in the debate, whose dialectical function is to do everything possible to block his attempt at "conversion." (This statement of the definition of failure is intended only to serve as a reminder of its content. My "official" statement of the definition remains the one presented in the book.)

Before I proceed to a discussion of Kelly and McGrath's essay, I wish to make two comments on the definition.

(1) The definition was intended to apply only to arguments for "controversial" conclusions—an *uncontroversial* conclusion being an argument that practically everyone would accept. (Not an entirely straightforward concept, I concede: Berkeley thought that practically everyone accepted his conclusions—whether they knew it or not; most philosophers reject this judgment.) Consider, for example, arguments against solipsism. Few philosophers, if any, who construct such arguments construct them with the intention of convincing the generality of humanity that solipsism is false. The purpose of

Concluding Meditation

those arguments must be something other than that (perhaps it is to display the basis of our “knowledge of other minds”; perhaps each of them is an attempt to provide an explicit statement of an argument that—its author supposes—almost everyone accepts tacitly). The criterion of success and failure for arguments (p.385) against solipsism or for the reality of change and motion or for the existence of truths about the past must, in my view, be of a different sort from the criterion of success and failure for arguments for (and arguments against) consequentialism or materialism or mathematical platonism. My definition is intended to provide a criterion of success and failure only for arguments of the latter sort.

In connection with this point, I will distinguish two senses in which a philosophical argument may be said to be an argument *for* a certain philosophical proposition ϕ . This may mean no more than that ϕ is the conclusion of the argument. But it may mean that and more. It may also mean that the argument has been constructed with an intended audience in mind, an audience of philosophers (and perhaps non-philosophers as well) who do not accept ϕ (who either accept the denial of ϕ or who accept neither ϕ nor its denial) and which was constructed for the purpose of convincing that audience to accept ϕ . I will use “argument for” in the latter, stronger sense.

(2) If I were given the opportunity to revise the account of philosophical success and failure that I presented in *The Problem of Evil*, I would no longer say that an argument was a failure just in the case that Desmond (the defender of its conclusion) “will ultimately be unable to use that argument to turn the agnostics into believers like himself”; I would say, rather, that an argument is a failure just in the case that he will ultimately be unable to do *either* of the following two things:

- turn the agnostics into believers like himself
- turn the agnostics into *positively weighted agnostics*.

And this is what italicized phrase means:

Concluding Meditation

At the beginning of the debate, the agnostics were *neutral* agnostics: either they assigned a subjective probability of 0.5 to the conclusion of the argument (and of course to its denial) or they refused to assign any subjective probability to either the conclusion or its denial (nor were they willing to ascribe to those propositions membership in any *range* of subjective probabilities—other than $[0,1]$, of course). A positively weighted agnostic (with respect to a proposition p) is a person who accepts neither p nor its denial and who either assigns to p a particular probability significantly higher than 0.5 or who (more likely) is willing to say that the probability of p , whatever exactly it may be, is significantly higher than 0.5.⁶⁸

Suppose, for example, that the outcome of the debate imagined in the discussion of the global argument from evil in *The Problem of Evil* was this: the formerly neutral agnostics all say something along the lines of, “We still don’t accept either atheism or theism, but now—now that the debate is over, now that both Theist and Atheist have said everything they have to say about the global argument from evil—we have been convinced by that argument to assign a significantly higher probability to atheism than to theism.” If that is what the outcome of the debate would be, then the argument would be a failure according to the definition in *The Problem of Evil* but a success according to the revised definition.⁶⁹

(p.386) In the sequel, however, I’ll keep to the original definition, if only because that is the definition that Kelly and McGrath were working with (and they have not criticized the definition on the ground that it entails that an argument with the power to turn neutral agnostics into positively weighted agnostics—but not into “believers”—would count as a failure).

I turn now to various matters that pertain to Kelly and McGrath’s critique of my argument for the Pessimistic Verdict.

17.7.1 Uniformity of Response?

In *The Problem of Evil*, I assumed “uniformity of response” from the audience of agnostics. That is, I assumed that, when the debate was finished, either they would all be convinced by

Concluding Meditation

the argument that had been advanced and would all, therefore, accept its conclusion, or they would all say, "For all I know, at least one of the premises of that argument is false."⁷⁰ (I conceded that this was an assumption I was making, and said that it would be interesting to look into the consequences of proceeding without it.) It still seems pretty reasonable to me, given the high degree of idealization involved in the debate. But, rather than defend it, I will look into the consequences of proceeding without it.

On consideration, I'm willing to accept the definition without the assumption. My reasoning is this. Suppose you did present a certain argument α for, let's say, the existence of abstract objects or platonism, to an audience of ideal agnostics in the context of an ideal debate between a platonist and a nominalist (a denier of platonism). Suppose the response from the audience of agnostics is, as they say, all over the shop. Some are convinced by α and become platonists. Others, remain agnostic on the question of the existence of abstract objects on the ground that that *for all they know*, some of the premises of α are false (some say premises (3) and (6) are, for all they know, false; others find those two premises luminously evident but are doubtful about premise (4)...). Perhaps a few of them remain agnostic owing to their conviction that one or two of the premises of α are simply false, false without qualification, and self-evidently so. Suppose further, that when α is presented (in the context of an ideal debate) to different (sizable) groups of ideal neutral agnostics, the non-uniform outcomes vary widely: in Group A, 32 percent remain agnostics and the remainder are convinced by α and become platonists; in Group B, it's 48 percent and 52 percent....

That certainly seems to me to be an excellent reason to say that α is a failure. Or put the matter this way. It seems to me that anyone who accepts the proposition

If, when α is presented to an audience of ideal neutral agnostics in the context of an ideal debate, they all, to a man and to a woman, respond by saying, "For all I know, some of the premises of α are false," then α is a failure.

(p.387) should also be willing to accept the proposition

Concluding Meditation

If, when α is presented to an audience of ideal neutral agnostics in the context of an ideal debate, their responses vary widely—are, in fact, all over the shop—then α is a failure.

That is to say, an argument (for a controversial conclusion) is a success only if, if it were presented to an audience of ideal neutral agnostics in the context of an ideal debate, the members of that audience would *uniformly* respond by accepting its conclusion. (We could weaken that principle, perhaps—say to ‘*almost uniformly* respond’. But we should, I believe, certainly count an argument as a failure if we somehow knew that the “ideal response” would be *very far from uniform*.)

17.7.2 Logical Omniscience?

Kelly and McGrath correctly identify a problem that would confront my definition of philosophical success and failure if my ideal agnostics were “logically omniscient”—if, for any proposition they accepted, they immediately saw, and accepted, all its logical consequences. Fortunately, I do not have to deal with that problem, since I did not say anything that implied that they were logically omniscient. I did want my agnostics to be human beings, and—obviously I would say—human beings can consider or grasp or hold before their minds only an infinitesimal proportion of the logical consequences of any proposition. When I said that the agnostics possessed “the highest degree of logical acumen,” I meant only that, for any argument Desmond or Ophelia might present them with (it would have to be simple enough for a human being actually to understand—but then Desmond and Ophelia are human beings, too, and they are unlikely to present the agnostics with arguments so complex that they themselves are unable to understand them), they would, given a reasonable amount of time to consider the matter, be able to determine, infallibly, whether that argument was logically valid. Few would deny “the highest degree of logical acumen” to Gottlob Frege, and yet for many years he was unaware of the fact that the three propositions (the language of this example is anachronistic)

Concluding Meditation

For any one-place open sentence x such that, for every object y , either y determinately satisfies x or y determinately fails to satisfy x , there exists exactly one set whose membership comprises exactly those objects that satisfy x

For every object y , either y determinately satisfies 'x is a member of x ' or y determinately fails to satisfy 'x is a member of x '

For every object y , either y determinately satisfies 'it is not the case that x is a member of x ' or y determinately fails to satisfy 'it is not the case that x is a member of x '

logically implied the proposition

There exists a set x such that x is a member of x and it is not the case that x is a member of x .

And yet, when Russell presented him with the argument whose premises were (essentially) those three propositions, and whose conclusion was (essentially) the fourth proposition, he immediately saw—owing to his possessing logical acumen in the highest degree—that that argument was logically valid.

I meant only to imply that my agnostics possessed Frege's sort of logical acumen.

(p.388) 17.7.3 What is a Positive Thesis?

Kelly and McGrath raise questions about what a positive thesis is, and whether, by any reasonable account, atheism is a positive thesis. (The latter point is relevant to the use I made of the Pessimistic Verdict in *The Problem of Evil*, where I suggested that it would not be surprising to one who accepted the Verdict if the argument from evil turned out to be a failure.) On this point, I will quote a scattering of passages from a paper Kelly and McGrath do not consider in their essay.⁷¹ What is said in these passages will not perhaps satisfy those who think that the concept of a positive thesis is insufficiently well defined to be a load-bearing member in a serious philosophical argument (or satisfy those who insist that an account of 'positive thesis' that had the consequence

Concluding Meditation

that atheism was a positive thesis would be ipso facto objectionable), but it is the best I can do. The first passage is,

What it means to say that a belief (proposition, thesis, conjecture, theory, hypothesis...) is positive or negative is hard to explain in any philosophically satisfactory way, and I will not attempt to do so. I shall have to be content to give a few examples of philosophical beliefs or propositions that are paradigmatically not positive: 'Formalism is not the correct philosophy of mathematics'; 'Utilitarianism is not an acceptable ethical theory'; 'Knowledge is not simply justified true belief.' And, by the same token, 'Knowledge is justified true belief,' although it is no doubt a false thesis, is a positive thesis, and to assent to it is to have a positive philosophical belief. Formalism and utilitarianism—assuming that these terms have been sufficiently well defined that they denote particular propositions—are positive theses, and anyone who accepts formalism or accepts utilitarianism thereby has a positive belief. (p. 16)

A few sentences later, speaking of the theses 'Change and motion are not real features of the world', 'One has no reason to suppose that there are minds other than one's own', and 'There are no material objects', I said,

All these theses, or my statements of them, contain some sort of negative construction. Nonetheless, all of them are what I would call 'positive' theses. As I said, 'positive' is a very hard term to explain. (pp. 16-17)

And to these two sentences I added a footnote (footnote 3, p. 17):

I would say that the negation of a negative belief must be a positive belief, but that the negation of a positive belief will in some cases also be a positive belief. An analogy is perhaps provided by the concept of positive and negative geographical information. That the spy whose whereabouts we should like to know is not in London is a negative piece of geographical information, and that he

Concluding Meditation

is in London is a positive piece of geographical information. That he is in the Western Hemisphere is a positive piece of geographical information, but so is the information that he is not in the Western Hemisphere—at least given that he must be either in the Eastern or the Western Hemisphere—, for the latter piece of information narrows down our range of possible specific hypotheses as to his location precisely as effectively as its negation does. I might put my point this way: ‘Theism is false’ is a positive philosophical belief because both theism and its negation, atheism, are philosophical theories or at any rate philosophical (p.389) positions. ‘Utilitarianism is false’ is not a positive philosophical belief because [the negation of utilitarianism] is not a philosophical theory or position. There are many philosophical theories—many ethical theories—that are incompatible with utilitarianism, but [the negation of utilitarianism] or the disjunction of all ethical theories (indeed, of all propositions) incompatible with utilitarianism, is not one of them: it’s incompatible with utilitarianism all right, but it’s not an ethical theory—and not a theory of any sort.

17.7.4 The Pessimistic Verdict

I turn now to the Pessimistic Verdict. I prefer to state it this way:

Every known⁷² argument for any substantive, positive philosophical thesis is a failure.

(Here, of course, the sense of ‘failure’ is to be provided by my definition. I do not deny that there may be other senses of ‘failure’, eminently defensible senses, in which some arguments for substantive philosophical theses are not failures. Anyone who thinks that this is so and who thinks the point an important one may feel free to substitute some other word or phrase for ‘failure’ in my statement of the Pessimistic Verdict; ‘failure_{PV1}’ perhaps. In philosophy, we must at all costs avoid *verbal essentialism*.)

But I am, just for the purposes of the present reply, going to make one modification of the Verdict.

Concluding Meditation

Kelly and McGrath rightly raise the question of how the fact that some arguments for philosophical conclusions have empirical premises (propositions whose truth-values can be determined only by observation and experiment) factors into the question of what it is for an argument to be a success or a failure. This is a very good question, and no doubt it is one I should have paid more attention to than I did. (I did remark that it was conceivable that an argument that was a success in one century might be a failure in another, owing to advances in empirical knowledge.) But for present purposes, I am going to do an end run around this very good question. I will accomplish this end run simply by replacing the Pessimistic Verdict with a pessimistic verdict whose scope is restricted to arguments *a priori*. (An ‘argument *a priori*’, for present purposes, is an argument none of whose premises is an empirical proposition.) And here is

The Restricted Verdict: For any empirical proposition p , and any substantive, positive philosophical thesis ϕ , every known argument *a priori* for the conditional proposition whose antecedent is p and whose consequent is ϕ is a failure.⁷³

I will use an example to explain the way in which I mean to apply the Restricted Verdict. Suppose that a philosopher offers an argument for the existence of a first mover, and that one of the premises of this argument is ‘The age of the cosmos is finite’—which I hope everyone will agree is an empirical proposition. (If you don’t [\(p.390\)](#) think it is, change the example.) Suppose further that this is the only premise of that argument that is an empirical proposition. Suppose, finally, that the truth of that premise has been established by empirical investigations that only people who have mastered some very difficult mathematical physics can understand. If we imagine Desmond presenting this argument to an audience of ideal agnostics in an attempt to convince them of the existence of a first mover, must we, in attempting to anticipate their reaction to the argument, suppose that they are all trained mathematical physicists—and Desmond and Ophelia as well? If not, how are we to apply our definitions of success and failure to this argument? I have no good answer to these

Concluding Meditation

questions. But *if* the Restricted Verdict is correct, I don't need to apply these definitions to his argument—not to *that* argument. Let me explain.

Although Desmond's argument does not fall within the scope of the Restricted Verdict, there is an argument that bears an important and intimate relation to his argument that does fall within its scope, an argument whose conclusion is the conditional proposition 'If the age of the cosmos is finite, then there is a first mover'. If the premises of Desmond's argument for the existence of a first mover were A, B, 'The age of the cosmos is finite', and D, the "new" argument will have premises A, B, and D (none of them empirical propositions) and the conclusion 'If the age of the cosmos is finite, there is a first mover'. We may call an argument obtained by this method from a philosophical argument with empirical premises the "*a priori* conditional transform" of the original argument—or, for short, its Transform. The Restricted Verdict implies that if an argument for a substantive, positive philosophical thesis has empirical premises, its Transform is a failure. And if the Transform of such an argument is a failure, so must the original argument be. Intuitively, an argument some of whose premises are empirical propositions may be said to have an empirical part and an *a priori* part, and it will be a success only if the agnostics find both parts convincing. The Restricted Verdict implies that *whatever* may be the case with the empirical part of the argument, the agnostics will find the *a priori* part unconvincing: *no* philosophical argument for a substantive, positive conclusion will be a failure *only* because it has empirical premises that the agnostics find unconvincing. *If*, therefore, the Restricted Verdict is correct, the original Pessimistic Verdict is also correct. Thus the "end run": if the Restricted Verdict is correct, the problem of specifying the methods Desmond is to use to convince the agnostics of his empirical premises is irrelevant to the question of the truth or falsity of the Pessimistic Verdict—for, however that problem is solved, the Pessimistic Verdict is correct if the Restricted Verdict is correct.

I hereby endorse the Restricted Verdict. In the sequel, however, I shall defend my argument for the original

Concluding Meditation

Pessimistic Verdict. That, after all, is the argument that Kelly and McGrath have criticized. And my argument for the original Pessimistic Verdict could easily be modified to produce a parallel and equally good (or equally bad, as the case may be) argument for the Restricted Verdict.⁷⁴

(p.391) And here is the argument for the original Pessimistic Verdict:

Suppose that there were an argument α for a substantive, positive philosophical thesis ϕ , an argument that was well known to philosophers, such that:

If α were presented to an audience of ideal neutral agnostics (*sc.* about the truth or falsity of ϕ) in the context of an ideal debate about whether ϕ was true or false, then, when the debate had run its course, they, the agnostics would *uniformly* agree that α (considered in the light of everything said for and against it during the debate) had led them to assent to ϕ .

If there were indeed a known argument that had that counterfactual property, then a *significantly* higher proportion of philosophers would accept ϕ than actually do accept ϕ . (This conditional statement, of course, is the premise on which the argument turns.) Therefore, no known argument has this feature. There may be arguments subsisting in the Platonic heaven (or even arguments actually written down somewhere—perhaps in Berkeley’s lost manuscript on ethics) that have it, but no argument with a substantive, positive conclusion that is known to philosophers has the power to produce uniform acceptance of its conclusion among an audience of ideal agnostics when it is presented to them in the context of an ideal debate. That is to say: Every known argument for any substantive, positive philosophical thesis is a failure (a *failure_{PVI}*, if you like).

Now what about that premise? Let us consider an example. According to the study David Chalmers cites in his essay in the present volume (the 2009 PhilPapers Survey), 39 percent of us analytical philosophers accept platonism, and 38 percent

Concluding Meditation

accept nominalism. (Or at any rate, this is how matters stood in 2009. I will assume that these figures are still accurate.) Now suppose that there exists, Platonically speaking, an argument α for platonism that has never been discovered in the actual world. α , moreover, is not an argument that only God or an archangel or the Wise Old Beings from the Galactic Core could comprehend: α is an argument that *might* have been discovered by some human philosopher (and it could have been discovered without philosophy's having gone in some entirely different direction; α employs only concepts and logical methods that have actually been available to analytical philosophers since, let us say, the middle 1960s). Let us suppose, in fact, that it is simple enough that it could successfully be taught to serious undergraduate students of philosophy. And suppose that α has the conditional property with respect to determining the outcome of an ideal debate about the truth-value of its conclusion described above—that it is a philosophical *success*.

If all that is true, the premise endorses the following counterfactual thesis:

Consider the closest⁷⁵ possible worlds that satisfy the following conditions:

In those worlds, α was discovered and published in 1970; in them, the state of analytical philosophy in 1970 was as much like the way it actually was in 1970 as is consistent with the discovery and publication of α in that year; by 1975, α was well known to all analytical philosophers interested in ontological questions.

(p.392) In those worlds, the following proposition is true:

Now, in the teens of the twenty-first century, the proportion of analytical philosophers who assent to platonism is *significantly* higher than 39%.

And isn't that counterfactual thesis at least highly plausible? Certainly the discovery of an argument with such marvelous powers of rational persuasion would not have the consequence that the proportion of philosophers who accept platonism is

Concluding Meditation

lower than it actually is, would it? And if α indeed has what we might call the *intrinsic rational power* to produce uniform assent to platonism among a population of ideal and initially neutral agnostics in the context of an ideal debate, can we really suppose that its having become known to philosophers—more-or-less-as-they-actually-were-in-1970 would not have significantly raised the proportion of those philosophers and their “professional descendants” that were platonists? Were philosophers active in the 1970s immune to reason to that degree? Were they then—and did they continue to be—so wedded to the convictions that were theirs before they encountered α that careful consideration of α would have no power to move them? And, if they were so wedded to their prior convictions, what about those “descendants” of theirs? What about the people whose philosophical formation included exposure to α , who first encountered the argument as “innocents”—when they were undergraduates—and who later became professional philosophers? (In the closest worlds that satisfy the above conditions, the majority of philosophers professionally active “in the teens of the twenty-first century” first encountered α when they were undergraduates.)⁷⁶

I have presented an argument for the Pessimistic Verdict. The argument, of course, is a failure. It is a failure by my standards, not to mention the more stringent standard devised, if not endorsed, by Alvin Plantinga, and that perhaps “maximally stringent” standard that, according Robert Nozick, he accepted as a young man.⁷⁷ It is not a “compelling” argument.⁷⁸ I think, however, that it is a good argument, an argument that should lead philosophers to take its conclusion seriously. And, in my view, that is the best that can be hoped for from a philosophical argument for a substantive, positive conclusion.

(p.393) References

Bibliography references:

Dennett D., and Steglich-Petersen, A. (2008), *The Philosophical Lexicon*, <<http://www.philosophicallexicon.com>>.

Concluding Meditation

Lewis, D. (1986), *On the Plurality of Worlds* (Oxford University Press).

Lilla, M. (2001), *The Reckless Mind: Intellectuals in Politics* (New York Review of Books).

McGinn, C. (1993), *Problems in Philosophy: The Limits of Inquiry* (Blackwell Publishers).

Paul, L. A. (2012), 'Metaphysics as Modeling: The Handmaiden's Tale', in *Philosophical Studies* 160: 1-29.

Van Inwagen, P. (1988), 'The Magnitude, Duration, and Distribution of Evil: A Theodicy', in *Philosophical Topics* XVI, no. 2: 161-87.

Van Inwagen, P. (1996), 'Review of *Problems in Philosophy* by Colin McGinn', in *The Philosophical Review* CV: 253-6.

Van Inwagen, P. (2003), 'The Compatibility of Darwinism and Design', in N. A. Manson (ed.), *God and Design: The Teleological Argument and Modern Science* (Routledge), 348-63.

Van Inwagen, P. (2004), 'A Theory of Properties', in D. Zimmerman (ed.), *Oxford Studies in Metaphysics*, Vol. 1: 107-38.

Van Inwagen, P. (2006), *The Problem of Evil* (Oxford University Press).

Van Inwagen, P. (2009a), 'Listening to Clifford's Ghost', in A. O'Hear (ed.), *Conceptions of Philosophy: Royal Institute of Philosophy Supplement 65* (Cambridge University Press), 15-35.

Van Inwagen, P. (2009b), 'Weak Darwinism', in L. Caruana (ed.), *Darwin and Catholicism* (T&T Clark), 107-20.

Van Inwagen, P. (2011), 'Relational vs. Constituent Ontologies', in *Philosophical Perspectives* 25: 389-405.

Van Inwagen, P. (2014), 'What is an Ontological Category?', in *Existence: Essays in Ontology* (Cambridge University Press).

Concluding Meditation

Van Inwagen, P. (2015), *Metaphysics*, 4th edn. (Westview Press). (p.394)

Notes:

(¹) I heartily agree with Loux's contention (p. 12) that there is no such thing as *the* problem of universals—and with his arguments in support of this contention. I would add that in my view, this contention is a special case of what I take to be a general law of philosophical terminology. Some other instances of this law: there are no such things as *the* mind-body problem, *the* problem of free will (and determinism), *the* problem of evil, *the* problem of other minds, *the* problem of knowledge (and skepticism), *the* (hard) problem of consciousness....In general: no phrase of the form 'the problem of X' that is in general philosophical use is actually a proper definite description. In every case, careful analysis of the texts in which treatments of the supposed problem are commonly supposed to be found shows that these texts comprise discussions of a family of problems that are in many cases closely interrelated but certainly distinct. (And, for any such family, it is rare for a given work that discusses some of them to discuss all of them.)

(²) Van Inwagen (2014).

(³) Or, better, in Wolfgang Pauli's fine phrase, they are *not even wrong*.

(⁴) Or of indeterminate truth-value: neither determinately true nor determinately false.

(⁵) Or of indeterminate application to that thing: neither determinately true nor determinately false of that thing.

(⁶) I trust that the reader has already read the essay on which I am commenting and thus has no need for a more complete description of the Pauline ontology. (That essay also contains answers to such "obvious" objections to its author's ontology as the following. According to Paul, parthood is not transitive. For suppose that Amber the electron is a part of Hydra the hydrogen atom; if Paul is right, Amber is a fusion of various properties, among them a part of the universal "negative

Concluding Meditation

charge.” But Hydra is not negatively charged and thus does not have a part of negative charge as a part.)

(⁷) I trust I need not explain this technical term or its foil ‘property_{PVI}’; I promise to employ them only when the sense of the word ‘property’ is not made clear by its context.

(⁸) Or perhaps I should say ‘in a strict and mereological sense’, for a color is not a part of an apple in the same sense of ‘part’ as that in which a seed is a part of an apple.

(⁹) For the whole story, see van Inwagen (2004).

(¹⁰) I mean...I mean how *could* you see a property—on *anybody’s* understanding of ‘property’? You can only see things if they absorb electromagnetic radiation and re-emit it. You can see a book or an apple because they, or parts of their surfaces, do just that. They can do that because they are made of atoms and molecules that have electronic structures: a photon kicks an electron momentarily into a more energetic atomic or molecular orbital, and another photon is emitted when the electron falls back into a lower-energy orbital; some of these emitted photons interact with the visual apparatus of the observer. Properties (on anybody’s account of properties) aren’t composed of atoms and molecules, are they? Granted, the *fact* that a certain book or apple has the property greenness (the fact that it’s green) has its formal cause in the electronic structure of the atoms and molecules that compose its surface layers, but that doesn’t mean that one can *see* the property. One can see *that* the book or apple is green, of course, but that doesn’t mean that one sees the *property*—any more than the fact that an appropriately placed observer can see that Laurie Paul is a biped means that that observer can see the number 2. “But one can see the property blueness when one looks up on a fine summer day, and when one is having that visual experience characteristic of that condition, there’s nothing above one such that one sees that *it* is blue.” This footnote is getting out of hand. For a reply to the “blue sky” objection, I refer the reader to van Inwagen (2004: 136).

(¹¹) See Paul (2012).

(¹²) Van Inwagen (2011).

Concluding Meditation

(¹³) See also my reply to Michael Loux in the present volume.

(¹⁴) Lewis (1986).

(¹⁵) Perhaps she will reply to this statement as R. M. Chisholm once replied to a similar statement: “I accept the disjunction.”

(¹⁶) I’m going to suppose that the reader is familiar both with her essay in the present volume and *The Problem of Evil*.

(¹⁷) A neutral agnostic either assigns a subjective probability of 0.5 to the proposition that God exists (and, of course, to its denial) or declines to assign it any subjective probability (or to place it within any range of subjective probabilities—other than $[0,1]$), and, of course, does the same for its denial.

(¹⁸) That’s what I said in *The Problem of Evil*. Conversation with several philosophers has convinced me that it would have been better if I had said, ‘...*either* cease to be agnostics and become atheists, *or*, failing that, should at any rate cease to be *neutral* agnostics—should conclude that atheism is significantly more probable than theism’.

(¹⁹) We assume that the argument under consideration is formally valid.

(²⁰) See, for example, the two paragraphs that precede her statement of the No Tolerance Principle. But I confess to an imperfect understanding of the meaning of these two paragraphs.

(²¹) The concept “for all one knows” obviously has some sort of important connection with the concept of (subjective) probability, but drawing that connection would not be a trivial exercise. If you are about to be dealt a card from a well-shuffled standard deck, then for all you know that card will be the four of diamonds, despite the fact that there is only one chance in 52 of your being dealt that card.

(²²) A bowdlerized adaptation of a colorful demotic illustration of the distinction between “a possibility” and “a real possibility.”

Concluding Meditation

(²³) Feature (a) is very close to being an intrinsic feature of the stories that have it (one might contend that it is technically a relational feature of those stories, since they might lack it if, say, the evils of the world were significantly different). But feature (b) is blatantly relational; moreover, Theist's hope may fail: the story may turn out not even to *have* feature (b). It would therefore be more perspicuous to speak of a story's being on a certain occasion *presented as* a defense than of a story's *being* a defense.

(²⁴) One is a story that entails the existence of both God and all the evils that have befallen human beings; the other is a story that entails the existence of both God and all those evils that have befallen non-human animals—throughout the hundreds of millions of years during which there have been organisms capable of suffering—that cannot be ascribed to the actions of human beings. (The first defense was at a later point supplemented with further elements intended to address the “local” argument from evil. In this reply, I will not consider that aspect of the defense.)

(²⁵) “It’s not good enough for a Defense to say that it is ‘at least very plausible to suppose’ that there’s *something* that makes it all right for God to behave in a way that would be morally wrong if a human being did it. The whole point of the Defense is to *show* us what that difference might be.” Why? Who says that’s the whole point of the defense? In her paper, Antony capitalizes ‘defense’ to remind her readers that the word is my technical term. Well, yes, it’s *my* technical term, and I have to insist that nothing I said about its meaning implies that the whole point of any defense is to show how something could be.

(²⁶) And, of course, he wants them so to react to the conjunction of all the propositions the defense comprises—to the “whole defense.”

(²⁷) Again, we assume that the argument we are considering is formally valid.

(²⁸) Unlike this possibility: An extraterrestrial civilization has surrounded the solar system with an enormous force field that

Concluding Meditation

reflects radio waves from extra-solar artificial sources but has no effect on radio waves from extra-solar natural sources. That *might* be true. And flaming monkeys *might*...

(²⁹) As I implied in the preceding paragraph, I did not put the “second” defense into the mouth of an imaginary theist; I rather presented it in my own voice. This was a matter of literary convenience. Readers were invited to imagine it as being presented in the same “debate” format as the first defense.

(³⁰) My imaginary defense counsel’s “C.I.A. defense” and the defense suggested in note 28 are examples of defenses that are *prima facie* “simply unbelievable.” There are also stories that are not themselves *prima facie* “simply unbelievable” but which have demonstrable entailments with that feature. The story of the Spanish Barber is a trivial example. Antony at one point mentions a very non-trivial example indeed: Frege’s unrestricted comprehension principle—a “story” about the existence of the extensions of concepts.

(³¹) See van Inwagen (1988: 176-7).

(³²) And yet at one point at least Antony does seem to understand that this was the way I was arguing. See her footnote 4. My project in *The Problem of Evil* was to attempt to convince my readers that the argument from evil was a “failure” in the precise technical sense that I gave to that word.

(³³) At one point (p. 182) Antony quotes a passage from *The Problem of Evil* and says of it, “I quoted this passage in full because I find it stunning.” (That is to say, she finds it a stunningly inept dialectical lacuna in Theist’s presentation of his defense.) Now I don’t think she should have found it “stunning,” and I think that the reasons she proceeds to give for regarding it as stunning are specious—in fact, stunningly so. My point here, however, is that it’s a matter of indifference to me and to my mouthpiece Theist whether *she* finds this passage stunning. What Theist and I care about is whether the audience of agnostics will find it stunning—that is to say, we care about how the agnostics will regard the passage after

Concluding Meditation

listening to an extended debate between Theist and Atheist about whether it is a stunningly inept (etc.). I am confident that they will regard Theist as having “won” this debate—although, quite possibly, Atheist will be convinced that she has won it.

(³⁴) For a defense of this position and an account of ‘substantive’ and ‘positive’ (and a defense of the thesis that atheism is a “positive” thesis), see van Inwagen (2009a), especially note 3, p. 17.

(³⁵) “Are you saying, then, that the Consequence Argument, is *not* a ‘cogent’ argument for incompatibilism? That it is a ‘failure’? If so, why are you an incompatibilist?” Well...it’s complicated. See the essay referred to in the previous note and my replies to the papers in Part IV (“Method”) of the present volume.

(³⁶) Van Inwagen (2009b).

(³⁷) It is something like this: according to Darwinism, the history of life has been a radically contingent process; and no one, not even an omnipotent being, can control the outcome of a radically contingent process. I have examined this kind of argument in great detail in van Inwagen (2003). See especially pp. 356–61.

(³⁸) See van Inwagen (2009a).

(³⁹) Van Inwagen (2009a: 16–17).

(⁴⁰) This definition may need a little Chisholming. Some philosophical propositions, even some quite trivial ones, might be said not to be “accepted by almost all human beings” simply because they involved concepts that a significant proportion of human beings do not have—‘Eliminative materialism is incompatible with reductive materialism’, for example.

(⁴¹) Generalizations about “philosophers” are generalizations about human philosophers unless otherwise noted—and we shall be discussing some (fictional) non-human philosophers.

Concluding Meditation

(⁴²) This is a broader class of arguments than the class of arguments (that have been given by philosophers) whose conclusion is p or its denial. Suppose, for example, that A is an argument (given by philosophers) whose conclusion is p , and that A' is an argument (given...) whose conclusion is the denial of one of the premises of A (or whose conclusion is that A is invalid or that it assumes the point at issue...). Then A' is an argument that is “relevant to the question of the truth or falsity of p .”

(⁴³) This is loose talk; strictly speaking every world is of 0 (or perhaps infinitesimal) probability; what is “improbable” is not w itself but, rather, that a world having the features ascribed to w (there are infinitely many) should be the one that is actual.

(⁴⁴) I mean this to imply that in w he believes that consciousness exists; and, moreover, in w , his beliefs on this matter are consistent: he does not *also* believe that consciousness does not exist.

(⁴⁵) Let's not quibble about words: I stipulate that Zombies would be human beings in the sense of ‘human being’ that these words have in the definition of ‘controversial philosophical proposition’.

(⁴⁶) See McGinn (1993). See also van Inwagen (1996), a review of McGinn (1993).

(⁴⁷) The idea of such a species first turned up in some speculations of Noam Chomsky and was further developed by McGinn. For a “real” science-fiction story (of course unrelated to the writings of Chomsky and McGinn; it was first published in 1984) that features such a species, or at least a species that *claims* to know the answers to all philosophical questions, see George Alec Effinger, “The Aliens Who Knew, I Mean, *Everything*.” The story—originally published in *The Magazine of Fantasy and Science Fiction*—can easily be found on line by Googling its title.

(⁴⁸) Dennett and Steglich-Petersen (2008).

(⁴⁹) Van Inwagen (2015: Chapter 13).

Concluding Meditation

(⁵⁰) The following case is more closely parallel to the case we are considering: Jane must draw a card, and has been informed that she will be shot if the card she draws is neither a red card nor a face card. The following two statements are not equivalent: ‘The card Jane draws will be either a red card or a face card’: ‘If the card Jane draws is either a red card or a face card, we should not find that surprising’. (We should not find that outcome surprising because its probability is about 0.615.)

(⁵¹) An analogy. Suppose I’ve been set down at a random place on a roughly circular island about 500 meters in diameter. I know that a roughly circular region that occupies about a third of the island is an unmarked minefield. Naturally, I’d like to know whether I’m inside the minefield—but I don’t know, and have no way of finding out, whether I’m inside it. I do, however, somehow know that a boulder that is only ten meters from me *is* inside the minefield. I reason as follows: if I’m outside the minefield, I’m just barely outside it. And that’s rather improbable—it’s *rather* improbable, given the size of the island and the size of the minefield, that my being randomly placed on the island should have resulted in my being outside the minefield but within ten meters of its boundary. (If we assume that every point in the minefield is more than ten meters from the water’s edge, a simple geometrical calculation gives a probability of 0.86 that I am in the minefield, conditional on my being ten meters from an object in the minefield. Of course, in the “dark energy” case, no calculation of the numerical value of the probability of our being unable to discover its nature, conditional on the Spicans being unable to discover it, is possible.) An oracle then unexpectedly resolves the matter for me: I am indeed inside the minefield. I ought not to find this revelation particularly surprising, given what I already knew before the oracle spoke.

(⁵²) Van Inwagen (2006).

(⁵³) If there are first-person arguments, there are probably also second-person arguments. If so, Keller’s point applies to them *mutatis mutandis*.

Concluding Meditation

(⁵⁴) A potential predilection is not a predilection. A and B may both have no predilection to accept p and yet it may also be that it would be much easier to cause A to have such a predilection than to cause B to have one. It may be, for example, that A has a predilection to accept certain propositions that logically imply p and that B has no predilection to accept any propositions that logically imply p . (Note added in the present essay.)

(⁵⁵) ‘The highest degree of logical acumen’ was not meant to imply “logical omniscience” on their part. See the reply to Thomas Kelly and Sarah McGrath, p. 383.

(⁵⁶) Lilla (2001). Available on line at <<http://www.worldcat.org/wcpa/servlet/DCARead?standardNo=0940322765&standardNoType=1&excerpt=true>>.

(⁵⁷) It can therefore hardly be the case that “Pyrrho” and “Carneades” are the historical Pyrrho of Elis and the historical Carneades of Cyrene. In any case, their dialectical procedures do not correspond to those of their respective namesakes.

(⁵⁸) I.e., arguments each of which, when considered individually, seems to the members of AP and to Pyrrho himself *obviously* to be a proof or demonstration of its conclusion. Pyrrho, of course, concedes that some of or all his arguments are unsound: he does not reject the law of the excluded middle. But he is convinced that he hasn’t the faintest clue as to what the flaw(s) in any of his arguments might be; and he is attempting to induce (simply by presenting those arguments and defending them against all known objections) a like conviction in the minds of the members of AP .

(⁵⁹) I trust that everyone who has got this far in the discussion will know what ‘in IDEAL CIRCUMSTANCES’ abbreviates.

(⁶⁰) Van Inwagen (2006: 47). My account *is*, therefore, as I said earlier, audience-relative. Following the quoted sentence, I went on to say, “...so limiting the jury pool, of course, relativizes our criterion of success to our time and our culture.” I direct the interested reader to the discussion of this “relativization” on pp. 47 and 48 of *The Problem of Evil*.

Concluding Meditation

(⁶¹) Or at least a probability that is *essentially* 1—a probability so close to 1 that we unreflectively treat it as 1 when we are engaged in practical reasoning.

(⁶²) Assuming that everyone really can be said to “have” a set of priors. If that assumption is wrong, of course, then Bayesianism is wrong and thus does not present Success_{PVI} with any sort of difficulty.

(⁶³) If you think that knowing that the consequent of an indicative conditional is true is not sufficient for knowing that the conditional is true, substitute this disjunction for the conditional: Either (it’s not the case that Alvin Plantinga has said that Genuine Modal Realism is not an acceptable modal ontology) or Genuine Modal Realism is not an acceptable modal ontology.

(⁶⁴) If you doubt whether I know that Genuine Modal Realism is false, I could state my conclusion hypothetically: If I *did* know that Genuine Modal Realism was false, this argument would not be a success for me.

(⁶⁵) Might it be that the argument I have offered as a counterexample to Keller’s definition “begs the question,” and is thus fallacious? That is hard to say, because it is hard to say what it is for an argument to beg the question. (It would certainly be *possible* for someone to know that the conditional second premise of the argument was true without knowing whether its conclusion was true.) But if Keller’s definition is modified in the way I have suggested, the proponent of the revised definition who undertakes to determine whether (e.g.) the argument from evil is a success for Eleonore Stump need not consider whether the argument from evil is guilty of the fallacy of begging the question—whatever, precisely, that fallacy may be. Note also that if God has revealed to Meghan Sullivan that Alvin Plantinga is always right about modal ontology, the argument in the text will be a success for *her* even if, in some impersonal sense, it begs the question. (One who accepts my revision of Keller’s definition and who believes that there is such a fallacy as begging the question should probably replace Keller’s ‘and non-fallacious’ with

Concluding Meditation

something like ‘and is non-fallacious except, possibly, in the matter of begging the question’.)

(⁶⁶) At least some versions of the argument from evil are very good philosophical arguments indeed.

(⁶⁷) Think of the relation between these two terms this way: to say that an argument is a failure is to say that it fails to be a compelling argument.

(⁶⁸) If one uses the revised definition, one should not say that an argument that is not a failure is “compelling.”

(⁶⁹) I touched on this point in *The Problem of Evil*. See pp. 50–1.

(⁷⁰) We assume throughout, for the sake of simplicity, that all the arguments we are considering are uncontroversially logically valid. (If there are questions about what *that* means, I should be willing to discuss them—but to attempt within the scope of this brief reply to anticipate and address such questions would be what one of my teachers called philosophical shadow-boxing.) In consequence, the only questions relevant to our purposes that can be raised about the arguments pertain to the truth-values of their premises.

(⁷¹) Van Inwagen (2009b).

(⁷²) ‘Known’ is short for ‘well known to philosophers’. An argument that exists only in the mind or the unpublished ms. of some reclusive philosophical genius is of course “known” to *her*, but it is not known in the sense I intend.

(⁷³) Note that the following sound argument, ‘The Earth is round; *hence*, if the Earth is not round, the will is free’ is not an argument a priori. The valid argument, ‘The Earth is both round and not round; *hence*, If Descartes died in Sweden, persons are immaterial substances’ is an argument a priori, but pretty clearly a failure.

(⁷⁴) The modification would be not only easy but not needed at all if it were true that any conditional proposition whose antecedent was an empirical proposition and whose consequent was a significant philosophical thesis was itself a

Concluding Meditation

significant philosophical thesis. But that is obviously not the case: the true proposition ‘If the earth is not round, the will is free’ is not a significant philosophical thesis. Cf. n 73.

(⁷⁵) The closest such worlds do not include worlds in which all life on earth ended in 1979 or in which all universities eliminated their philosophy departments in the 1980s or worlds in which in that decade all philosophers happened to suffer head injuries that caused them to lose their memories or in which philosophers perversely conspire to ensure that certain arguments they know of “die with them”—conspire to keep them secret from their students....

(⁷⁶) I cannot, alas, claim that the “Consequence Argument” has been a philosophical success. But consider the following two “populations”: professional philosophers with an interest in the free-will problem who first encountered the Consequence Argument when they were undergraduates or graduate students; professional philosophers with an interest in the free-will problem who first encountered the Consequence Argument more than ten years after receiving their doctorates. At every time at which both populations were large enough for this judgment to be meaningful, the proportion of the former who were then incompatibilists has been significantly higher than the proportion of the latter population who were then incompatibilists.

(⁷⁷) See *The Problem of Evil*, pp. 37-9.

(⁷⁸) I must protest a statement in Kelly and McGrath’s essay, which I will quote. Speaking of some passages they have quoted from the version of the argument presented in *The Problem of Evil*, they say, “In these passages, the fact that the conclusion of a philosophical argument is typically denied by a significant number of philosophers is treated as compelling evidence that the argument would not convince the ideal agnostics” (p. 335). No, not *compelling* evidence. Good evidence, however—evidence that should lead one to take the Pessimistic Verdict seriously.