DEMYSTIFYING THE APPLICABILITY OF MATHEMATICS

NICOLAS FILLION

ABSTRACT. Essential tensions remain in our understanding of the reasons underlying the striking success achieved in science by applying mathematics. Wigner and many likeminded scientists and philosophers conclude that this success is a miracle, a "wonderful gift which we neither deserve nor understand." This essay seeks to dissipate that aura of mystery and bring the factors underlying the success of applied mathematics into the fold of scientific rationality.

AUTHOR BIO. Nic is an assistant professor of philosophy and member of the center for scientific computing at Simon Fraser University, Canada. His work focuses on philosophy of science, philosophy of mathematics, logic, and numerical analysis.

Inquiries into the nature of mathematics as a science of its own and into its role in empirical science have a venerable tradition. Given that mathematics displays a kind of exactness and necessity that appears to be in sharp contrast with the contingent character of worldly facts, the problem that is perhaps the most unsettling examines how mathematics can be used to adequately represent the world. For instance, Einstein argued that "[t]he laws of mathematics, as far as they refer to reality, are not certain, and as far as they are certain, do not refer to reality" [6]. Similarly, Russell maintained that "[t]he exactness of mathematics is an abstract logical exactness which is lost as soon as mathematical reasoning is applied to the actual world" [8]. And yet, since the scientific revolution, efforts devoted to writing the book of the world in the language of mathematics have been resoundingly successful.

In light of this tension, many scientists and philosophers maintain that the applicability of mathematics is condemned to remain intrinsically mysterious. For instance, Wigner famously claimed that the "miracle of the appropriateness of the language of mathematics for the formulation of the laws of physics is a wonderful gift which we neither understand nor deserve" [10]. Dirac has similarly claimed that "[t]here is no logical reason why [the method of mathematical reasoning to study natural phenomena] should be possible at all, but one has found in practice that it does work and meets with remarkable success" [5]. It is true that, due to resilient tensions in our

understanding, the applicability of mathematics is surrounded by an aura of mystery, but the present essay seeks to bring it back into the fold of scientific rationality.

1. A MOSAIC OF PROBLEMS

Failure to make significant progress towards solving a foundational problem often results from a clumsy understanding of that problem. In the case of the applicability of mathematics, it is also the case that part of the mystery stems from gathering many problems that require different types of solutions under the same umbrella. The striking achievement we wish to explain is the *success* of our use of mathematics in scientific practice. Yet many of the most widely discussed themes are only tenuously related to the explanandum.

Some such themes focus on mathematics qua language. There is bewilderment that it is even possible to use the language of mathematics to describe the world. In oder to see that no mystery lies here, we must regard the activity of mathematical modelling as any other modelling practice. Constructing a model always involves the choice of a medium for the representation. Yet regardless of whether the medium chosen is plastic, wooden sticks, a picture, or statements in some language, models will succeed in capturing some aspects of a system, while other aspects will be idealized away. Each medium has its strengths and weaknesses. The main advantage of mathematics qua language is its considerable *expressive power* and *versatility*. If we consider the generality of foundational approaches to mathematics such as set theory or category theory, it would be difficult to imagine possible states of affairs that could not be somehow describable in mathematical terms. Hence the possibility of using the language of mathematics to describe the world is not in itself very surprising. But more importantly, the expressive power of mathematics qua language should not be conflated with our explanandum, for many mathematical expressions do *not* successfully apply. So what needs explaining are the circumstances that make *some* of the mathematics apply so successfully.

Another such theme focuses on the unexpected applicability of mathematical concepts developed in epistemic contexts in which no conceivable applications were anticipated. Yet such questions do not seek the actual reasons that underlie the successes of various applications. Instead they require an account of the conditions of possibility of such successes. However, in order to account for these actual successes, we would simply assume by fiat (based on the recent history of science) that mathematics *is* applicable, and then seek the causes of successful applications.

Thus, the problem is not one of characterizing mathematics as a language. Rather, it is one of explaining how to compare the virtues of different mathematical representations in a way that accounts for the success of those with a comparative advantage. It is common to identify truth as the theoretical

virtue fulfilling this role. But even if we are willing to grant the broad point that the fundamental goal of science is the pursuit of truth, it cannot be denied that few of our successful models and theories are exactly true—indeed some are not even remotely true.

As I will argue in the next two sections, idealizations and other epistemic shortcuts play an indispensable role in the effective use of scientific rationality. For this reason, I take the main problem to be this: Given that the construction and manipulation of our successful mathematical models of reality is riddled with uncertainty, measurement error, modelling error, analytical approximations, computational approximations, and other forms of guesses and ignorance, how can their remarkable accuracy be explained? On the basis of the commonsensical "garbage in, garbage out" rule, this accuracy appears rather baffling, and accordingly I call this the problem of the *uncanny accuracy* of mathematics.

2. TOO TRUE TO BE GOOD

As I have pointed out, an explanation of the actual success of applied mathematics in scientific practice is unlikely to be grounded in the literal truth of models, since very few models have this property *de facto*. However, to understand the nature of the success we seek to explain with more precision, it is important to acknowledge that this failure to be exactly true is a *feature*, not a bug. Indeed, any good theory idealizes away aspects of a physical system. Truesdell [9] elegantly make the point that "[o]ne good theory extracts and exaggerates some facets of the truth. Another good theory may idealize other facets. A theory cannot duplicate nature, for if it did so in all respects, it would be isomorphic to nature itself and hence useless." Theories and models play such a prominent role in physics because untangling the world is beyond the reach of our unmediated reason. We do not build theories to duplicate this complexity, but to set it aside as much as possible.

To illustrate this point, consider three different kinds of "idealized bodies" (or, idealized "building blocks") employed in classical mechanics to represent physical systems: mass-points particles, perfectly rigid bodies, and perfectly continuously deformable bodies. Despite their fundamentally idealized character, each type gives rise to a specific approach to classical mechanics. Articulating different idealized perspectives that complement each other enables us to efficiently get a grasp on the inner workings of physical systems. On the other hand, insisting on a *sub specie aeternitatis* true apprehension would be a path toward certain failure. A model or theory that contained "the whole truth and nothing but the truth" would quite simply be *too true to be good*.

Even so, if idealizations are to lead to any success, not any distortion can be warranted. Models should be *true enough* in order to be good. Mathematical modelling is a question-driven endeavour, so that the success has a
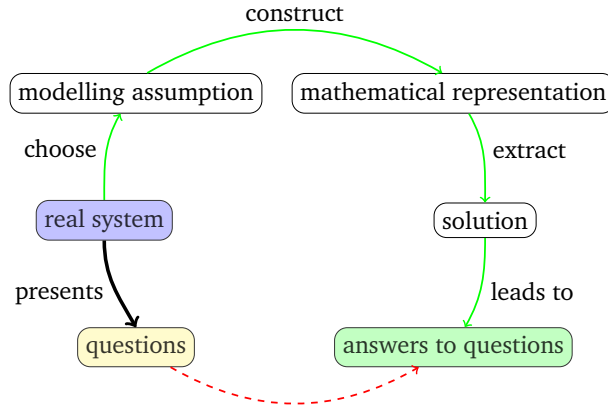
FIGURE 1. Constructing mathematical models to answer questions about the world

pragmatic dimension. In applied mathematical practice, one considers real systems, i.e., systems as we actually encounter them in the universe we live in. So, in contrast to abstract models, a real model is not populated with mass-points, rigid bodies, or continuous media, but rather with things like tennis balls, blocks of concrete, wood studs, steel beams, planets, galaxies, impure water, etc. In the presence of such real systems, we formulate specific questions that determine what aspect of the system is the behaviour of interest. Here are examples of such questions: Would this structure break under a typical load? Would a certain solution containing likely impurities remain stable under a certain increase of temperature? Can the observed trajectory of Uranus be explained by the presence of another heretofore unobserved planet? The task of mathematically modelling real systems is to derive a mathematical representation of the system that will allow us to correctly capture some of its physical properties. From this point of view, a good model does not have to correctly capture all aspects of the system, but only those relevant to the questions that concern us in the first place. Moreover, the question-driven character of applied mathematics makes it clear that representations do not have to be true in order to lead us to correct answers—*selective accuracy* is sufficient.

It is hard to give a completely general account of the way in which mathematical representations are constructed in order to answer our questions about the behaviour of interest, but figure 1 perhaps comes close. Starting from a raw, non-mathematical real system, we choose idealized building blocks as our representational medium and attempt to articulate what is mathematically relevant to the problem. The selection of modelling assumptions is a crucial step, which is often plagued with error and uncertainty. It is sometimes possible to say whether modelling assumptions are factual or not, but it is typically hard to directly assess such claims in a comparative way. In other words, it is hard to determine whether one assumption is as

far from the truth as another. Making such judgements is even more difficult for sets of modelling assumptions. Moreover, whether this error will invalidate our answers cannot be determined at this point. While we only have particular facts about a system, the evolution of that system is underdetermined. Without an underlying theory providing general kinematic principles, such as a geometrical structure and general conservation laws, no modelling equations could be derived to characterize the behaviour of the system. Similarly, a general theory does not have sufficient specificity to predict anything about the behaviour of systems without being supplemented with specific modelling assumptions [7]. It is this interconnected collection of hypotheses that faces the tribunal of experience.

Indeed, this collection of hypotheses determines dynamical equations characterizing the temporal behaviour of the system (i.e., equations describing the evolution of points or regions in a state space through time). The evolution rule is typically a differential equation (continuous time) or a difference equation (discrete time). Finding the trajectory in the state space prescribed by the rule amounts to solving the system, and it is a very crucial step in extracting the information needed for empirical tests. Without effective solution methods, there is no prediction nor explanation, only speculation. However, this step often involves significant analytical and computational challenges, and we return to this theme in next section. But presuming that information has been accurately extracted, we would then obtain answers to our questions and evaluate the successfulness of our model.

3. Success as a balancing act

The ineliminable need to set aside complexity puts us in a situation in which the sets of modelling assumptions from which we derived model equations are extremely simplified compared to what would faithfully capture real physical systems. Needless to say, when we build a model for a system of real bodies, the inaccuracy and incompleteness of the modelling assumptions could very well lead us to incorrectly answer questions about the behaviour of interest. To establish whether this is the case, a traditional view enjoins the modeller to compare the idealized model to a de-idealized model derived from an accurate and complete set of modelling assumptions. This would allegedly guarantee that the model equations thus derived would correctly answers our questions about the system. Batterman [1] pinned down the idea nicely: "The aim here is to effect a kind of convergence between model and reality. One tries, that is, to arrive at a completely accurate (or 'true') description of the phenomenon of interest. On this view, a model is better the more details of the real phenomenon it is actually able to represent mathematically." However, to experienced applied mathematicians, it is clear that the more details are built into the model, the more mathematically intractable the mathematical equations representing

FIGURE 2. The fundamental balancing act at the core of applied mathematics

the behaviour of interest are likely to be. That means that, even if we can somehow derive model equations from our accurate and complete set of modelling assumptions, it is likely that we will not be able to use them to make predictions and to obtain answers to our questions concerning the behaviour of interest. That would be a representation that is not manageable, no matter how true, accurate, or complete it is—it would be useless to us.

The improvement of the accuracy of the modelling assumptions brings about a decline in the tractability of the model. Thus, counter-balancing the view of the role of applied mathematics as a language for formulating true representations of systems, there is the view that mathematics is "the art of finding problems we can solve," as Hopf said (cited in [2]). Since in applied mathematics there is in addition a question of accuracy, there is always a cost-benefit analysis to perform. The most important contribution of mathematics to modelling is that it provides the tools to do just that. What makes mathematical modelling difficult is that above all we must find a balance between accuracy, completeness, and tractability, as in figure 2. Finding this balance with respect to the behaviour of interest is the *true measure of success* in applied mathematics.

But how can we know whether we have reached this balance? How do we distinguish accidental positive results from models that truly capture the essential features of the system? Perhaps this is the mysterious part that will resist our efforts to bring it into the fold of reason.

4. RATIONALIZING THE UNCANNY ACCURACY OF MATHEMATICS

We have seen that the complexity of the world is such that models will typically not be exact representations of physical systems. Moreover, even simplified models typically have a level of complexity such that extracting information from modelling equations will lead to an additional layer of error. Thus, the key to successful applications of mathematics is to establish that a description of the behaviour of a system is in fact approximately true,

without causing an overflow of information that would undermine our ability to assess the situation. In this respect, an essential virtue of mathematics is that it can be applied to itself in a way such that *finding whether a representation is close to the truth is easier than finding what the truth is*. I will explain this slogan below based on general insights from perturbation theory and numerical analysis.

The concepts of *sensitivity to* and *robustness under perturbations* play a crucial role in any perspective on error management. There are many rigorously defined concepts in applied mathematics which capture aspects of this very general idea (e.g., Lyapunov exponents, condition numbers, Lipschitz constant, etc.). But for the sake of this essay, an intuitive illustration will suffice. Most of us have at some point had to live in an apartment of questionable quality. Taking a shower in apartments of this ilk is not always without danger. Indeed, a very slight push on the shower knob—technically, we call this a perturbation of the knob's position—might lead to quite dramatic changes in water temperature. In such a case, we can say that the water temperature is very sensitive to perturbations. Robustness under perturbations is just the opposite. If you were so lucky to have air conditioning in this apartment, odds are that however much you cranked the knob, the ambient temperature would not change much. So, the ambient temperature was robust under perturbations. Great accuracy in the shower knob position would be required to correctly predict water temperature, but large errors in the AC unit's knob position could be tolerated in order to accurately predict ambient temperature.

The idealization, simplification, error, and uncertainty contained in models we construct to characterize some behaviour of interest can also be understood as perturbations. Let me first illustrate the point with a modern approach to understanding the impact of computational error as it occurs in computer simulations. Suppose that a dynamical system specified by an ordinary differential equation $x' = f(x)$ and an initial condition $x(0) = x_0$ has been derived as in figure 1 to represent a given physical system. Using some computer algorithm, we find a trajectory $\hat{x}(t)$ that will hopefully describe the behaviour of the system accurately. However, there is no *a priori* guarantee that it will. We need to first analyze the various sources of error. The applied mathematical toolbox offers many ways of talking about error. In what follows I will use the notion of *residual error* as it is easiest to interpret in physical contexts [4]. If we somehow knew the exact solution $x(t)$ to our dynamical system, we would find that $x' - f(x) = 0$ just by re-arranging the terms. However, since the simulated solution $\hat{x}(t)$ contains some degree of computational error, $\hat{x}' - f(\hat{x})$ would not be equal to 0. The quantity Δ given by $\Delta = \hat{x}' - f(\hat{x})$ is, what we call, the residual error. Now, we can reverse-engineer the situation. Instead of saying that $\hat{x}(t)$ is hopefully approximately solving the equation from the dynamical system, we can say that it is an exact solution to the dynamical system $x' = f(x) + \Delta$. With this change of perspective, we can now treat Δ as a perturbation of the original

dynamical system. It could be thought of as a breeze, a vibration, a small gravitational effect, or anything relevant. If the magnitude of the computational error, reinterpreted in physical terms, is smaller than the expected modelling error and uncertainty, then the computed solution is deemed *true to our modelling assumptions*.

In fact, for all we know, such a solution could exactly represent the physical system. As mentioned, the representation $x' = f(x)$ is not in general exact due to various sources of modelling and experimental error. However, as the expressive power of mathematics is virtually unlimited, one can presume that some other equation exactly represents the system, say $x' = f(x) + \varepsilon g(x)$, where ε is a small term and $\varepsilon g(x)$ acts as a correcting factor. We could once again verify the residual error of the computed solution $\hat{x}$ mentioned above, but this time with respect to the 'true' equation. Of course, we will not in general have an exact characterization of the correction factor and, as a result, will not exactly know the value of the residual error. However, we can study the amplitude of the residual using *qualitative methods* over various intervals of time, such as the limiting behaviour of the residual error as t goes to infinity. The intervals and parameters of choice for the error analysis will once again be determined by the behaviour of interest. Often, we will find that the residual error is vanishingly small. We then have a precise and effective method to rationalize the fact that even if the construction and manipulation of our successful mathematical models of reality is riddled with uncertainty, measurement error, modelling error, analytical approximations, computational approximations, and other forms of guesses and ignorance, they can be remarkably accurate.

5. Explaining miracles away

The uncanny accuracy of mathematics has been claimed to be miraculous in the sense that it does not seem to admit any rational explanation. To decide whether such a claim can be defended, it is necessary to have a conception of what might be received as a rational explanation. This, in turn, requires a correct understanding of the "logic" of model construction and model assessment. Precisely articulating such metatheoretical concepts is the province of epistemology. As a consequence, a solution to the problem of the applicability of mathematics will be of an epistemological nature, rather than of a metaphysical one.

In addition to the *de facto* presence of falsehoods, errors (intended and not intended), approximations, and uncertainty (including both known and unknown unknowns) in science, there are other elements that we have not yet mentioned. Indeed, cases of fortunate mistakes, aesthetic preferences, and personal idiosyncrasies of influential figures are also integral parts of real science. However, it does not follow that all those factors play an

equally important role in epistemology, as its point is to explain the reliability of scientific knowledge and to delimit its scope. As a result, epistemology does not take the actual thought processes of scientists as its objects, or the actual words used by scientists, or even what scientists take their own activity to be. Rather it envisages a better scenario in which the claims, hypotheses, models, theories, and methods are accounted for not by fortunate mistakes, idiosyncrasies, etc., but by a rationally compelling presentation they ought to have. To use the term introduced by Carnap [3], the object of scientific epistemology is a *rational reconstruction* of science.

The dimension of the rational reconstruction process that generates an object of study suitable for a properly epistemological analysis is often presented as an invective to distinguish the context of discovery from the context of justification. The distinction between the contexts is one between processes of discovery and methods of justifications. The phrase "methods of justification" denotes what satisfactorily establishes knowledge claims, independently of the beliefs of the historical actors. It is important to emphasize that which methods of justification are rationally admissible is not god-given, as there is room for disagreement. It is nonetheless clear that what is to be included in the context of justification is determined by what methods and tools are considered rational. Different choices might result in different organizations of what belongs to what context. The stakes are clear: if the methods of justification we are willing to admit are too restricted, then some essentially successful scientific practices will appear to be miraculous, without any rational grounds.

Scientists and philosophers alike often depict the scientific method as containing two essential methods of justification. On the one hand, there are the methods of probability theory and statistics which are meant to underly inductive inferences from observed phenomena. On the other hand, logic and axiomatics are meant to capture the deductive structure of scientific theories. Probability and statistics are essentially about making precise judgements about the likelihood of hypotheses. Deductive logic is essentially about truth-preserving inferences (i.e., inferences such that if their premises are true, so will their conclusion). These methods are undoubtedly rationally admissible when properly utilized, but it is essential to emphasize that they do not exhaust the field of rational justifications. It is therefore necessary to revise and supplement our "rational reconstruction toolbox," for otherwise significant parts of applied mathematical sciences would be wrongly considered methodologically unsound.

The successes of applied mathematics crucially depend on the methods of perturbation theory. The type of questions they address are not about probability, likelihood, or truth-preserving inferences. Instead, they concern questions of this type: if causal factors were slightly changed or if parameters were tweaked in various ways, what impact would it have? To see the contrast with deductive logic even more sharply, one could say that the methods of perturbation theory are essentially about determining the circumstances

in which arguments with false premises lead to accurate conclusions. Deductive logic cannot address such questions, for even if its inference forms preserve truth, they do not in general preserve approximate truth. Perturbation methods give us the resources we need to learn how to live with falsehood, and this is key to understanding the factors that make so many mathematical models and theories uncannily accurate.

To conclude, persistent failures to unravel the mystery of the applicability can be attributed to an insufficiently rich way of rationally reconstructing scientific and mathematical knowledge. To the extent that the problem of uncanny accuracy is concerned, we need to suitably enrich the catalogue of methods admissible for the rational reconstruction of the concepts of science and mathematics. We should not contemplate elaborate counterfactual constructions about pristine theories that contain no error and uncertainty, but learn how to live with them, and love them, for they are the conditions of possibility of successful science. Then, and only then, the allegedly miraculous character of the applicability of mathematics will be demystified.

References

1. Robert W. Batterman, *Asymptotics and the Role of Minimal Models*, British Journal for the Philosophy of Science **53** (2002), 21–38.
2. Armand Borel, *Mathematics: Art and science*, The Mathematical Intelligencer **5** (1983), no. 4, 9–17.
3. Rudolph Carnap, *The logical structure of the world*, University of California Press, Berkeley, 1928, Translation Rolf A. George, 1967.
4. Robert M. Corless and N. Fillion, *A graduate introduction to numerical methods, from the viewpoint of backward error analysis*, Springer, New York, 2013, 868pp.
5. Paul A. Dirac, *The relation between mathematics and physics*, The Collected Works of P.A.M. Dirac (1924–1948) (R.H. Dalitz, ed.), Cambridge University Press, 1939, 1995, pp. 122–129.
6. Albert Einstein, *Geometry and experience*, Sidelights on relativity, Dover Publications, 1923.
7. Nicolas Fillion and Robert M. Corless, *On the epistemological analysis of modeling and computational error in the mathematical sciences*, Synthese **191** (2014), 1451–1467.
8. Bertrand Russell, *The art of philosophizing, and other essays*, Philosophical Library, Inc., New York, 1968.
9. Clifford Truesdell, *Statistical Mechanics and Continuum Mechanics*, An Idiot's Fugitive Essays on Science, Springer-Verlag, New York, 1980, pp. 72–9.
10. Eugene Wigner, *The Unreasonable Effectiveness of Mathematics in the Natural Sciences*, Communications on Pure and Applied Mathematics **13** (1960), 1–14.

(Nicolas Fillion) DEPARTMENT OF PHILOSOPHY, SIMON FRASER UNIVERSITY, 4604 DIAMOND BUILDING, 8888 UNIVERSITY DRIVE, BURNABY BC, CANADA, V5A 1S6. WEBSITE: WWW.NFILLION.COM. EMAIL: NFILLION@SFU.CA.