Reorienting Realism

M. Chirimuuta

Contribution to Symposium on Reconceiving Realism.
Meeting of the Philosophy of Science Association, 10th November 2022

SLIDE 1
My stance on scientific realism is ambivalent. I am not reconceiving realism in order to defend it, but in order to characterise a tendency that I think we need to do away with. At the same time, there is an element of realism which I would like to retain. I'll say more about that towards the end of the talk.

SLIDE 2
But firstly, I will discuss the reorientation; secondly, an alternative to realism; thirdly, what's minimally realist about this alternative; and finally, some comparisons with the other views on the table this afternoon.

SLIDE 3 – START SECTION 1

SLIDE 4
Once upon a time Philosophy of Science was all about Physics. Debates about Laws of Nature, Deductive Nomological Explanation and Theory Change all grew up around questions that occur naturally when considering physics, and were retrofitted for the biological and social sciences. The debate over Scientific Realism centred on the question of the Existence of Unobservables, one that occurs naturally for physics and chemistry, but seems irrelevant to many branches of biology.

SLIDE 5
Nowadays, Philosophy of Science is as much about the special sciences as it is about physics. A consequence of this has been the Downfall of D-N Explanation, and the Rise of Mechanistic Explanation. Should we think that there is no realism debate for most of the special sciences, because there are not enough unobservables? And that in these domains realism wins by default? No, what we need to do is reorient the realism debate so that we can better see what the issues are for the special sciences. At the heart of the realism debate is a basic curiosity about whether *the best science tells us how the world <u>really</u> is*? This question *can* meaningfully be raised across the board, as I'll show in the course of this talk, after a few more preliminaries.

SLIDE 6
What's peculiar about Physics is that it investigates simple stuff: nonliving matter, when not undergoing chemical reactions, is quite homogenous and unchanging, and insensitive to context. As well as dealing with the simplest parts of nature, there does seem to be a *belief* in simplicity that animates the history of physics. I call this the *Deep Platonic Assumption*. It's the idea that behind the messy, changeable appearances of nature, there is a simple, unchanging order that can be represented mathematically.

This is a realist intuition. As Ernst Mach said, the aim of physical theory is to represent the appearances in as economical a way as possible. The realist goes further than the empiricist in taking the compact representation – for example an equation describing gravitational attraction – to be converging on an order or pattern in nature that underlies the very variable instances in which gravity is observed.

As a historical point, realism about modern science is as much indebted to Ancient Greek atomism as it is to Platonism. Atoms, as originally conceived, were unchanging simples whose movements and patterns of coalescence were posited to explain the changing appearances of the natural world. As much as with Platonic forms, the guiding idea is of a simple order, that is not manifest in the perceivable, empirical world, but which we need to know about in order to make sense of the empirical world. For reasons that come later in my commentary for the round table, I'm happy to give this tendency the label, "Platonic".

SLIDE 7
The objects of the special sciences are much more complicated than the objects of physics. In what follows, I use neuroscience as the source of my examples, but I think that the story about reorienting realism applies to all the sciences of very complicated things.
My work on realism is part of a book project on modelling in neuroscience, which takes it that the central challenge of the scientists' research is to find effective ways to simplify the brain. The brain is an ultra-hyper-complex object! It is highly heterogeneous, its parts are always changing, and it is highly sensitive to context. One definition of complexity, due to Murray Gell-Mann, is that these are systems in between randomness and order: there is pattern there, but never exact repetition. The weather is like that. So are financial markets. This is a useful way to understand complexity in the brain. The most challenging way that the brain is complex is in its changeability. Exact science needs fixed targets but the brain is constantly changing itself in response to how things are in the body and the rest of world – which is necessarily so for it to have a role in sustaining intelligent behaviour. To act intelligently is to adapt effectively to what is going on in the changeable, complicated world around you.

All neuroscience theories and models are a drastic abstraction away from this complexity.
A perennial challenge for neuroscientists is to try to figure out which of the details matter to their explanatory and technical projects.
ALL neuroanatomy is an abstraction. Even this dense reconstruction based on electron microscopy is leaving out neuroglia. These are the more numerous cells in the brain, which probably DO play a role in cognition but are normally just ignored because they are not electrically excitable. In most of neuroscience, glial cells are treated as irrelevant details.

SLIDE 8
Moving from anatomy to physiology, here are some simple models of neuronal activity. The Hodgkin-Huxley model abstracts away from the molecular mechanisms involved in electrical excitability in neurons. The Point Neuron model assumes that the actual structures of branching neural fibres are irrelevant to information processing. (Incidentally, all artificial neural networks in deep learning AI make this assumption.) Models like this are false representations of empirical facts about neurons, neglecting or distorting many details that are observable to the naked eye or through light microscopy.

The question of realism turns precisely on whether you think that abstractions and idealisations, in their departure from apparent truth, are actually hitting on a deeper, hidden truth about how the brain works. This is how the Deep Platonic Assumption plays out in neuroscience….

SLIDE 9
"What makes a model good? Clearly it must be based on biological reality, but modeling necessarily involves an abstraction of that reality. It is important to appreciate that a more detailed model is not necessarily a better model. …. No matter how detailed, no model can capture all aspects of the phenomenon being studied.… [M]odeling is the lie that reveals the truth."

To be a realist about neuroscience is to believe that in the best cases, in virtue of their abstractions and idealisations, the theories and models reveal hidden but essential features of brain processes.

SLIDE 10
Therefore, to resist realism is just to reject this Deep Platonic Assumption. It is to assert that nature is irredeemably protean and complicated, and that *things are <u>not</u> more simple than they look.*

The point of this talk is to demonstrate how the issues play out in this reorientated debate. At the same time, I'm making it clear where my sympathies lie, even though I don't have time to give you the full justification or even say much about the motivations for my opposition to realism. What I want to convince you of today is that this is a useful way to reorient the realism debate for the special sciences so I will address some issues that may come up at this point.

SLIDE 11
One concern is that I am confusing the Appearance/Reality distinction with a distinction between Complexity and Simplicity.

My response is that the traditional realists I'm opposing (but not all of the self-declared realists in company here today) are committed to the idea that science can in principle discover absolute or mind-independent properties of objects, rather than human-relative ones. Absolute properties are inherently simpler than relational ones. This notion of a mind-independent reality characterised by absolute properties is de facto simpler than a world of objects bearing relational properties, including but not restricted to human-relative ones. My point is that the Platonic (and atomistic) appearance-reality distinction is fundamentally about the contrast between the human-relative, complex, changeable, manifest world, and an absolute, human-independent, simpler order. It's beside the point to mention that many of the entities and relationships now described in theoretical science are themselves very complicated.

Another objection is that the question of complexity does not divide Realists and anti-Realists because realists also think that nature is very, very, complex, but that the best theories and models may accurately represent just a little bit of it.

The opponent to realism in this new debate is going after something more radical than partial realism. Their idea is that changeable, ultra-hyper complex objects are not simple enough to be accurately represented in intelligible models – ones that make sense to finite, human cognisers. The partial realist assumes that nature can be broken down into small enough chunks for the human intellect to swallow. Their opponent is claiming that nature is, in itself, even when taken piecemeal, far more complicated than our conceptual schemes can handle, and that we only make sense of things in virtue of abstractions and idealisations that depart from the truth of how things are from a God's eye point of view. If that's the case, the anti-realist says, we should not think of the models as giving true but restricted representations of mind-independent reality.

SLIDE 12
You'll see that the alternative to realism rests on the thought that nature – the object of scientific enquiry -- is immeasurably complicated. You might then worry about the status of this claim. Is it Scientific or Metaphysical? What justifies it?

We can think of it as a ground level empirical fact that can be observed about the natural world. Let's call things *empirically accessible* when they can be observed and monitored through microscopes, imaging machines, and other devices, as well as perceived through ordinary senses. The more is learned about empirically accessible nature, the more complicated it appears. I think this is a faithful report of what is being found out in neuroscience and other branches of biology: the more probes are used, the more changeability and context sensitivity becomes apparent. Nature – as far as it comes across through empirical observation – is immeasurably complicated.
Realism says that there's more to the world than the empirically accessible world: there is, in addition, a simpler order that scientists represent in their elegant theories and parsimonious models. The anti-realist denies this.

SLIDE 13
It is helpful to compare my reorientation of the realism debate with how Catherine Elgin describes the dispute between realists and 'constructive nominalists'.
She writes,
"The issue that divides [traditional] realists and nominalists is whether—or perhaps to what extent—such scientific work is creation or discovery. Do the scientists draw the lines, or discover where nature draws them? Realists maintain that nature draws the lines. The scientists and other investigators simply find what is there to be found. The problem, though, is that there are too many things there to be found." Elgin (2019:523)

Like me, Elgin motivates her alternative to traditional realism with the problem of complexity, that there are "too many things to be found".
One refinement is that I want to resist a dichotomy of creation *versus* discovery. More will be said in Section 3 about why I assert that scientific representations are not only constructed or created, but also constrained by their objects of investigation.

First, I'll give you an illustration of how the reoriented realism debate works when we consider the theoretical claims of neuroscience. Here, a large part of the theoretical content of the science consists in computational models of the brain.

SLIDE 14 – START SECTION 2

SLIDE 15
The old realism debate asked questions such as the following:

- *Do the entities posited in physical theory constitute a mind-independent reality standing behind the observable appearances?*
- *In short, do unobservables exist?*
- *Are relationships amongst entities actually as described in the theories?*

The new, reoriented realism debate asks instead:

- *Do the simpler kinds and relationships depicted in scientific theories and models constitute a real order hidden behind the manifest complexity of empirically accessible entities and processes?*

Since Psillos (1999) it is customary to think of the realism debate as divided up in to three strands, an ontological, epistemic, and semantic one. It's helpful to do this here.

SLIDE 16
The ontological question is whether at least some of *the simpler kinds and relationships depicted in the best scientific theories and models have a mind-independent existence. For example, are there basic patterns or essential principles of neural coding that the best work has discovered, that exist independently of the simplifications that come about through experimental practice and through modelling techniques such as idealization?*

I call the positive answer to this question "formal realism", and the negative answer "formal idealism".

SLIDE 17
Here's an example to show what's at stake between "formal realism" and "formal idealism" in computational neuroscience.

This is a deep convolutional neural network model of visual processing in the ventral stream of the primate brain.
Note that these scientists don't think they've arrived at a highly accurate model yet. But the realist idea is that as models improve, this improvement consists in it converging on the actual computations performed by the primate visual system. Many suppose that the structure represented mathematically as a computation (i.e. a series of state transitions) is there in the brain processes – IF the model is accurate. I call this supposition of a structural similarity (or homomorphism) between model and target, *formal realism.* This stands in contrast to an alternative way to interpret the model which I call *formal idealism.* The view here is that whatever processes exist in the brain are vastly more complicated than the

structures represented in the computational models. The aim of modelling is to achieve an acceptable simplification of those processes, which helps the scientists to achieve their explanatory, predictive, and technical goals. Thus, the success of the research is more a matter of *structuring* a much too complex set of processes than of discovering pre-existing structures. A successful model is a compromise between the vast complexity of nature and the limitations and needs of human enquiry.

SLIDE 18
We'll now move on to the semantic strand. According to the formal realist, neural computational models should be interpreted literally as attempts at representation of a real mathematical order hidden behind the manifest complexity of neural anatomy and physiology. That is, a representation of computations actually performed by the brain.
The formal idealist holds, in contrast, that neural computational models should be interpreted **analogically**: there is a *constructed* similarity between brains and computations, that provides neuroscientists with a useful simplification of the brain.
The computational model represents an *ideal pattern.* This pattern is to some extent dependent on the activity of the scientists and is not a structure or phenomenon that exists in the brain independently of experiment and data processing.

SLIDE 19
This slide illustrates the contrast between the literal and analogical interpretations of neuro-computational models. The literal interpretation acknowledges that the computational model does not purport to represent the full causal structure of any brain area. The idea is that it represents a subset of the causal relationships thought to be responsible for a cognitive process.

SLIDE 20
This slide illustrates what is meant by an "ideal pattern". It is the content of the neuro-computational model, according to the analogical interpretation. It is not a pattern which exists independently of the scientists' experiment and modelling practices. It is a pattern which helps scientists make sense of brain data, by offering them an analogy with processes that occur in artificial computers.

SLIDE 21
And finally, the epistemic strand. The Formal Realist asserts that there is knowledge of the brain to the extent that the forms and relationships depicted in simplified neuroscientific models correspond to a real order hidden behind the manifest complexity of neural anatomy and physiology.
In contrast, the Formal Idealist holds that there is knowledge of the brain to the extent that scientists are able to strike a compromise between the unfathomable complexity of the brain, and their limited human capacities, developing theories and models that allow them to solve problems at hand (including biomedical applications) and progress their research. The emphasis is on the practical utility, and indeed necessity, of simplification. In a Kantian vein, an implication is that the *Brain-in-Itself (in its full complexity) is not an object of knowledge.*

SLIDE 22 – START SECTION 3

Formal idealism does not suppose that the structures given in scientific models are fully constructed or mind-dependent, but that there is an ineliminable human component in all scientific representations, due to the fact that they can never depict the full complexity of their target systems. As such, they are the result of human decisions about how to simplify. I now want to discuss how formal idealism is still a realism in an elementary sense.

SLIDE 23
The ineliminable human component of scientific representations is due to the fact that these representations are the product of an *interaction* between human investigators and a target system.
I like to use the sensory metaphor of touching (or haptics) for the investigative process. Models are the products both of constraints imposed by objects of investigation, and the constructive activity of scientists.
This means that I can deny full constructivism. Also, Formal Idealism is not Berkeley's Idealism!
Reality comes in the picture as what pushes back against human agency. It's better (more "realistic" in Hasok's sense) to construe reality this way, rather than as what is absolutely mind-independent.

SLIDE 24
The key Idea is that scientific models are like hands. Hands are both sensory organs for touch, and our means for manipulating the world.
These two roles are not separable. How we sense (or know) is conditioned by how we desire to make changes (or act), and vice versa.

Tools like this one extend the capabilities of hands. Scientific models are, likewise, tools. They are extensions of human sensing and acting capabilities.

SLIDE 25
With a notion of reality now in place, we can offset a worry about formal idealism, that it treats all pattern or form as a projection of the scientists. I am <u>not</u> saying that nature, independent of human categorization is 'formless', that there is no regularity whatsoever. For if the world was like that, science would be impossible because induction would be useless.
What I am saying is that the categorizations that scientists come up with (at least in the sciences of ultra-complex systems like the brain), are the result of a very large abstraction away from whatever patterns are 'out there', and this is necessitated by the task of rendering nature comprehensible to finite human knowers.

SLIDE 26
This position chimes with some recent work on natural kinds, following the 'practice turn' in philosophy of science. Kendig (2016) describes how there has been a break from the tradition of taking mind-independence to be the standard for the reality or objectivity of natural kinds. Reydon (2016:59) proposes an account in which there is "co-creation" of

natural kinds, so that "both the contributions from nature and from us to the making of classifications and kinds should equally be taken into account…".[1]

SLIDE 27
Conclusions

SLIDE 28
Thanks and Contact info – Please leave up for people to take down details.

PLEASE SAVE THIS SECTION FOR THE ROUNDTABLE PART OF THE SYMPOSIUM

SLIDE 29 – START SECTION 4
I will now touch on some comparisons with my fellow symposiasts, and regret not being in the room to properly discuss these points of connection and divergence.

SLIDE 30 –
I'll begin with the view furthest from my own, which is *Future Proof Science.*

I agree with Peter that the old realism debate is problematic because tied to the question of the existence of unobservables like electrons. My proposed reorientation is intended to speak to this worry. In a way, I have not completely abandoned the observable/unobservable distinction, but I want it morph into a division between what's empirically accessible and complicated, and what's represented in elegant theories and parsimonious models. I prefer this to recasting it as a debate between "believers" and "sceptics" about well-established scientific claims.
Peter calls his opponents "science sceptics". I don't want that label for the anti-realist position I'm articulating, because it is not a kind of scepticism. Instead, it is a critique of a realist claim for scientific knowledge – as knowledge of human-independent reality.

> Compare: Kant's epistemology was not "sceptical" about things in themselves. His point was that the wrong way to think about human knowledge is to take the objects of our knowledge to be things in themselves.

The colour example in Chapter 1 of *Future Proof Science* illustrates the point nicely. Peter tells us that strawberries are actually uncoloured, and that our visual experience is universally and systematically misleading. He is an 'absolutist' about colour: if colours exist, they must be perceiver-independent properties instantiated by strawberries, etc. I'm a relationist about colour: I think that colours exist, and that our visual experience is not systematically misleading, but that we must accept that colours are perceiver-dependent properties. The wrong way to think about human perception is to take it as aiming to reveal to us how things are in themselves, instead of their human-relative properties.

And the same goes for science. My position is relationist about science: there is scientific knowledge but we must understand this as knowledge of things insofar as those things

[1] See also Sullivan (2016) chapter from same book, *Natural Kinds and Classification in Scientific Practice*

interact with people and their equipment.[2] The kind of realism I oppose is absolutist about science: it says that science, at its best, gives us knowledge of human independent reality.

I also have one general observation about the value of future proofing. It's interesting that texts in the history of philosophy have better longevity than historical works in science. Aristotle's *Ethics* is relevant to today's ethicists in a way that no work of ancient natural philosophy is relevant to today's scientists. Does this mean that philosophy can be future proof? We have to say 'no' because you cannot pull out a list of 'facts' that Aristotle nailed down, and contemporary virtue ethicists still subscribe to, and so the example doesn't fit Peter's criteria. However, I bring this up because the failure of old work to go out of date is normally held *against* philosophy -- it is evidence that philosophy is not *progressing.* There's a popular view that scientific knowledge is special precisely because it's always progressing, which means you can't bank on any theory or fact being so well established that it couldn't be overturned. Does future proof science accommodate this idea about what makes science special?

SLIDE 31
Perspectival Realism seeks to establish perspectivism as a distinct and credible version of scientific realism. Earlier presentations of perspectivism, like Ron Giere's, faced criticism for not being realist enough. Since my position also grew from considering Giere's perspectivism, and probably stays closer to it, I'd like to say something about this comparison. The analogy between colour perception and scientific knowledge is of course Giere's. This is a construal of perspectivism that Michela now calls perspectival$_1$ and she considers it to have the worrying metaphysical implications that there are no non-perspectival facts, and that all properties are dispositional (or relational) as Chakravartty (2017) has pressed.[3]

I take Michela's account to be in the middle of a spectrum between absolutism and relationism: it insists that scientific knowledge is bound to the situated human perspective, but at the same time seeks to avoid the more radical implications of relationism and therefore recover a realism in which perspectival representations provide a "window on reality".

I think that Giere's original kind of perspectivism is better motivated and more defensible if we begin our enquiry with the thought that nature is immeasurably complex. The idea is that it is so complex and changeable that scientists must interact with it in careful ways in order to stabilise phenomena, and get representable facts. This is a way of saying that there are not perspective independent facts - i.e. facts that just hold independently of choices made about how to investigate certain portions of the world. That nature is not a passive, stable thing, waiting to be mirrored in scientific representations, but something that actively responds to processes of investigation is a point made salient in actor-network theory. It may seem like a weird, post-modern thing if your paradigm science is physics, but in neuroscience it is just stating the obvious. Brains of all animals plastically change with the

[2] See Giere 2006, Rueger, A. (2005) 'Perspectival models and theory unification', *British Journal for the Philosophy of Science* 56, 579–594.
Discussed in Massimi 2022:59ff
[3] See Massimi 2022 §3.4 and §3.5

experience of the whole animal or person in the lab, and neurons behave differently when subjected to invasive procedures and examined in vitro. It makes sense to think of all of the brain's properties as relational, and to ignore the question of how the brain is intrinsically, independently of its interactions with the scientist.

My concern about invoking the metaphor of the "window on reality" is that it reverts perspectivism to an account of scientific knowledge in which scientific activity is at its best a transparent medium through which a human-independent world can be viewed. The metaphor of touch is intended to break with that ideal: we have knowledge of nature in virtue of our activity, which means that the objects of knowledge are things that are affected by us. This is why it is fine to deny perspective-independent facts.

SLIDE 32
I'll now move on to *Realism for Realistic People.* This quotation from an earlier paper by Hasok really spells out the contrast between 'relationist' and 'absolutist' philosophies of science. Even though in this talk I've emphasised that our image of science changes if we start with a science like neuroscience, rather than physics, it's gratifying to find someone beginning with physics and chemistry, and still converging on many of the same concerns about traditional realism. [here for reference, no need to read out]:

- "what I am calling for is a change of viewpoint, toward a humanism in epistemology: instead of taking 'natural' as something transcending humans and seeking forever to erase all traces of humanity in our concepts, we can embrace the humanity of our concepts and assess their merits in terms of how well they enable human scientific inquiry. Nature conceived in the anti-humanist manner is impossible for us to access, though the aim of learning about such 'nature' may make sense as a sort of forlorn regulative ideal worthy of Sisyphus or Tantalus."
- Chang (2016:43) 'The rising of chemical natural kinds through epistemic iteration'

This quotation also brings up a small point of disagreement that I thought it would be interesting to mention. Here, Hasok is dismissing the goal of attaining knowledge of the human-independent world, even as a regulative idea. I'm not so sure about this.

SLIDE 33
Realism for Realistic People is in lots of ways close to *haptic realism,* but let me point out that the aims of our two books are quite different. Hasok aims to give an account of scientific knowledge that is useful not only to philosophers and broader communities, but to scientists themselves. The suggestion is that by giving up on absolutist ideals, and thinking differently about their acquisition of knowledge, scientific communities will be more empirically active and pluralist, and therefore more successful. My book is about neuroscience, and this brings up something of a niche aim which looms large for me: to figure out how philosophers should *interpret* scientific knowledge claims for the purposes doing philosophy of mind. The question of interpretation is also important for non-academic outsiders to science, and scientists when they popularise their work, but I'm not treating it as something that needs to, or should even, impact the working scientist. In other words, I'm not developing this account in order to prescribe it to scientists themselves. In fact, I think that scientists themselves might be better off sticking with absolutism. I suspect that the search for absolute truth, and adherence to the aim of unification and not pluralism,

have served scientists by giving them focus and conviction, and providing a heuristic to help them find regularities that are more stable and therefore useful for interventions, amongst all the patterns that are empirically accessible. I agree with Hasok that there is something other-worldly and impractical about traditional realism. Think back to what I called the 'Deep Platonic Assumption'. In Plato's philosophy, knowledge of the forms, of an absolute reality, was a condition of spiritual elevation. Modern science, insofar as it adheres to the Deep Platonic Assumption takes over the ambition of attainment of knowledge of an absolute reality, but its aims are worldly and not spiritual – they are to do with prediction and control of the material world. I think that what Max Weber said about the protestant ethic and capitalism is an apt way to think about absolutism and science: we have "worldly asceticism" – what were originally spiritual ambitions, and modes of discipline for suppression of our human impulses and interests, are now directed towards materialistic concerns. After pointing out that the God's eye view is unobtainable for science, we should ask ourselves whether achievement of the God's eye view was ever really the point of upholding realism about science. Maybe absolutism is the noble lie that science needs to tell itself.[4] I will leave it there.

SLIDE 34
Thanks and Contact info – Please leave up for people to take down details.

[4] As in this *methodological maxim*: 'Proceed as if there is a single world with a unique structure' (Giere 2006b, p. 36)