

Protecting rainforest realism

James Ladyman, Don Ross: Everything must go: metaphysics naturalized, Oxford: Oxford University Press, 2007, pp. 368 £49.00 HB

**P. Kyle Stanford · Paul Humphreys ·
Katherine Hawley · James Ladyman · Don Ross**

Published online: 30 March 2010

© The Author(s) 2010. This article is published with open access at Springerlink.com

P. Kyle Stanford

James Ladyman and Donald Ross's broadside against traditional analytic metaphysics embodies the most admirable characteristics of a good slap across the face: it is forceful, frank, and delivered in response to sufficient provocation. Ladyman and Ross are quite right to point out that much of analytic metaphysics, when it seeks to take the findings of the empirical sciences into account at all, appeals instead to what they call a "domestication" of those findings into a more intuitively comfortable picture of the physical world that appeals to "little things", "microbangings", and a "containment metaphor." And they are right to suggest that

P. Kyle Stanford (✉)

Department of Logic and Philosophy of Science, University of California, Irvine, CA 92697-5100,
USA

e-mail: stanford@uci.edu

P. Humphreys

Department of Philosophy, University of Virginia, Coker Hall 105, Charlottesville, Virginia, USA

e-mail: pwh2a@virginia.edu

K. Hawley

School of Philosophical, Anthropological and Film Studies, University of St. Andrews, St. Andrews,
Fife KY16 9AR, Scotland, UK

e-mail: kjh5@st-andrews.ac.uk

J. Ladyman

Department of Philosophy, University of Bristol, Bristol, UK

e-mail: james.ladyman@bristol.ac.uk

D. Ross

School of Economics, University of Cape Town, Cape Town, South Africa

e-mail: don.ross@uct.ac.za

there is little point in trying to guide our metaphysics by or reconcile it to this unholy chimera of antiquated science, popular science, folk science, science fiction, and commonsense or philosophical intuition, an enterprise whose products they deride as “the philosophy of A-level chemistry.” Worst of all, Ladyman and Ross are surely also right to point out that much of what happens in contemporary analytic metaphysics takes place with a complete lack of concern for the findings of the sciences at all, proceeding instead to attempt derive substantive results about ways the world is, might be, or must be simply by reflection on our intuitions and concepts.

Of course it is not an *a priori*, analytic, or necessary truth that such conceptual analysis is a misguided or hopeless strategy for investigating the most fundamental features of existence, but its use would nonetheless seem to require a convincing philosophical rationale. Descartes thought he had one, insofar as he took himself to have decisive reasons for believing his concepts to have been created by an omnipotent, omniscient, and benevolent God who would (therefore?) not deceive him so long as he was careful to reason judiciously from and about those concepts. A more recently influential rationale was offered by Kant, whose Copernican revolution is widely taken to insist that much of the structure of the external world as we experience it is imposed upon it by the central concepts we use to construct that experience of the world from the raw materials of sensation. On these sweeping views of the universe and our place within it, such conceptual analysis is a weighty business indeed, but for good and familiar reasons few contemporary analytic metaphysicians defend the interest and significance of what they are doing by declaring allegiance to these grand philosophical systems. It is sometimes suggested more prosaically that conceptual analysis tells us about the world because it is, after all, only by means of our concepts that we engage the world—it is then, in some sense, our concepts and the meanings we assign to the associated terms that determine what *counts* as a “material object”, an “injustice”, or a “peanut butter sandwich”. But bereft of something like the Kantian rationale, this is a very cheap respect in which our ideas determine anything about what the world is like. Instead, it seems more natural to say that such conceptual analysis at best tells us something about how we think and talk about the world, and not about the characteristics of the entities and events that make it up. And the significance of such inquiry is further diminished for those who think that these very concepts and meanings are subject to interpersonal variation and/or change over time at even the most fundamental levels. The picture of much contemporary analytic metaphysics that emerges is that of an admittedly sophisticated practice that has nonetheless almost literally lost its bearings—it continues apace despite the progressive deterioration into a crumbling ruin of the rationale that once grounded its interest and significance.

On the other hand, the conception of metaphysics that Ladyman and Ross seek to erect in place of the Cartesian idols they throw down also demands a convincing rationale, and here I think they are less surefooted. On their account of the matter the point of metaphysical inquiry is simply to unify what we have already (and independently) learned about the world from science itself, for “the *raison d’être* of a useful metaphysics is to show how the separately developed and justified pieces of

science (at a given time) can be fitted together to compose a unified world-view” (45). An early statement of their “Principle of Naturalistic Closure” tells us that “Any new metaphysical claim that is to be taken seriously should be motivated by, and only by, the service it would perform, if true, in showing how two or more specific scientific hypotheses jointly explain more than the sum of what is explained by the two hypotheses taken separately, where a ‘scientific hypothesis’ is understood as an hypothesis that is taken seriously by institutionally *bona fide* current science” (30).¹ Accordingly, the only standard of success for a metaphysical claim or proposal is its ability to unify what the various sciences have *already* told us: “one metaphysical proposal...is to be preferred to another to the extent that the first unifies more of current science in an more enlightening way” (66). Thus, although Ladyman and Ross are quite clear that we cannot simply read our metaphysical commitments off our best scientific theories but must instead first examine how those theories are “practically put to work” (119), it seems that they nonetheless intend for us to take quite literally their suggestion that metaphysical inquiry properly conceived is fundamentally derivative from and subordinate to that of scientific investigation. Indeed, they tell us that by the very term “metaphysics” they will “refer to the articulation of a unified world-view *derived* from the details of scientific research” (65, my emphasis).

I am not entirely sure how to pursue metaphysics so conceived despite the example they go onto set, but I suspect there is a better approach in any case. Quine long ago suggested that there is only a single project of inquiry, and thus at least implicitly that any metaphysics worth its salt was neither prior nor posterior to the scientific investigation of the world, but indistinguishable from it. On this view we discover features of the world only by entheorizing it most successfully using *whatever* evidence is deemed relevant by the lights of the general picture of the world and our place in it that we inherit to begin this process and the one that results from developing it further in response to those evidential inputs. And on such a view, the only relevant categories into which reasons for holding a belief fall are *good* and *bad*, not scientific versus unscientific or philosophical or some other kind. Of course, this account no more suggests that we can simply read fundamental ontology off our best scientific theories than do Ladyman and Ross, for evidence from other sources—evidence from the history of scientific inquiry, for example, concerning characteristics of our theories and ourselves as theorizers no less than the results of controlled scientific experiments—must feed back to inform our view of what we are doing when we theorize about the world and of the character, reach, and epistemic status we should ascribe to the various products of that theorizing in different circumstances: a view of all this, too, must be part of any integrated account we accept of the world and our place in it. What is central here is neither radical holism nor any number of other Quinean idiosyncracies, but simply the insistence that *all* the (good) evidence matters, and the consequent refusal to begin by elevating science, history, philosophy, or any other restricted part of the overall

¹ Ladyman and Ross seek to finesse the acute demarcation problem thereby created with the unpromising suggestion that we can let such institutional factors as the likelihood of obtaining funding stand proxy for any more direct (and inevitably less reliable) philosophical or epistemic evaluation of what counts as genuine or *bona fide* science (see esp. 30–38).

project of inquiry into a privileged position. (The spirit of this suggestion is well-represented by Penelope Maddy's (2007), though Maddy and I differ sharply concerning what the historical evidence reveals about the reach and limits of our scientific theorizing.) This approach is straightforwardly precluded by Ladyman and Ross's two-stage model of inquiry, in which we must first ascertain what bona-fide science has independently concluded about nature before undertaking the distinctively metaphysical unifying project that tells us what these "separately developed and justified" findings really imply concerning the most fundamental constitution of the world.

The deep contrast between Ladyman and Ross's "frank scientism" (61) and this broadly Quinean alternative is somewhat obscured by their claim to find confirmation for their Ontic Structural Realism (OSR) in the consilience of a wide variety of sources of support: by their lights, OSR is not merely the metaphysical view recommended to us by the cutting edge of contemporary physical theorizing, but also one that solves a wide variety of persistent problems in the philosophy of science itself, offering convincing accounts of laws, causes, classification, and explanation, while simultaneously resolving the dispute between scientific realists and their opponents in a way that is responsive to the central considerations advanced by both sides of the debate. Though I harbor concerns about its selectivity, I will not second-guess Ladyman and Ross's complex technical discussion of whether OSR represents the view of ontology that fundamental physics counsels us to adopt. I do want to suggest, however, that there is no unequivocal notion of "structure" that can do all the jobs they want and need it to do. More specifically, I doubt there is a single *kind* of structure that is simultaneously recommended to us by fundamental physical theorizing, preserved in the transitions between all or most suitably successful past scientific theories, and sufficient to answer the realist demand to genuinely explain the success of our scientific theories.

Consider this final demand first. Ladyman and Ross suggest that there are aspects of the success of our scientific theories that really do demand or at least favor some form of scientific realism as their explanation: most importantly, the ability of our best scientific theories to (sometimes) predict novel phenomena whose existence we would have little reason to suspect absent the theory. They argue, however, that the historical record of radical ontological discontinuity between increasingly successful generations of scientific theories simply precludes us from embracing the traditional scientific realist's explanation that contemporary theories succeed in this way because they are approximately true. Happily, they suggest, such radical discontinuity does not afflict well-confirmed theories' claims about modal structure; instead "all the well-confirmed modal relations expressed by old theories are approximately recovered in their successors" (123). "By modal structure," they tell us, "we mean the relationships among phenomena...that pertain to necessity, possibility, potentiality, and probability" (153–154), and this supposedly gives us all we really need to explain even the novel predictive success of our best theories: "Since some theories have achieved novel predictive success our overall metaphysics must explain how novel predictive success can occur, and the explanation we favor is that the world has a modal structure which our best scientific theories describe" (79).

In the very next paragraph, however, Ladyman and Ross identify this same modal structure with van Fraassen's "relations among the phenomena", but conceived modally or nomologically, rather than as "extensional, occurrent regularities". This seems to me to involve a subtle sleight-of-hand. The idea seems to be that, contra van Fraassen, our successful theories don't describe brute regularities in the relations among the phenomena, instead they describe modal structure, and it is because they describe such modal structure correctly that when they suggest the existence of a new phenomenon and we look for it, we (sometimes) find it. But it seems for all the world that to conceive of van Fraassen's "extensional, occurrent regularities" among the phenomena *as modal* would *just mean* to regard them as persisting throughout circumstances that are presently counterfactual, merely possible, or untested (see esp. Sect. 2.3.2.3), even where this would require the existence of novel phenomena. If so, then to say that our theories correctly describe the modal structure of the world (in this sense) is simply *to redescribe or restate and not to explain* the fact that they enjoy novel predictive success (to whatever extent they do). Perhaps Ladyman and Ross are even within their rights to insist that such modal relationships between the phenomena are "ontologically basic" (128), but if there is a legitimate demand in the first place to *explain why* our best theories enable us to predict novel phenomena,² simply appealing to a modalized version of van Fraassen's "relations among the phenomena" does nothing to satisfy it—indeed this seems to have more the character of an incantation than an explanation.

Even if we suppose that modalized structural relations between the phenomena really do explain rather than simply redescribe novel predictive success, however, Ladyman and Ross must face an even more serious problem. According to their "Rainforest Realism", all that genuinely exists is "real patterns", and this is ultimately all there is to say about ontology: it is "real patterns all the way down" (228), we typically "track" or "locate" these real patterns by our agent-relative, merely pragmatic division of the world into individuals, causes, events, and processes which do not truly exist (see 153–154 and Chap. 4 *passim*).³ I confess it is not clear to me whether Ladyman and Ross think we can explain the novel predictive success of our best theories simply by appeal to the modalized structural relations uncovered by fundamental physics or whether those uncovered by the various special sciences are required as well: compare, for instance, their claim that "*From the point of view of those engaged in special science activity, fundamental physics gives the modal structure of the world*" (288, original emphasis) with "If science tells us about objective modal relations between the phenomena (both possible and actual), then occasional novel predictive success is not miraculous but

² Of course, Ladyman and Ross concede that novel predictive success is only "occasional" even for theories that correctly describe objective modal relations among the phenomena (153). And we might doubt that we know *what* we are trying to explain unless we know what rate of successful novel prediction we should expect not from fundamentally false theories *simpliciter*, but from just those fundamentally false theories *that have nonetheless managed to best explain, predict, or accommodate all the systematically related known phenomena at a given time in a given field of inquiry*.

³ As in Leibniz's monadology, the everyday world of objects and causes with which we are familiar reemerges from the metaphysical menagerie imposed on us by strict adherence to Christian theodicy or fundamental physics, but can only capture how things seem to us rather than how they really are. Matthew 6:24 tells us truly that no man can serve two masters.

to be expected” (153) and “The basis for our confidence that the special sciences often successfully track real patterns is the no-miracles argument...” (298). But neither option is ultimately attractive in any case. Let us see why.

The problem with taking the modalized structural relations described by fundamental physics to explain novel predictive success in other sciences is that the “scale-relativity” of ontology ensures that the real patterns identified or “tracked” by sciences operating at scales very different from those of fundamental physics are not connected by type identities, smooth reductions, or even supervenience to those recognized by fundamental physics itself (see esp. Sect. 4.4). Thus, even if we stipulate that the modalized structural relations we discover between phenomena in fundamental physics are genuinely explanatory, this will do little or nothing to explain how *other* sciences, by tracking or locating different scale-relative real patterns, which simply cross-classify those of fundamental physics in a wide variety of heterogeneous ways, manage to achieve novel predictive success when they do.

It seems, then, that to explain novel predictive success in sciences other than fundamental physics we must appeal to modalized structural relations between the phenomena uncovered by those very sciences themselves. To serve the needs of Ladyman and Ross’s OSR, however, these structural relations will also have to be of a sort that persist through the kinds of dramatic theoretical discontinuities that led Ladyman and Ross to demur from the traditional realist’s explanation of those novel predictive successes in the first place, and such a broad claim of historical continuity for well-confirmed modalized structural relationships between the phenomena seems empirically implausible. To take just one example, well in advance of its experimental confirmation J. F. Meckel made the novel prediction that gill slits must appear at some point in human ontogeny on the basis of his fundamentally recapitulationist conception of the structural relationship between the phenomena of ontogeny and phylogeny. But as Gould (1977) emphasizes, such recapitulationism (famously embodied in Haeckel’s Biogenetic Law that ‘ontogeny recapitulates phylogeny’) envisions ontogeny as the successive terminal addition of further developmental stages to the *adult* forms of phylogenetically earlier organisms and the “condensation” of those adult forms farther and farther back into the course of development. And this recapitulationist conception of the structural relationship between the stages of embryonic development and the history of life on earth is simply rejected by contemporary biology: contra Ladyman and Ross, the modalized structural relationship it describes cannot be “approximately recovered” (123) or “recovered as a limiting case” (157) from that which contemporary biologists take to hold between the phenomena of ontogeny and phylogeny. Instead, to accept anything like the modern view of this matter is simply to hold the distinctively *structural* relationship between the phenomena of ontogeny and phylogeny asserted by its recapitulationist predecessor to be fundamentally mistaken, as the scientists evaluating them recognized clearly: Morgan (1916, 18) described the widespread acceptance of the competing idea of “germinal variation” simply as “the death of the older conception of evolution by superposition”, noting (1916, 23) that

To my mind there is a wide difference between the old statement that the higher animals living today have the original adult stages telescoped into their

embryos, and the statement that the resemblance between certain characters in the embryos of higher animals and corresponding stages in the embryos of lower animals is most plausibly explained by the assumption that they have descended from the same ancestors, and that their common structures are embryonic survivals.

Thus the historical record promises a rude reception to Ladyman and Ross's claim that "we know that well-confirmed relations among phenomena must be retained by future theories" (157).

We have by now, I think, begun to see *why* there is no univocal sense of "structure" that can do all the jobs Ladyman and Ross need it to do, for these various jobs pull in competing directions. Only a very thin conception of "structure", for example, seems to have any chance of surviving the challenge posed by historical discontinuity, while a much thicker one is needed to explain novel predictive success. But of course, this is simply an especially robust version of a challenge that remains unsolved by structural realisms more generally. It seems perfectly reasonable to suggest that in every transition from one scientific theory to another we can always find *something* plausibly described as "structural" that remains continuous between the two theories, but further historical investigation reveals it to be quite implausible to suggest that *one and the same* structural element remains continuous in all such successions, or (therefore) that we can project into the future the "structural" elements of current theories that will survive. In a similar fashion, it seems plausible enough to suggest that something fairly described as "structural" would suffice to explain the novel predictive success of our theories, and perhaps even that something fairly described as "structural" is what we find described by our most fundamental physical theories. But it seems very implausible, in ways suggested by the challenges above, to suggest that the very *same* structural element that remains continuous through theoretical transitions in the historical record (if there were one!) *also* suffices to explain the novel predictive successes of those theories and is *also* the one described as fundamental by our best physical theories. Seen in this light, even the modest claim that structural realism is supported by the historical record seems to trade on the inherent vagueness and ambiguity in what we are prepared to describe as "structural", while Ladyman and Ross's much more ambitious claim that their Ontic Structural Realism is supported by a consilience of many different sources of evidence (*including* its ability to answer this historical challenge) threatens to simply turn the case for structural realism into a bad pun.⁴

Paul Humphreys

Over the last few years a growing divide has emerged between the fields that are often called 'analytic metaphysics' and 'scientific metaphysics'. A priori methods and conceptual analysis are the primary tools of choice for analytic metaphysics and

⁴ Acknowledgements: Many thanks to Pen Maddy and to Tom Pashby for useful discussion of many of these issues.

empirically grounded scientific results are not a part of its evidence base. (With the exception that actual existence entails possible existence.) In the past, a core part of metaphysics consisted an effort to establish modal truths. Necessary truths, of course, hold in our world but at some point parts of analytic metaphysics began to make claims that were specifically about the metaphysics of our universe. A well-known example of this kind of project is David Lewis's Humean Supervenience program: 'I have conceded that Humean Supervenience is a contingent, therefore an empirical, issue. Then why should I, as a philosopher rather than a physics fan, care about it?... Really, what I uphold is not so much the truth of Humean Supervenience as the *tenability* of it. If physics itself were to teach me that it is false, I would not grieve' (Lewis 1986, xi).⁵ As an exercise in theorizing, this attitude would be unobjectionable were the results seriously put to an empirical test. The problems here are twofold. First, Humean Supervenience is factually false. It has long been empirically established that entangled states in quantum mechanics exist and do not supervene on what would in classical cases be called the states of the components. The second problem is that science, and physics in particular, long ago outran the conceptual abilities of most metaphysicians. Consider this passage: 'Many philosophers use 'ontology' as a name for the study of the most general structures displayed by objects... the questions about the structure of a chair that I can understand are questions to be answered by carpenters, chemists, and physicists... What I cannot see is how a chair could have any sort of structure but a spatial or a mereological structure...' (van Inwagen 2001, 1–2).

As the authors of *Everything Must Go* (hereafter EMG is used, as appropriate, for either the book or its authors) repeatedly and correctly note, contemporary scientific results, including the chemical and physical answers mentioned but not used by this author, have revealed a much more subtle and interesting world than mereology gives us.⁶ This is the domain of scientific metaphysics, an activity that is primarily interested in the ontology of our world, and its practitioners hold that contemporary science, and physics in particular, serves as a primary source of information about that ontology. Because the domains of scientific metaphysics and analytic metaphysics overlap it is inevitable that conflict has arisen between them. Results from scientific metaphysics have shown that what were thought to be necessary truths are not even contingently true. Various versions of Bell's Theorem have given us insights into, and evidence for, the existence of indeterminism thus undermining various principles of sufficient reason; evidence from quantum statistics shows that the identity of indiscernibles principle is false; knowledge about the nature of human minds that has been gained from cognitive psychology and

⁵ Lewis's position is, of course, more nuanced than this short quote can indicate. In Lewis (1994, 474), he wrote 'The point of defending Humean Supervenience is not to support reactionary physics, but rather to resist philosophical arguments that there are more things in heaven and earth than physics has dreamt of.' An illuminating discussion of Lewis's attempts to balance ordinary intuitions and scientific claims is Weatherson (unpublished, 'Why care about humean supervenience?' Available at <http://lewisblog.weatherson.org/archives/004549.html>).

⁶ There are scientifically sensitive advocates of more sophisticated forms of mereology. For one example, see Arntzenius and Hawthorne (2005) although, as I understand it, the structural realism of EMG is incompatible with those versions of mereology as well.

neuroscience has, at least in some areas, improved philosophical discussions of our interior lives. On the other side, advocates of analytic metaphysics point to the pessimistic induction and argue that appeals to scientific authority are too flimsy a basis for metaphysical truth.⁷

The first chapter of EMG contains a slashing attack on analytic metaphysics. The attack is wide-ranging and it is primarily aimed at ridding metaphysics of individuals and mereological claims. Although the criticisms of analytic metaphysics expressed in EMG will attract considerable attention, the majority of the book is devoted to an exercise in naturalistic metaphysics, the elaboration and defence of an information-theoretic version of ontological structural realism. This material is dense, challenging, and creative. It repays serious study and, my disagreements with their position notwithstanding, this book constitutes a major contribution to the structural realist movement. Moreover, this much about EMG's first chapter is correct: some of the more prominent literature in contemporary analytic metaphysics is addressing manufactured problems, it is dismissive of well-established empirical research, and it is not even factually true.

Individuals

Let me begin with a constructive suggestion. Principal amongst the things that must go, according to EMG, are individuals: '[In the structuralist view] individuals are nothing over and above the nexus of relations in which they stand.' (138) Their core argument against the existence of individuals is the ample empirical evidence that at or near the fundamental level, the correct statistics are Fermi–Dirac or Bose–Einstein and this fact requires us to abandon the position that these systems are composed of individuals. (I leave aside parastatistics and other alternatives here.) There is no question that this conclusion is correct when we are dealing with phenomena that lie within the scope of quantum mechanics. What happens when we move away from that domain? When $e^{-\mu/kT} \gg 1$ (where μ is the chemical potential), the particle density distribution for a Fermi gas approaches the Maxwell–Boltzmann distribution and since the quantum concentration is inversely proportional to temperature, at moderate temperatures and moderate to low densities, the Maxwell–Boltzmann distribution applies to most macroscopic gases. This condition $e^{-\mu/kT} \gg 1$ is equivalent to the situation in which the average distance between the components of the gas is much greater than the de Broglie wavelength, a condition that lends itself to the interpretation that the spatial position of the components can be used to individuate the particles. (Bloch 2000, 236–237, 244–245) The distinguishability of the entities is thus retained and this seems to allow our ontology to revert to one using individuals in certain domains. However, the distinguishability position falls foul of the Gibbs paradox. This paradox hinges on the empirically established fact that if we mix two gases with the same number and

⁷ EMG are not alone in their critical attitude towards analytic metaphysics. van Fraassen's (2002) is severely critical of traditional metaphysics, although he is not an advocate of ontological structural realism (e.g. van Fraassen 2005). One thing that EMG does not mention is the remarkably convenient fit between the conclusions of some contemporary analytic metaphysicians and specific religious beliefs. Of course, on the other side, physicalism is an attractive resting place for atheists.

kind of particles the entropy does not increase under the permutation of particles, as it should if the particles were distinguishable. That is to say, there is no mixing entropy even with classical particles. (See Saunders 2006 for a clear presentation of some other arguments that indistinguishability must be extended to classical ‘particles’ as well.) These arguments could serve as a useful strengthening of EMG’s position because these are not philosophical arguments for or against individuals, but empirical arguments; even in domains outside those where quantum effects dominate, traditional, individual-based Maxwell–Boltzmann statistics fail to give the correct empirical results.

Fundamental physics and physicalism

My first disagreement with EMG’s position concerns its two central principles. Fundamental physics is given a privileged position by the authors via the Primacy of Physics Constraint (PPC):

Special science hypotheses that conflict with fundamental physics, or such consensus as there is in fundamental physics, should be rejected for that reason alone. Fundamental physical hypotheses are not symmetrically hostage to the conclusions of the special sciences. (44)

The PPC is accompanied by a unification principle, the Principle of Naturalistic Closure (PNC):

Any new metaphysical claim that is to be taken seriously at time t should be motivated by, and only by, the service it would perform, if true, in showing how two or more specific scientific hypotheses, at least one of which is drawn from fundamental physics, jointly explain more than the sum of what is explained by the two hypotheses taken separately...

Despite this appeal to fundamental physics, EMG are not physicalists. They reject a hierarchical, layered, view of ontology and hold there is no reason to think that there is a bottom level of physical ontology. By fundamental physics they mean ‘...that part of physics about which measurements taken anywhere in the universe carry information.’ (55) It is essential when assessing these principles to not make an unintentional slide from fundamental physics to the entire domain of physics. EMG invites confusion by often using the term ‘physics’ as shorthand for ‘fundamental physics’. Although this is noted (f.n.2, p. 190), the effect of this contraction is to make the constraints seem different than they actually are. (And in fact one such crucial switch is made on pp. 43–44 right before stating the PPC.) The Primacy of Physics Constraint should, therefore, properly be called the Primacy of Fundamental Physics Constraint—hereafter PFPC. This clarification is important because EMG relies on two kinds of inductive evidence to support the PFPC. The first appeals to a lack of downward influence from special science phenomena to physical phenomena and the second appeals to progress in understanding many areas of the special sciences in terms of physics. In the overwhelming majority of cases, this understanding is gained through work in biology, chemistry, and non-fundamental physics, not through fundamental physics. If fundamental physics in

their sense is to be connected with these successes, EMG must find some other basis for the transitivity of explanation than the traditional reductive hierarchy that exists within physicalism. One difficulty for such a project is that well before we get to fundamental physics, physics itself has decoupled certain length scales from others through the use of effective field theories. As Batterman (2002) has noted, the explanations of physical phenomena that satisfy universality criteria are independent of the details of finer grained theories. This means that a considerable part of fundamental physics is simply irrelevant to other areas of physics, which can be treated on their own terms. Thus, in many areas of physics itself, the inductive evidence in favor of the PPC does not count in favor of the PFPC.

This focus on fundamental physics is partly due to EMG's dismissive attitude towards ontologically emergent phenomena, preferring a variant of Dennett's anemic 'real patterns' to account for non-fundamental phenomena. If we shift attention from the focus on individuals, there is ample empirical evidence for entangled states, phase transitions involving long range correlations, topological entropy, the quantum Hall effect, and other features of non-fundamental physics that use ontologies which are not present in fundamental physics as construed by EMG and that can be treated as autonomous subjects.

Unification

The PNC is thus also inappropriate, because if an ontology is introduced for a non-fundamental theory, and that non-fundamental theory is decoupled from fundamental physics in the way just noted, it is unreasonable to require that it work together with fundamental physics to provide an increased degree of unification. The problems run deep because EMG's position is internally so coherent. Structural realism initially emphasized the importance of mathematical structures in order to preserve continuity across large-scale theory replacement.⁸ Many of these mathematical structures can then serve as the basis of an argument pattern in Kitcher's (1989) sense, these argument patterns being the basis of EMG's account of unification. This position is consistent with some widely canvassed information-theoretic formulations of physics (e.g. Haken 2006), with digital physics (e.g. Fredkin 1990), and a specific information-theoretic version of structural realism has also been defended by Floridi (2008).

Yet there is tension in the position. As well as denying the existence of individuals as part of a legitimate scientific ontology, a structuralist position has to deny the scientific relevance of intrinsic properties, something that in computational theories of the mind lead to one of the most serious splits in the field, the split between those who deny the importance of intrinsic properties of individuals (among them qualia and the first person perspective) and those who see them as ineliminable. The issues for chemistry, biology, physics, and neuroscience are different from those for the human mind, because there is no first person perspective

⁸ Current ontological structural realism is entirely dissimilar to Worrall's original version of structural realism.

to account for in those areas. And here, I think, we have the clue to what is truly distinctive about EMG's version of ontological structural realism.

Enter the human sciences, exit humans

Regarding the PNC, as a realist I take the world to be the way that it is and if it turns out not to be unified in certain respects, then that is what realists must accept. Whether the domains of the various sciences are unified or not should reflect matters of fact and this is not a situation to be legislated by methodological principles or 'stances'. For example, individuals are dispensable in some areas of the human sciences. The German sociologist Luhmann (1997) has argued that an individual-free sociology better captures the facts about contemporary capitalist society than do individual-based approaches. Within consumer economies, it is irrelevant who purchases the pack of cigarettes—the purchaser can be male, female, Chilean or Chinese, middle-aged or old, white collar or blue collar—all that matters is that the relevant economic communications take place. Indeed, Luhmann's work is a striking example of a research program within which the importance of humans as individuals is severely diminished and the emphasis placed on the autonomy of higher-level features. But it does not follow from this that individuals are eliminable as relata in other social relations. The man whose life has been saved by another can accord that individual a unique place in his world; honour, respect, and love are not relations that are indifferent to the identities of their relata.

What is the status of the relations in EMG's version of structural realism? EMG remains deliberately non-committal: 'What makes the structure physical and not mathematical? That is a question we refuse to answer. In our view, there is nothing more to be said about this that does not amount to empty words and venture beyond what the PNC allows. This 'world-structure' just is and exists independently of us and we represent it mathematico-physically via our theories.' (158) This is an intellectual evasion. We are supposed to be realists about structures but we are not told what they are. The evasion is also at odds with the scientific approach that the authors advocate and I look forward to the authors significantly clarifying the nature of these structures in their response to this review essay.

The fact that existing philosophical arguments for massive disunity in the sciences are unconvincing does not require us to shift to the opposite extreme of taking unification as a basic principle in a philosophical position via the PNC. Successful unification, cashed out by EMG in terms of explanations following Kitcher's argument patterns, has to be based on the superior empirical success of the unifying theories. If not, the program runs the risk of being outmanoeuvred by the superior unifying power of simplistic analytic metaphysics. And the current empirical evidence, especially from condensed matter physics, is that you have to take on their own terms the kind of non-fundamental physical phenomena mentioned earlier. Nor does it follow from the fact that a domain has informational universality that it is epistemically, or scientifically, privileged.

Returning to the PFPC, humans are epistemically located with respect to the rest of the world in a way that makes some aspects of the world transparent to our scientific investigations, whereas other aspects are not. The difficulty of access can

be due to the kind of mathematics required to capture the structure of the system, the degree of difficulty of converting a system's properties into those accessible to human perception, the degree of amplification needed, and many other factors. There is no reason to believe that these degrees of accessibility correspond in any straightforward way to whether information resulting from measurements is available only in selected parts of the universe, or everywhere. Fundamental physics is both epistemically highly opaque to humans and highly speculative, whereas some less fundamental sciences, perhaps because of their localization, are epistemically transparent and stable. It is no slight to physicists working on quantum gravity to point out that knowledge in that area is less secure and successful than are the knowledge that doxycycline is currently a safe and effective treatment for Rocky Mountain spotted fever and the knowledge that the American dog tick is a principal vector. I make this epistemic comparison because once the usual mereological fundamentality of physics has been rejected—and rightly so—an opposition to metaphysics and a concomitant deference to science must be grounded in the epistemic superiority of various parts of science. The epistemic ordering does not follow the traditional empiricist's epistemic ordering, but neither does it entitle fundamental physics, in EMG's sense, to a privileged place in a realist's ontology.

So what the appropriate type of realism for us to adopt? What I conclude from this analysis is that the evidence for ontological structural realism would be persuasive in the context of the kind of automated science that I discuss in Humphreys (2004) and (2009). It does not have the same kind of evidential power in the context of contemporary, anthropocentric, science.

Science as institutional

One final issue concerns identifying what counts as scientific. EMG adopts a firm stand on what is good and bad science: 'There is no such thing as scientific method...Thus, science is, according to us, demarcated from non-science solely by institutional norms...' (28) and '...we assume that the institutions of modern science are more reliable epistemic filters than are any criteria that could be identified by philosophical analysis and written down.' (37) The problem is that these filters do not work very well for some sciences and for reasons that can be stated by philosophers very concisely. For many years there were no institutional pressures on theoretical economists to subject their models to serious empirical tests. As a result, a domain of highly stylized techniques in neo-classical economics dominated the field for decades, supported by only minimal empirical evidence and resisting attempts to develop empirical data to test their adequacy. Complaints by what are often called heterodox economists that they were excluded from mainstream academic life are not simply the result of academic paranoia. There is now overwhelming empirical evidence that large parts of micro-economics are either inapplicable to real economics systems or are simply false. Henrich et al. (2005a) demonstrate that in ultimatum, public goods, and dictator games, models based on self interest failed to make correct predictions of subjects' behavior in *every single society they studied*, covering a broad spectrum of economic and cultural situations. What we can conclude from this is that institutions by

themselves do not always properly regulate science. They must have appropriate normative principles that are brought to bear on the activities within the field. Any automated science will inevitably have to operate on explicitly stated rules about what is scientifically acceptable and what is not and it is a joint task for scientists and philosophers, not for sociologists of science, to identify what these are.

Conclusion

EMG is a provocative book that overstates its case. Despite this, the first chapter should be required reading for every analytic metaphysician. Philosophers of science should read the book for a different reason: it will, if read with an open mind, challenge them in an intellectually serious way to defend their current methodological and metaphysical positions. It will not be an easy task. There are stunning developments in various sciences, the world is being revealed as a very strange place indeed, and having been distracted for decades by historical and social studies, contemporary philosophy of science has still not adapted to the conceptual challenges. My disagreements aside, the authors are to be commended for taking on the challenge to develop a systematic, scientifically informed metaphysics for the twenty-first century.⁹

Katherine Hawley

Every Thing Must Go is wildly ambitious. It advances substantive views on the proper scope of metaphysics (unifying science), the nature of reality (things subservient to structures), the current state of play in quantum gravity (fragmented), and the connection between fundamental physics and the rest of science (hard to summarise). It is both fascinating and infuriating. A key theme is the dismissal of ‘neo-scholastic’ metaphysics and the promotion of ‘naturalised metaphysics’. I fear my own work qualifies as neo-scholastic, and although I’m reassured to have ‘some extremely intelligent and morally serious people’ as company, I’d hate to think we were ‘wasting [our] talents—and, worse, sowing systematic confusion about the world, and about how to find out about it’ (vii). So I will focus my attention on this theme.

Broadly speaking, *EMG*’s objection to contemporary analytic metaphysics is that it is typically conducted in ignorance of contemporary science. Now, whilst metaphysicians vary significantly in their scientific knowledge, it is certainly true that most of us don’t know much more than we can read in *Scientific American*, and many of us don’t even know that much. Those who *do* know plenty of science tend to be classified as philosophers of physics (or biology), not metaphysicians. But why does all this matter?

⁹ Acknowledgement: Thanks to the participants in my Spring 2008 philosophy of science seminar for helpful discussions on these issues. Comments from Jim Cargile and Jorge Secada were particularly helpful. Thanks also to Margaret Morrison for critical reactions that led to sharpening some of the points made here and to Anouk Barberousse and Robert Batterman for comments that improved a previous draft.

The authors of *EMG* think that it matters for a number of reasons. They argue that scientifically-ignorant metaphysicians often (mistakenly) think of the world as fundamentally composed of tiny billiard-ball-like particles, and that this gives rise to errors about causation, composition, extension, objecthood and so on. Relatedly, they argue that metaphysicians often rely on thought experiments concerning simple worlds, and underestimate how unlike actuality these possibilities are. They argue that metaphysicians pronounce *a priori* on matters concerning which scientists have empirical evidence. They assert that metaphysics is worthwhile only when it aims to unify different scientific hypotheses, and that this, of course, cannot be done in ignorance of science. “Metaphysics, as we will understand it here, is the enterprise of critically elucidating consilience networks across the sciences.” (28)

These are serious challenges, worthy of serious responses. Moreover, though the details are new, they fit into a venerable tradition of anti-metaphysical thought stretching back through the logical positivists to Hume. Yet it is striking that this challenge from science is not the methodological issue which most preoccupies metaphysicians right now. Metaphysics as a sub-discipline of analytic philosophy has flourished—or run amok—over the last few decades, drawing inspiration from the work of David Armstrong, David Lewis and D H Mellor on topics like causation, laws of nature, time, persistence and modality. But the sub-sub-discipline *du jour* is ‘metaontology’ or ‘metametaphysics’: friends and foes of metaphysics have devoted conferences, collections and collaborations to thinking directly about the nature and scope of metaphysics.

On such occasions, the main threat to standard metaphysics is taken to be some form of Carnapian conceptual-scheme relativism about existence claims. Do middle-sized objects—salamanders or samosas—exist, in addition to the more fundamental entities out of which these things are apparently composed? Various authors have argued that, strictly speaking, there are no composite objects; others have disagreed. The metaontological challenge to both sides focuses on that notion of ‘strictly speaking’—is there a metaphysicians’ sense of ‘exist’ which is in some sense more basic, more revealing than the everyday sense of the word? This challenge is raised, and met, using tools drawn primarily from the philosophies of language and of logic.

My guess is that the authors of *EMG* would be as impatient with this debate about the nature of ontology as they are with its target, the ontological debate about composite objects. For them, the problem with standard metaphysics and ontology is not that it tries to go beyond everyday language and common sense; the problem is that it sticks much too closely to everyday language and common sense, more closely than many of its practitioners realise, because they do not realise how very far modern physics has already moved away from common sense. The claim that there are only fundamental, indivisible particles, not composite objects, looks like a radical one. But it’s not as radical as rejecting the very coherence of the idea of big things’ being made out of little things, and denying that there are any such things as fundamental, indivisible particles.

For me at least, this vigorous reassertion of the challenge from science is a welcome redressing of the balance in the metametaphysical debate—after all, it is science which is the ‘great epistemic enterprise of modern civilization’ (310), the

closing sentence of the book), and it is with respect to science that we need to position ourselves as philosophers, not just as metaphysicians. That said, I want now to argue that neo-scholastic metaphysicians are useful members of philosophical society. In particular, I will argue that even if we accept *EMG*'s story about the primary goal of metaphysics—unifying science, or at least critically elucidating its consilience networks—analytic metaphysics as it is currently practiced, even by those who know little of science, has much to contribute to this project.

As a preliminary, I note that *EMG* often over-estimates the degree of consensus amongst metaphysicians, and the degree of confidence with which metaphysical beliefs are held. For example, on p. 151 we are told that 'standard metaphysics' assumes, amongst other points, the truth of both Humean Supervenience and the Principle of Identity of Indiscernibles, whilst on p. 261 'neo-scholastic metaphysicians' are to be found collectively embracing mereological atomism. On p. 202 'conservative metaphysicians' assert that 'a thing cannot have causal efficacy over and above the summed causal capacities of the parts with which it is allegedly identical'. Yet all of these doctrines are up for debate in contemporary metaphysics, where consensus about substantive theses is rare. (How could we hog so much space in the journals if we just nodded quietly in mutual agreement?)

Much of what has lasting value in the metaphysical literature, as in many areas of philosophy, consists in the careful, detailed work of distinguishing conflated questions and issues, investigating the space of possible theories, and establishing relationships of confirmation or even entailment between different claims. Can you be both a four-dimensionalist about time and a three-dimensionalist about objects? Do four- and three-dimensionalism exhaust the options? Can you be an anti-realist about numbers without being a realist about spacetime points? David Lewis is perhaps the most widely-admired and influential metaphysician of recent times. But he is admired and influential not because he is generally agreed to have established that there are many concrete worlds, across the closest of which all facts supervene upon point-by-point property-instantiations. Instead, his contribution was to show what could be done with the resources of concrete possibilities and similarity relations, how these could be parlayed into accounts of properties, causation, persistence and the like, what sorts of explanation are available against such a backdrop, how we might think of chance and determinism, and so on.

Similarly, Peter van Inwagen hasn't brought many people round to his view that the only composite objects are living organisms. But he has clearly shown us the differences between asking what relations hold amongst objects when they compose something, what relations hold between whole and its parts, and what property an object has when it has proper parts; moreover he has shown us how answering any one of these questions can leave the others unanswered (this point is relevant to the dismissal of mereology on p. 21). Hugh Mellor hasn't turned us all into B-theorists about time, but he has shown us how to tease apart questions about the indispensability of tensed language from questions about the reality of tensed facts (a distinction endorsed as if obvious on p. 163).

Well, who cares about any of that? The authors of *EMG* should care, because, like Lewis, van Inwagen, Mellor, and the rest of us, they are interested in the nature of objects, properties, relations, time, modality, causation and so on, even if

sometimes they prefer different categories and terminology. Much of *EMG* is concerned with developing a positive metaphysical view, one based on the findings of current science and aimed at unifying those findings, so far as is possible. And there are many points at which the conceptual tools and distinctions developed by neo-scholastic metaphysicians would have come in handy. In what follows, I will outline a few examples piecemeal. Inevitably this involves quoting out of context and uncharitably highlighting what look like errors, but this is a method the authors themselves use in their attack on neo-scholasticism.

Time crops up on several occasions. In Sect. 4.3 a ‘block universe’ view (taken, I think, to be the combination of realism about nonpresent events with the denial of ‘temporal becoming’) is taken to be at least superficially incompatible with the existence of temporally asymmetric influences, or real causal processes. It is said:

And if causal processes are taken to require actualization of effects after actualization of their causes, then there may be an important sense in which nothing is a causal process at the level of abstraction suitable for metaphysical unification, since physics motivates the hypothesis that we may live in a block universe (as we argued in Chap. 3), and then an acceptable metaphysics will have to be set within such a universe. (211)

I find this quite baffling: if denial of temporal becoming undermined temporal asymmetry, or the distinction between ‘before’ and ‘after’, we’d all accept temporal becoming. If on the other hand it’s the notion of ‘actualization’ that’s important, then I’d like to know what this means, and how to understand it without turning to neo-scholastic or even scholastic metaphysics.

Jonathan Lowe doesn’t get a great press in *EMG*, though he is co-opted to the side of the angels when he challenges Trenton Merricks, whose philosophical sins are mortal (22–23). Yet Lowe (along with Kit Fine) has done significant work distinguishing different ways in which one entity can depend upon another—he clarifies notions such as ‘identity dependence’, ‘ontological dependence’, and ‘essential dependence’, and distinguishes an object’s essential properties from the properties it has necessarily. This is surely as neo-scholastic as it comes, but distinctions like these would have helped the authors enormously in their extensive discussions of the ways in which higher-level entities depend (or don’t) upon lower-level entities, and the ways in which structures are somehow prior to objects. The authors are concerned to deny that there are any ‘self-subsistent’ entities, but it is deeply unclear what this denial amounts to. Moreover, I often found myself wondering how there could be *naturalistic* facts about what depends upon what, going beyond extensional facts about what exists.

On pp. 254–255 we read that the authors ‘are sanguine about ontological vagueness...[and] regard puzzles about where exactly mountains stop and start, and whether or not the table is the same if we remove a few particles from it, as pseudo-problems *par excellence*, of no scientific or factual relevance.’ A similar point is made on p. 229. But nobody thinks we can work out exactly where mountains start and stop, nor that you destroy a table when you sand it. It’s a big step from there to accepting ontic vagueness; the interesting, genuine problems about vagueness concern its origins (in thought, language or elsewhere), and the question of how to

reason correctly using vague terms, why slippery-slope arguments are fallacious, and so on.

The positive metaphysics advanced in *EMG* crucially involves the notion of ‘real patterns’, an idea developed out of Daniel Dennett’s work. ‘From the metaphysical point of view, what exist are just real patterns’ (121). A key goal is to integrate the traditional categories of object, event and process (much as neo-scholastic four-dimensionalists about objects do). I have failed to work out whether or where individuals (i.e. objects, I think) fit into the overall picture. Sometimes, individuals seem to be real patterns, i.e. real as can be. But sometimes ‘Individual things, then, are constructs built for second-best tracking of real patterns’ (242), or ‘epistemological book-keeping devices’ (240), or then again ‘All individuals, we will argue, are second-order real patterns’ (243), i.e. real patterns, i.e. existents. Things, on the other hand are ‘locators’, where ‘a locator is to be understood as an act of ‘tagging’ against an established address system’ (121), i.e. (I think) something at the level of sense or language, not reference or the world. The authors defend a form of ontic structural realism—this must surely involve the claim that structures are real—yet on p. 299 we read that ‘structures describe real patterns’, which suggests that structures themselves are merely a representational device.

Now, this is a long, detailed book, and the authors take themselves to be advancing a revolutionary metaphysical view, one which cannot easily be expressed in traditional terms. So it is hardly surprising that I can juxtapose brief quotations from different chapters and thereby create a puzzling impression. Moreover, it is hardly surprising that I can find a few errors relating to my own areas of expertise in a book which draws upon so many areas of philosophy and science. But my point is that the literature of contemporary analytic metaphysics contains a wealth of resources, distinctions and concepts which the authors might profitably have used in developing and setting out their positive views, even if just by way of contrast. After all, one standard, very modest view of the role of philosophy with respect to science is that of conceptual clarification, and this is a role which seems compatible with the conception of metaphysics as the unifier of science.

So, for example, I would be interested to know what relationship there is between the metaphysical view advanced in *EMG* and the following sketch:

The fundamental facts about the universe are facts about which properties are instantiated where; these can include irreducible facts about properties instantiated across extended regions, not just at points. We humans find certain regions interesting, and think of these as things; often these are regions which have some kind of thermodynamic stability, but beyond this our interest does not reflect a fundamental ontological feature of the world.

Whilst of course there’s much more than this in *EMG*, this basic picture is familiar to neo-scholastic metaphysicians, though by no means universally accepted. Indeed it’s not so very different from what you’d get if you took Lewis’s picture of actuality and subtracted his Humean Supervenience, a subtraction that Lewis himself was willing to make in the face of quantum entanglement (as noted on p. 148).

I have offered only a very limited defence of contemporary metaphysics here—I have not argued that we metaphysicians possess a special method of discovering the

nature of reality which can rival or even conflict with the methods of science. But I think it is evident that the actual work of many contemporary metaphysicians provides conceptual resources and tools which can be of great use to anyone who is attempting—admirably—to draw metaphysical conclusions from the detailed study of current science.

So should we all go off and study physics? Yes, indeed: it's always worth finding out more about the world, modern physics can be mind-blowing, and those philosophers who have the skills and background to really engage with contemporary science have the potential to teach the rest of us an enormous amount. But studying the details of science is not the only way in which we can contribute to the great epistemic enterprise of modern civilization. Even on the modest 'conceptual clarification' model of metaphysics, there is a place for blue-skies thinking, for ranging freely through the possibilities without an eye to what can immediately be applied to the critical elucidation of consilience networks. We mustn't let the short-term demands of commercialisation undermine our long-term potential to contribute to the intellectual economy.

Authors' response

James Ladyman & Don Ross

It is ironic that Katherine Hawley begins by commenting that our book is 'wildly ambitious', given that she represents analytic metaphysics and its project of describing the fundamental nature of reality by the power of pure reason. From our perspective the ambition of contemporary metaphysicians is much greater than ours. Where we seek to synthesise into a unified picture insights from fundamental physics and the special sciences, based on the empirical knowledge accumulated by a vast collective effort, many metaphysicians imagine they can discover deep general truths while paying no attention to the results and conclusions of their colleagues in the lab. This is not just ambition but hubris. It says a great deal about the state of the subject that Hawley feels the need to address the question of why it matters that many metaphysicians know very little about science.

As Hawley points out, metaphysicians are these days worrying a good deal about the nature of their subject and its methodology. Their deliberations to date are not reassuring. One of the most popular apologies for the debates about whether ordinary objects such as tables exist has been to invoke a special metaphysical or fundamental notion of existence, according to which tables might not EXIST even though they clearly exist. This is evidence of the degeneration of the subject, reminiscent of postmodernist nonsense about truth with a capital 'T', and invites popular parodies of philosophy as so much irrelevant wordplay. While Hawley is right that our claim that it is time to abandon the 'made of' metaphor in metaphysics is radical, the difference between it and the claim that only simples EXIST is that the former is motivated by what we take to be the lessons of our best science while the latter is not. Of course, we may well be wrong in the conclusions we draw from science, but that is another matter.

Despite our ‘infuriating’ polemic, Hawley has the good grace to welcome the challenge it contains and we are grateful for her engagement with our arguments. Rhetoric aside, the differences between her and us are in the end largely matters of emphasis and degree. Generously granting us the correctness of our characterization of metaphysics as the project of unifying the sciences, Hawley argues that metaphysicians can contribute to this project. Of course we agree, since we do not urge the elimination of metaphysics. The question is whether *a priori* metaphysics has any prospect of uncovering objective knowledge. This we deny. Hawley hints, without quite saying, that she’s not convinced of this either. The role in knowledge-building she thinks the analytic metaphysician can retain is that of drawing careful conceptual distinctions. *Up to a point* that must be right; all the observation you like doesn’t get you any science if you can not make abstract distinctions. The problem with most concept-mongering in analytic metaphysics, however, is that it is done without regard to science, yet in the expectation that scientific observations and generalizations will turn out to fit the templates it conjures. They almost never do. All too often, when the failures of fit emerge, the frustrated metaphysician rejects the recalcitrant science as being ‘conceptually incoherent’. We cite several instances of this in the book, all from leading philosophers.

It is for this reason that in our book we *don’t* borrow the carefully chiselled distinctions that Hawley says—no doubt rightly—would help her understand the positive part of our view. They’d even help *us* understand it better. But our point is that the duty of the naturalistic metaphysician is to work in the intuitively vertiginous conceptual space into which contemporary physics forces us. This is indeed a terrific workout for our intuitions, which is to say one inevitably feels bewildered at times. But that is the cross that must be borne by metaphysicians whose minds were built by natural selection, which was under no pressure to make good quantum physicists.

Hawley is probably right that we overstate the agreement among metaphysicians about the theses we are concerned to undermine, such as Humean Supervenience. However, we do think that even those who disagree with it take it far more seriously than is warranted. We remain convinced that the containment metaphor exerts a pervasive and pernicious influence on metaphysical thought, and in the book we provide a range of examples we cannot use scarce space to reiterate here. When it comes to Hawley’s account of what is of lasting value in the literature we are bemused by the examples: “Can you be both a four-dimensionalist about time and a three-dimensionalist about objects? Do four- and three-dimensionalism exhaust the options? Can you be an anti-realist about numbers without being a realist about spacetime points?” These are utterly dis-similar to any questions scientists take seriously. What empirical predictions, however, indirect, arise from antirealism (or realism) about numbers? Only if antirealism about mathematics inspires a programme in mathematics that contributes to the subject is it worth taking seriously.

Hawley’s concern that we work with a still incomplete account of structure is surely justified. Our book implicitly calls for new battalions of naturalistic metaphysicians to engage exactly that project. Thus we welcome Hawley’s incisive critical questions on this score. Our other two critics raise similar issues. So, with apologies to Hawley and excuses about the constraints of space available to us here, we move onto them.

Paul Humphreys is sympathetic to our basic orientation, and helpfully suggests empirical arguments against individualism even as applied to classical particles. We welcome his lovely example of the Gibbs paradox. Humphreys also endorses our opening attack on the a priori nature of analytic metaphysics. We strive instead for a metaphysics that is based on empirical evidence. Some of Humphreys's critical remarks on our book suggest that he falls short of appreciating quite how seriously we take that commitment.

As empiricists, we do not take for granted that the world is unified. Therefore, we don't take for granted that it is unifiable by any models that evidence will turn out to best support. We don't take for granted that naturalistic metaphysics—and therefore metaphysics in general—will ultimately prove useful. We do aim to show, however, that van Fraassen and Friedman have pronounced it hopeless prematurely. Our starting point is that there is a legitimate role for metaphysics just insofar as the world is unified. This, and *not* the history of cases of explanations of special-science phenomena by reference to physical structures, is the basis of our privileging of fundamental physics in the PPC. The history of cases isn't intended to support an inductive generalization to the primacy of physics. It is merely intended to remind us that scientific practice that has assumed an asymmetric relationship between physics and special sciences has often worked out very well. Scientific practice that has actively *denied* this asymmetry has worked out terribly. We do not, and need not, deny Humphreys's point that plenty of good science *ignores* the asymmetry. That is why we are not physicalists, as Humphreys notes.

It is true that we think, as he says, that what is primary for purposes of metaphysics is fundamental physics. For us, this comes close to being a merely semantic point, since what we mean by 'fundamental' physics is that which is relevant to (partially) unifying *everything* because every measurement at every scale implicitly tests its generalizations. We do not know which parts of physics will turn out to be fundamental in this sense and which won't. We think that it's possible to identify some overwhelmingly *unlikely* candidates—special relativity, for example—but we do not think we'll ever positively identify the boundaries. We do not expect that fundamental physics will turn out to be equivalent to the domain of quantum physics, for example. It's interesting to discover what might be fundamental and what demonstrably isn't, and that's the kind of empirical discovery that metaphysicians should care about most, because it's directly relevant to modality.

We may not know that the world is unified, but we provide reasons for working with the assumption that it is. Here our verificationism, which Humphreys doesn't mention, plays an important role. Suppose the following very simple disunity hypothesis were offered: there is one (obviously non-fundamental, in our sense) physical structure **S*** that is radically autonomous in the sense that no discernible principles link it to what would count as fundamental physics if we re-defined that to refer to structures that constrain every measurement taken in the *rest* of reality. ('Almost' fundamental physics, if you like.) For this thought experiment to make sense, we need to face the question: what supports the flow of information about **S*** to observers elsewhere in the universe? If the answer is 'nothing' then **S*** is a perfectly idle conjecture, beyond science and therefore also beyond metaphysics. If some structure *does* link information about **S*** with information available in the rest

of the universe, then fundamental physics should constrain the characterization of this structure and S^* 's autonomy is only partial. Recall that our verificationism isn't supposed to itself be something established as true. It's just a stance that seems very sensible because if you don't adopt it you can conjecture worlds without limit and indulge in analytic metaphysics that will appear useless to everyone else. This is why we start EMG with the attack on analytic metaphysics: you're supposed to be persuaded to adopt our verificationist attitude because of the foolishness you risk indulging in otherwise.

If we thought that every useful scientific claim about partially or relatively autonomous systems would be explained by fundamental physics then we *would* be physicalists. We don't expect a fundamental physics that explains the specific nature of chairs, and so we don't expect metaphysics to explain that either. (We do think metaphysics should help to tell us why its useful for some observers to abstract things like chairs out of the underlying dynamic structures, though.) We take the same attitude to Humphreys's heroic life-saver and to the American dog tick. We *agree* with Humphreys that metaphysics doesn't contribute much of interest to understanding anthropocentric science. If our metaphysical story, or something like it, is best, then it blocks inferences, such as Dupré's, from anthropocentric science and highly autonomous systems to metaphysical disunity. We don't expect that to make a big difference in your life if you're an ecologist or an epidemiologist. We think it might be more interesting to you if your science is highly abstract, like economics, and you're afraid that it might all be empirically empty mathematics. Then you might be engaged in Humphreys's "automated" science, a context in which he grants that ontological structural realism is "persuasive".

Humphreys reads us as supposing that the institutional filters of science always work well. We believe no such thing. Rather, we believe that no *other* institutional filters on would-be objective knowledge work reliably *at all*. That said, we think that Humphreys's example of a filter gone wrong, standard microeconomics, is very badly chosen. As argued in Ross (2005) and many other places—Binmore (2009) is especially clear on the point—economists' sometime political commitment to self-interest as the univocal underlying motivation for all choice has *never* done any real work in economic theory, and the results reported in Henrich et al. (2005b) imply no revision in that theory. Space prohibits pursuing that argument here, however.

Finally, after accusing us of "intellectual evasion", Humphreys challenges us to "significantly clarify" the nature of structures in our reply to him. Stanford presses us hard on the nature of structures too, so with luck our answers to Stanford might go some way to assuaging Humphreys's sense of frustration. But on the specific point that Humphreys raises here: we simply see no plausible reason to believe that someone will one day come along with a good empirical motivation for identifying a clear boundary between the domain of non-physical mathematics and the domain of fundamental physics. (There are parts of mathematics that clearly don't *seem* to be describing anything physical, but let us never forget Kant's embarrassment over non-Euclidean geometry.) That point doesn't bring us closer to saying what structures are, but it explains why we don't think it's helpful to try to come down hard on one side of a debate over whether they're mathematical or physical. We

grant that it would be nice to know, but then we'd like to know where they put Jimmy Hoffa too. We just don't.

P. Kyle Stanford also endorses our critique of analytic metaphysics and we welcome his denunciation of it. However, Stanford is just as eloquent and sharp in his criticism of our positive views. We are grateful for his penetrating and challenging discussion and can only begin the task of responding to it here.

Stanford's first complaint is that we offer a conception of metaphysics according to which it ought to be derived from science rather than being simply indistinguishable from it. He explicates the latter idea in terms of Quinean naturalism as represented recently by Penelope Maddy, and argues that "this approach is straightforwardly precluded by Ladyman and Ross's two-stage model of inquiry, in which we must first ascertain what bona-fide science has independently concluded about nature before undertaking the distinctively metaphysical unifying project that tells us what these "separately developed and justified" findings really imply concerning the most fundamental constitution of the world." Clarification is needed here however. Some interesting problems descend from a period before modern fundamental physics stabilized, such as whether all physical influences propagate locally, and are commonly regarded as metaphysical merely by historical courtesy; yet the evidence relevant to their answers lies wholly within the domain of physics. Had the course of scientific development vindicated physicalist reductionism, it would follow that *the whole* of metaphysics was either dissolved or supplanted by fundamental physics. However, the failure of physicalist ambitions leaves room for a kind of metaphysics distinct from any of the special sciences. This is the project of unifying science, that is, of articulating a picture of the world on the basis of science as a whole. The specific metaphysics that we propose for the job, namely ontic structural realism (OSR), may stand or fall—and Stanford, does a good job of trying to knock it down, on which more below—but we contend that either way we have offered a positive conception of naturalized metaphysics and its distinctive task that is not bound to the fate of our particular execution of it.

Stanford's main problem with OSR is that "there is no unequivocal notion of "structure" that can do all the jobs they want and need it to do". Stanford is right that we assign three such jobs: first, to replace objects and individuals in the ontology of fundamental physics; second, to be that which is preserved across theory change in the history of science; and finally, to replace unobservable entities in the explanans of the no-miracles argument. He clearly explicates our claim that past scientific theories whose ontologies have been superceded—such as caloric theory, the ether theory of light and phlogiston theory—nonetheless described modal relations among the phenomena that have been retained by subsequent science. As he says, the novel predictive success of such theories over and above their systematization of known regularities, is what motivates our realism about modal structure and unwillingness to rest content with van Fraassen's empiricist structuralism according to which brute regularities are all that are described by scientific theories. However, Stanford objects that "to say that our theories correctly describe the modal structure of the world (in this sense) is simply *to redescribe or restate and not to explain* the fact that they enjoy novel predictive success", and so our appeal to modal structure to explain novel predictive success

amounts to nothing more than word play. This is a line of attack to which van Fraassen would of course be very sympathetic, since he has long argued that realist ‘explanations’ of the success of science are empirically empty. Our response is to point out that the commitment to more than regularities among the phenomena has methodological implications for the practice of science, namely, that scientists who believe that the theoretical laws of past theories describe the modal structure of the phenomena, in addition to the phenomenological regularities that are derived from them, have a motivation that a pure empiricist lacks for seeking to recover those theoretical laws as limits or approximations of successor theories. Since, actual science has progressed by doing just this, for example Einstein’s retention of the classical Poisson equation for gravity as a limit of the new equations of General Relativity, we take our view to better fit the data.

A related, but distinct, problem raised by Sanford echoes Humphreys’s main criticism. This is that the failure of physicalism blocks inferences from fundamental structures to the domains of special sciences. The point of Sanford’s example from the history of developmental biology is to suggest that, on independent grounds, OSR poorly describes the dynamics of theory change in special sciences that study emergent phenomena. These two problems are taken to leave us two options: (1) restrict the scope of OSR to the philosophy of fundamental physics, surrendering its unificationist pretensions, or (2) retain the Quinean programme for unifying the sciences (and metaphysics into the bargain) but abandon the hope that OSR is up to the job.

We have sketched a reply to the first issue, as far as space allows, in our discussion of Humphreys’s objections above. On the second issue, we think that Sanford’s example misses the target. We do not take conjectures about highly specific and contingent historical processes, such as “ontogeny recapitulates phylogeny”, as plausible descriptions of structures. Accounts of structures relevant to biology are mathematically formulated generalizations of information theory, ultimately to be partially explained by reference to fundamental physics, that tell us what kinds of information can and can not be transmitted by the types of processes of which biological processes are instances. The generalization about ontogeny and phylogeny *might have* been motivated in this way, but it *wasn’t*; it was an induction on a limited (and selective) range of observations, for which Haeckel conjectured a Principle by way of an explanation. It is indeed because special sciences are relatively autonomous that they routinely generate such conjectures, but the conjectures in question usually turn out to have been over-generalizations. By contrast, the rejection of doctrines such as vitalism is plausibly based on reflection on fundamental structure, and in that sense on what is physically possible. This is the only, limited, extent of unifiability to which our OSR and Rainforest Realism are committed.

Stanford’s more general worry about whether there is a single notion of structure of which we can avail ourselves is also important because it “threatens to simply turn the case for structural realism into a bad pun”. We think that this challenge can be met and remain convinced by the consilience case we developed in the book, but we cannot adequately respond to it here. It is worth pointing out however, that even van Fraassen agrees that there is structural continuity across theory change in

science, and that contemporary physics, and increasingly other sciences, do seem to represent the world by mathematical structures without clear ontological interpretations. Our scientism requires us to work harder to articulate the notion of structure we deploy, and we look forward to doing so with renewed vigour, having been inspired by the constructive engagement and penetrating criticism of Hawley, Humphreys and Stanford, for which we are extremely grateful.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Arntzenius, Frank, and John Hawthorne. 2005. Gunk and continuous variation. *The Monist* 88: 441–465.
- Batterman, Robert. 2002. *The devil in the details*. New York: Oxford University Press.
- Binmore, K. 2009. *Rational decisions*. Princeton: Princeton University Press.
- Bloch, Felix. 2000. *Fundamentals of statistical mechanics*. London: Imperial College Press.
- Gould, Steven J. 1977. *Ontogeny and phylogeny*. Cambridge: Harvard University Press.
- Floridi, Luciano. 2008. A defence of informational structural realism. *Synthese* 161: 219–253.
- Fredkin, Edward. 1990. Digital mechanics. *Physica D* 254–270.
- Haken, Hermann. 2006. *Information and self-organization*, 3rd ed. Berlin: Springer.
- Henrich, Joseph, et al. 2005a. ‘Economic man’ in cross-cultural perspective: Behavioral experiments in fifteen small-scale societies. *Behavioral and Brain Sciences* 28: 795–855.
- Henrich, J., R. Boyd, S. Bowles, C. Camerer, E. Fehr, and H. Gintis. 2005b. *Foundations of human sociality: Economic experiments and ethnographic evidence from fifteen small-scale societies*. Oxford: Oxford University Press.
- Humphreys, Paul. 2004. *Extending ourselves: Computational science, empiricism, and scientific method*. New York: Oxford University Press.
- Humphreys, Paul. 2009. The philosophical novelty of computer simulation methods. *Synthese* 169: 615–616.
- Kitcher, Philip. 1989. Explanatory unification and the causal structure of the world. In *Scientific explanation*, ed. Kitcher Philip, and Salmon Wesley. Minneapolis: University of Minnesota Press.
- Lewis, David. 1986. *Philosophical papers*, vol. II. Oxford: Oxford University Press.
- Lewis, David. 1994. Humean supervenience debugged. *Mind* 103: 473–490.
- Luhmann, Niklas. 1997. *Die Gesellschaft der Gesellschaft*. Frankfurt: Suhrkamp.
- Maddy, Penelope. 2007. *Second philosophy*. Oxford: Oxford University Press.
- Morgan, T.H. 1916. *A critique of the theory of evolution*. Princeton: Princeton University Press.
- Ross, D. 2005. *Economic theory and cognitive science: Microexplanation*. Cambridge: MIT Press.
- Saunders, Simon. 2006. On the explanation for quantum statistics. *Studies in History and Philosophy of Modern Physics* 37: 192–211.
- van Fraassen, Bas. 2005. Structure: Its shadow and substance. *British Journal for the Philosophy of Science* 57: 275–307.
- van Fraassen, Bas. 2002. *The empirical stance*. New Have: Yale University Press.
- van Inwagen, Peter. 2001. *Ontology, identity, and modality*. Cambridge: Cambridge University Press.