Reality Realism

Sean M. Carroll

1. *Introduction*

Morality & Mathematics (Clarke-Doane, 2020) opens with a characterization of "a common *naturalist* position among philosophers and scientists," which "combines realism about the sciences with anti-realism about value." As evidence of the latter, a quote from a representative physicist (me) is provided, blithely proclaiming that "There are not objective moral truths."

Then on the following page, we are told that "an empirical scientific realist would seem to need to be a mathematical realist as well," since "a typical empirical scientific theory, rigorously formulated, presupposes pure mathematical facts as well." Indeed, "If a naturalist like Carroll were to declare that he is a realist about, for example, the standard model of particle physics, but *not* about mathematics, then it would not even be apparent what he meant."

Strictly speaking, I am not a realist about "the standard model of particle physics." The standard model is a partial description, a representation that captures some aspects of how reality behaves, and only an approximate representation at that – extremely accurate within a certain domain, but completely inapplicable in others. What I am a realist about is *reality*, by which I mean the totality of the physical universe. The standard model of particle physics, like general relativity or Newtonian mechanics, provide useful ways of talking about reality in certain circumstances, but I would not describe them as fundamentally "real."[1] The same, I would argue, goes for mathematics generally, about which I am not a realist.[2]

The purpose of this paper is to make apparent what might be meant by that.[3]

2. *Reality By Itself*

[1] The view I am defending is thus related to a kind of structure realism, as elucidated for example in Ladyman et al. (2007). But my claims are narrower than a defense of any specific variety of structural realism.

[2] I do not mean to claim that almost all or even most physicists think this way, although I suspect many of them do. Most working physicists try not to think about philosophical issues regarding "reality" at all. And some are undeniably realists about both mathematics and the laws of physics. One of these is Nobel Laureate David Gross. I once asked him if he could point to Einstein's equation of general relativity out there in the world. He picked up a book, let it fall to the ground under the force of gravity, and announced, "There you go."

[3] There are of course others who have defended scientific realism while denying mathematical realism. See e.g. Field (1980), Azzouni (2004), or Szabó (2022). The present essay is offered as the perspective of a particular physicist, rather than a comprehensive examination of the issues.

Let me start by laying out the picture to be defended, before getting into possible objections and responses to it. The intuition we're trying to develop is that there is something called "reality," or "the world," or "the universe," or "nature" – the totality of existence. And the world is *sui generis* – it exists, in a way that no other kinds of things exist. (We can discuss along the way whether we should think that other kinds of things exist in other ways, or whether it's better to save "existence" as something that only applies to the world.) The world being described is supposed to be our actual world, but we will talk about it using language that is intentionally noncommittal concerning fundamental ontological questions, as those remain incompletely answered. There can of course be subsets of the world, or collections of such subsets, which may rightfully thought of as existing as well. I'm on the side of saying that tables and chairs exist, recognizing that is contentious in some circles.

Consider the world from the perspective of an omniscient outside observer. No such observers are imagined to actually exist, but thinking about what they would say if they did provides a convenient stepping-stone to the more realistic perspective of observers inside the world.

Our external observer notices that the world is not structureless; there are patterns relating different parts of it. The world can be fully described in a more efficient way than simply listing every part of it and that part's relation to other parts. The state of the world is algorithmically compressible; given the patterns within it, facts about certain parts of the world imply facts about other parts. These patterns might be called the "laws of physics."

We'll be a little vague about what the laws of physics actually are, because physics isn't done yet and there are presumably laws yet to be discovered. But we don't need to restrict the label to some aspirational theory of everything that applies at the most fundamental level and in all circumstances; there are patterns that emerge at higher levels that still deserve being called laws of physics (Dennett 1991; Carroll 2016).

One such pattern might usefully be labeled "evolution through time." The world can be divided into "moments of time," which have a relationship of being either nearby or far away. Nearby moments are related to each other in somewhat predictable ways; indeed, the persistence of certain kinds of collective structures over time is what inspires us to think of those structures as "existing" in their own right. There is nothing necessarily fundamental or absolute about time or evolution; at this level, it is simply a useful, perhaps approximate and limited, way of relating some parts of the universe to other parts.

In this picture, there is nothing *extra* that exists. There are no extra-physical beings or essences. The laws of physics, in particular, are not existing things that bring about the world or govern its behavior; they are simply convenient summaries of what exists. We are therefore being Humean about the laws of nature, in contrast with an anti-Humean perspective that would grant the laws both existence and powers (Maudlin 2007). The world as we have described it is something like David Lewis's "Humean mosaic" of events (Lewis 1986), although Lewis granted primacy to spatiotemporal relations in a way that seems unnecessary (and likely just wrong, at a fundamental level, if spacetime itself is emergent).

There may be any number of useful ways of *describing* or talking about the world, whether from an external perspective or from within the world. Some of these ways might involve mathematics, even though mathematics does not exist in the same way the physical world does. The reality-realist view considers mathematical statements in a way that a Humean considers laws of nature: as compact summaries of things that happen in the world. We might then want to extend our discussion of such statements beyond their world-based origins, but we should be cautious about overinterpreting the results of such exercises.

To see how this picture might work, let's abandon the external perspective and ask what the world seems like from the inside. There can be subsets of the world that are usefully identified as "agents" – persistent patterns over time that interact with the rest of the world in such a way as to construct within themselves (partial, imperfect, fallible) models of the world, models that include the agent itself. That is, there is some substructure of the agent that represents patterns that an external observer would recognize as characterizing the world. Such an agent can "perceive" and "think," by which we indicate processes within the agent that gather information (leading to correlations between part of the agent's internal state and the state of the external world) and process it (physically manipulate information-bearing parts of itself so as to generate predictions about itself and the world). All such operations are imagined to be purely physical, compatible with the patterns we recognize as the laws of physics.

An agent might, for example, draw the distinction between having an object (let's say an apple) in their possession, and not having any such object. Then a further distinction could be drawn between the state of having one apple, and the state obtained by acquiring another apple when already in the possession of an apple. The agent could recognize similarities between these situations and analogous situations involving oranges instead of apples. This might inspire them to label each situation with a "number," where possessing the same number of apples and oranges implies that the groups of fruits can be put into correspondence with each other. These physical occurrences could inspire our agent to abstract away from the notion of collections of things to an abstract notion of quantity. And the numbers, measures of how many objects one possesses, could be found to obey patterns of their own: every time an agent possessed two things, and was given two more of the same kind of things, they now possess four of those things, no matter what kinds of things are involved. That kind of awkward construction can be simplified to "2+2=4." Then the agent (or many agents, working over time and sharing knowledge) could invent a set of axioms and rules that would generate reliable statements about numbers, as well as about other abstract mathematical objects.

All of this new mathematical knowledge, however, starts out as no more or less than a convenient way of talking about physical things happening in the world. The abstract proposition that 2+2=4 is not a reflection of an independently existing truth. It is a bit of formalism that happens to find use as a way of indicating certain features of physical reality.

Similar considerations apply to morality. Those agents that tend to survive and pass on their heritage (biological and cultural) to subsequent generations will tend to, in some sense, *care* about the world and what happens in it. (Caring about what happens confers greater survival probability than complete indifference; it leads one to expend energy to run away from hungry tigers, for example.) Such caring manifests itself in correlations between internal states of the agents and events in the outside world: states that we can recognize as "approval" and "disapproval" as well as related feelings. Intelligent agents will naturally work to systematize these feelings, assigning properties of "right" or "wrong" to different actions that occur or might occur. All of this is nearly inevitable given the basic facts of biology and evolution as we find them in our own world, without any need to assume any independent reality for such moral judgments. The burden would be on those who believe in such independent reality to explain why those moral facts often align so conveniently with the personal judgments we would naturally expect agents to develop over time.[4]

3. *How Mathematics Fits In*

Let us relate this reality-realist picture to standard ideas about mathematical realism. In the picture sketched here, "2+2=4" is not meant as a true statement about objectively-existing objects "2," "+,", "=," and "4"; rather, it is a summary of multiple facts about the

[4] I am not explicitly commenting on the "evolutionary challenge" to moral or mathematical realism, but I agree with Clarke-Doane (2012) that the two cases are highly analogous.

physical world, along the lines of "every time I have two apples and someone gives me two more apples, I wind up with four apples." The mathematical realist, by contrast, tends to think of this statement as an objective truth about things called "addition" and "the integers," which exist independently of the physical stuff from which the universe is made.

In *Morality & Mathematics*, Clarke-Doane provides a list of features that a field should have in order that we should be realist about it. These include Aptness (typical sentences are true or false), Belief (sentences conventionally express beliefs), Truth (some atomic sentences are true), Independence (truths are independent of human minds), and Face-Value (sentences should be interpreted at face value). Interpreting "2+2=4" and analogous statements as summaries of physical facts would seem to be straightforwardly compatible with Aptness, Belief, and Truth. (It is true that two apples, augmented by two more apples, leaves us with four apples.)

One could worry about Independence – the convenient summaries provided by mathematical statements are convenient for humans, after all. An omniscient and omnipotent observer wouldn't need recourse to any kind of summaries; they could just know and reason about all of the physical facts. But the facts that underwrite the summaries are still there, whether any agents notice them or not. An absence of human minds doesn't affect the combination properties of apples. So Independence appears to be satisfied.

Face-Value is another matter. This criterion requires that mathematical statements about certain kinds of things (e.g. "even numbers") actually refer to honest things of that form. But we are suggesting that there are no "things" corresponding to the category of "even numbers." There are things corresponding to the category of "apples," and we might consider a collection of apples including an even number of them, but the number itself is not a thing. This view therefore fails at being realist about mathematics, as expected.

In this picture, then, mathematics comes pretty close to satisfying the conditions for realism, but falls just a bit short. This is why it is reasonable to contemplate invoking different senses of "reality." Mathematics is incredibly useful in describing the world, and we think that conclusions derived mathematically can be extremely reliable, but it's not real stuff in the way the world is.

This, in turn, is related to the Benacerraf-Field problem (Benacerraf 1973, Field 1989), which highlights the lack of causal influence of mathematics on the physical world. From our perspective that makes perfect sense; descriptive tools don't necessarily affect the things they are describing. The world simply *is*, whether or not scientists and philosophers and mathematicians ever come along to talk about it in mathematical terms. At the same time, we do believe that mathematical statements express truths, and do so in a way that is independent of human minds. One can decide for oneself whether those properties qualify for a kind of "existence."

4. *Mathematics and Physics*

With that in mind, we can turn to some conventional defenses of mathematical realism. One, alluded to in the opening, is the "unreasonable effectiveness" of mathematics in physics, to use Wigner's (1960) phrase. Empirically, successful scientific theories are formulated in mathematical language, and the most precisely tested ones tend to be the most mathematical. To accept the success of something like the Standard Model of particle physics, which is quintessentially mathematical to its core, might naturally imply that math itself is real.

But as mentioned above, there is an important distinction between "the Standard Model is real" and "the Standard Model represents real things." The scientific realist must be committed to the reality of nature, not to any particular representation of it. The Standard Model is not reality, it is – as the name indicates – a model of it. It would be a mistake to attribute reality to any tools we might use to describe reality.

Nevertheless, theories of modern physics are so *very* mathematical that one is tempted to wonder whether they even make sense in the absence of mathematical realism. As noted in Section 3.5 of *Morality & Mathematics*, following Putnam (2012), when we talk about quantum states as superpositions of possible experimental outcomes (like the number of electrons in a box), that description is so extremely far from our everyday experience of stuff in the world that the mathematics seems indispensable.

Here I would blame our paltry human intuition, rather than perceiving a need for an entirely new ontological category. When we say "the state of the electrons in the box is in a superposition," what nature hears is simply "the contents of the box are in this particular quantum state." There is nothing metaphysically special about having definite numbers of electrons in the box, nor metaphysically suspect about superpositions thereof. Once again, we find it convenient to reach for mathematical descriptions when talking about such quantum states, but there is nothing intrinsically more "mathematical" about them than there is about more familiar classical configurations.

In this context, it is worth mentioning that the Standard Model, like other successful theories of modern physics, is not presumed to be completely fundamental, nor completely comprehensive (i.e. there are physical situations, like near the singularity inside a black hole, where the model isn't even supposed to apply). It is an "effective" theory, applicable only within a well-defined domain, and that domain doesn't include important phenomena such as dark matter and for that matter gravity (Carroll 2022). Most physicists expect the Standard Model to be eventually superseded by a more fundamental theory. That more fundamental theory is likely to also be formulated in mathematical terms, but the specific math being used could be utterly different. (Perhaps it will be a discrete theory, in which case real and complex analysis might be beside the point.) In that case, it's hard to know which mathematical concepts should be granted "reality." Is it only the most fundamental ones, or do concepts that are useful in emergent approximations qualify as well? How precise does an approximate theory have to be before it qualifies? If scientists happen not to stumble on a particular theory, is its mathematics rendered unreal?

Presumably the mathematical realist doesn't want the reality of certain parts of mathematics to depend on their usefulness within science. Rather, the idea should be that the usefulness of some mathematics provides evidence for the reality of mathematical ideas more generally. Indeed, there are examples where incompatible mathematical ideas find homes in different parts of science – Euclidean geometry is central to some theories, and non-Euclidean geometry to others. So it must somehow be the general idea of mathematical realism, rather than specific concepts, that gains credence from the success of quantitative empirical science.

But it is unclear how the usefulness of, say, matrices in quantum mechanics provides evidence for the reality of, say, homotopy groups. I would argue that what actually gains credence is the reality of the subject matter of quantitative empirical science – i.e., the world – rather than the tools we use to describe it. This view saves us from awkward decisions about which parts of mathematics are supposed to be real, and which parts of science support them.

5. *Mathematics Without Physics*

If one argument for mathematical realism leans on the importance of mathematics to scientific theorizing, another takes an orthogonal tack: we should take mathematical realism seriously because mathematical truth *doesn't* depend on physical facts. The statement "there is no largest even number" is true whether or not there are collections with an even number of apples, and it was true before any person invented the concept. In a sense this strategy tries to argue for realism solely from Independence, the fact that mathematical truths are independent of human minds.

This argument leans on the intuition that certain mathematical claims are simply *true*, whether or not they are useful or whether anyone has ever thought of them before. The truth-value of mathematical statements isn't located in time, or contingent on human imagination. "There is no largest even number" bears truth, and bore truth before anyone formulated the relevant concepts. Indeed the same goes for "there is no largest multiple of an integer N," even though there are specific values of N for which no human being (or alien) has ever formulated this statement. It is hard, according to this line of reasoning, to make sense of such facts absent a commitment to mathematical realism.

In response, we may distinguish between two different kinds of mathematical claims. One category is the set of claims that are (or were, or will be) relevant to describing some part of physical reality. The concept of an even number is useful when we are discussing collections of apples, and the concept of a section of a fiber bundle is useful for field theory. In such cases, mathematical statements can be translated into statements about physical reality. Propositions about even numbers stand in for facts about apples and other collectible objects. Therefore, these examples don't provide evidence for a separate reality for math itself; they merely reflect the reality of the underlying physical world.

But there are also mathematical statements that (at least for a time being) seem to be irrelevant to our descriptions of physical reality, and yet there is a strong temptation to think of them as "true." We don't need the concept of a largest even number, or the absence thereof, to actually do physics. But it seems to follow from the rules we have set up for dealing with integers more generally. Some mathematical truths can be thought of as *extrapolations* away from physical reality, but still based on the same principles that were suggested to us by considering that reality.

Should we consider such truths as evidence for a separate reality for mathematics? One consideration might be that if we didn't accept mathematical realism, we might expect there to be different, incompatible "truths" that are based on the same (or equivalent) axioms that we deploy in our discussions of physics.

But of course, there are such incompatible "truths," which appear in a variety of ways. One way is simply to consider sets of axiomatic systems that are completely equivalent in those cases when they are describing physical reality, but which might diverge when we extrapolate them farther. As a somewhat trivial example, we can consider the difference between ordinary addition and addition modulo some integer N. If N is sufficiently large, there are no physically realizable numbers (or numerical values for the quantity of a collection of objects) that we could ever add together to reach it. In that case, the two theories would be physically equivalent. There would be no Platonic truth of the matter concerning what answer one would get when adding N-1 to itself. One might object that “addition modulo *N*” isn’t what one meant by “addition,” but that is a choice you are choosing to make when defining terms, not something that is decided for us by reality.

Perhaps a more conventional example is the Continuum Hypothesis, which is famously undecidable in ZFC (Zermelo-Fraenkel set theory with the Axiom of Choice). The Continuum Hypothesis, which states that there is no set whose cardinality lies strictly between that of the integers and the real numbers, sounds like something a realist would like to be either true or false (as apparently Gödel did, for example). But it turns out that CH can neither be proved nor disproved within ZFC; one is free to add it or to add its negation. Unlike the mathematical realist, the reality realist has no trouble with this observation. The physical world either does or does not contain structures that can usefully be described by numbers with a cardinality in between that of the integers and the reals. We don’t know whether it does or not, but in neither case are we forced to decide about the pre-existing reality of the mathematical idea.

The case for the reality of mathematical propositions might seem most straightforward when those propositions are theorems that could be proven from axioms. That is the case for things like the non-existence of a largest even number, or analogous but less trivial statements. The process of theorem-proving is a purely logical operation, not relying on any features of the physical world. But which axioms we are tempted to label as “true” does depend on physical reality. The fact that the interior angles of triangles add up to 180 degrees may once have seemed like an immutable truth, but we now recognize that it depends on the axioms of Euclidean geometry, and fails to hold in other equally legitimate systems. We are then left with a form of “*if-then*ism” (Putnam 1967), in which true statements are of the form “if these axioms are accepted, this theorem can be proven,” but a typical mathematical realist wants the proposition in the theorem to be real, not merely the conditional statement.

It is therefore more common for realists to point to statements that cannot be proven from agreed-upon axioms. The Continuum Hypothesis is an example, as are various statements that cannot be proven or disproven according to Gödel’s incompleteness results, as are the “standard” models of something like Peano Arithmetic. In the last case, for example, there is a feeling that we know what the “true” model of the integers looks like, even though we have theorems guaranteeing the existence of non-standard models that satisfy the axioms perfectly well.

As Putnam (1980) points out, this isn’t really a problem for the hard-core mathematical realist, who posits an ability (somewhat ill-defined) by which we can “grasp” which model is the real one. But it is also, as he goes on to emphasize, not a problem for the hard-core anti-realist, who doesn't believe there is any one "correct" model of any particular axiomatic system. For purposes of this issue, reality realism is in the latter camp. There is, for sure, what the physical world does; some models might be useful in describing that, and some might not. But the set of all models of arithmetic or any other axiomatic system are not divided into the intrinsically "true" ones and the "false" ones. There is merely the question of which models are useful in talking about reality.

6. *Physics and Consistency*

A final argument for mathematical realism comes from the conviction that our mathematical theories are consistent, or at least there is a fact of the matter about the consistency of such theories. Such an argument is given in Sections 3.5 and 6.2 of *Morality & Mathematics*. The idea is roughly this (paraphrased):

> If we believe some theory *T*, we should also believe that *T* is consistent – that it does not lead to a contradiction. But to have opinions about what leads to what requires metalogic, which commits us to a theory at least as strong as the natural numbers. But we can't prove that such a theory, say Peano Arithmetic, is consistent. In fact it's much worse: if PA is consistent, so is PA + ~Con(PA), where "~Con(PA)" means PA is *inconsistent*. So if we don't want to attribute objective reality to one system or another, rather treating all axiom schemes as equal, then there is no objective fact of the matter as to whether PA, and therefore classical logic, is consistent. And that would be bad.

The brief response to this is that the question of whether a given formal theory is consistent is not really the important one. What matters, at least to scientists, is whether there is a model of that theory that accurately represents reality (or at least, part of the model accurately represents part of reality).[5]

To be clear, *reality* is consistent, essentially by construction. The real world can be thought of as a collection of "things" that "happen," where scare quotes remind us that we need to be a bit vague in the absence of the correct fundamental ontology. As mentioned earlier, quantum mechanics might fool us into supposing that reality is not really a set of things that actually happen, because e.g. the number of electrons in a box might be described by a superposition rather than a definite value. But that's just a consequence of

---

[5] As recently argued by Azzouni (2022), there is also the issue that theories become meaningful only when interpreted in terms of a model, and ~Con(PA) will have different meanings in a model of PA + Con(PA) as opposed to a model of PA + ~Con(PA).

using antiquated notions of "things" and "happen." There is a quantum state describing the box, full stop. Superpositions are perfectly legitimate things.

A formal theory, by contrast, generally posits some axioms and then derives some theorems. Sufficiently powerful formal theories will generally have multiple models associated with them. The thing that represents reality is an appropriate model, not the set of axioms. The existence of *inappropriate* models shouldn't bother us all that much.

In the case of PA, as explained in *Morality & Mathematics*, the worry about consistency is partly assuaged by how it actually works. The intuitive concern is that an inconsistent theory can prove a contradiction, and from a contradiction we can prove anything at all, so the theory is useless. On p. 82 we read:

> A model of PA + ~Con(PA) is a model in which there is an infinitely long ``proof'' of a contradiction from PA. I put ``proof'' in quotes, because a proof must be finite. The model is wrong about finiteness. Or that is what we would like to say.

Clarke-Doane goes on to argue that we can't say that "the model is wrong about finiteness" if we think that both PA + Con(PA) and PA + ~Con(PA) are equally valid sets of axioms, and the escape is to be a realist who thinks that one of those sets is true and the other is not. But there is another escape: not that the model is wrong about finiteness, but that it expresses a notion of finiteness which is not the one relevant to describing our real world. A model of PA + Con(PA) is more useful for that purpose.

One way to make the consistency worry seem more bothersome is to relate it to allowed physical processes. If we physically construct an automated theorem-proving machine that starts with seemingly reasonable axioms, will it ever prove a contradiction? Like most people, I'm happy to believe that it would not. But by making the worry more vivid, we've changed it into a question about physical reality. And physical reality is going to do what it's going to do, we just have to live with the consequences.

Although it is somewhat tangential to our concerns here, talking about reality and the laws of physics naturally leads one to ask why there are laws of physics at all. This is an especially sharp question for the Humean, who thinks that laws of physics are merely convenient summaries of sets of facts, rather than independently-existing concepts with some causal or governing powers. If the world is a collection of facts, why do those facts have so much structure and compressibility?

I have no idea. I bring it up only because it does seem like a legitimate concern for the Humean, or for reality realism more generally. It may be that this is a "why" question without a distinct answer other than the brute fact of the matter, but that seems unsatisfying. The universe never promised to satisfy us, but it's reasonable to ask whether it could before we entirely give up.

7. *Conclusions*

Thinking through these issues has caused me to reflect on the extent to which many working physicists are liable to resist thinking about what is "real." They care about what works, and what can be measured, but will actively avoid questions about what really exists. (This attitude has been expressed to me by multiple physicists in more or less just these words.)

This is a shame, and is a reflection of the unfortunate divergence between science and philosophy. Physicists have, no doubt, been extraordinarily successful at constructing models of the world that work and make accurate predictions, even without caring too much about the underlying reality. But I would suggest two shortcomings of this perspective.

First, reality is intrinsically interesting. I would wager that most physicists first became interested in science because they wanted to better understand the real world, not because they simply wanted to make successful predictions. The latter attitude is inculcated during their training as scientists. There may be an instrumental reason for this, focusing attention on practical/solvable problems, but something of the initial motivation is lost.

Second, reality is potentially useful. That is, even if one just wants to be a hard-headed model-building scientist, there are angles and insights that might only come from thinking hard about what is real. The fact that physicists don't agree on the fundamental ontology of quantum mechanics, nearly a century after the formulation of the theory, is a case in point. We don't know for sure, but this lack of agreement could be holding us back in the search for a theory of quantum spacetime and ultimate unification. It wouldn't be completely surprising if the nature of reality played an important role in the invention of such theories.

The project of *Morality & Mathematics* is therefore a crucially important one, for reasons in addition to the obvious importance of understanding the status of moral claims. We need to be able to separate what is real from what is not, and what precisely that means. Although I am an anti-realist about both morality and mathematics, I do appreciate the force of some of the arguments for realism. I look forward to changing my mind if a sufficiently convincing argument comes along.

Acknowledgements

I am enormously grateful to Justin Clarke-Doane for prompting this paper, nudging it into existence, and offering many extremely helpful conversations along the way. Thanks also to Jody Azzouni and Tim Maudlin for very useful discussions.

Departments of Physics and Philosophy, Johns Hopkins University,
and Santa Fe Institute
Baltimore, MD, USA
seancarroll@gmail.com

References

Azzouni, J. 2004. *Deflating Existential Consequence.* Oxford: Oxford University Press.

Azzouni, J. 2022. Clarke-Doane on Gödel, unpublished.
https://jodyazzouni.com/articles/

Benacerraf, P. 1973. Mathematical Truth. *Journal of Philosophy*. 70: 661–79.

Carroll, S.M. 2016. *The Big Picture: On the Origins of Life, Meaning, and the Universe Itself.* New York: Dutton.

Carroll, S. M. 2022. The Quantum Field Theory on Which the Everyday World Supervenes. In *Levels of Reality in Science and Philosophy* (pp. 27-46). New York: Springer, Cham.

Clarke-Doane, J. 2012. Morality and Mathematics: The Evolutionary Challenge. *Ethics* 122: 313-340.

Clarke-Doane, J. 2020. *Morality and Mathematics.* Oxford: Oxford University Press.

Dennett, D.C. 1991. Real Patterns. *The Journal of Philosophy* 88: 27-51.

Field, H. 1980. *Science Without Numbers: A Defence of Nominalism*. Princeton: Princeton University Press.

Field, H. 1989. *Realism, Mathematics, and Modality.* Oxford: Blackwell.

Ladyman, J.J. and D. Ross, with D. Spurrett, and J. Collier 2007. *Every Thing Must Go: Metaphysics Naturalised.* Oxford: Oxford University Press.

Lewis, D. 1986. *Philosophical papers* Vol. **2**. Oxford: Oxford University Press.

Maudlin, T. 2007. *The metaphysics within physics*. Oxford: Oxford University Press.

Putnam, H. 1967. Mathematics without foundations. *The Journal of Philosophy*. 64: 5-22.

Putnam, H. 1980. Models and reality 1. *The journal of symbolic logic*. 45: 464-482.

Putnam, H. 2012. Indispensability Arguments in the Philosophy of Mathematics, in H. Putnam, *Philosophy in an Age of Science: Physics, Mathematics and Skepticism.* Cambridge: Harvard University Press.

Szabó, L.E. 2021. Physicalism without the idols of mathematics [Preprint] URL: http://philsci-archive.pitt.edu/id/eprint/21383.

Wigner, E. P. 1960. The unreasonable effectiveness of mathematics in the natural sciences. Richard Courant lecture in mathematical sciences delivered at New York University, May 11, 1959. *Communications on Pure and Applied Mathematics*. 13: 1–14.