A dialogue concerning undecidability, uncomputability and unpredictability

by Gabriele Carcassi
Assumptions of Physics (https://assumptionsofphysics.org)

Written for the FQXi
"Undecidability, Uncomputability, and Unpredictability Essay Contest" (2019-2020)

My attention focused on a discussion at the far end of the hall.

A: … they are different things. The only thing they have in common is the realization you can't do something. We have to look at the details.

In her presentation, stomping out vague ideas was a recurring theme… and she often made sense. So I approached the group to listen in.

B: Proceed! I am also curious about this and whether it is at all relevant to physics.

He seemed one of those experimentalists that has very little patience for mathematical gibberish or grandiose arguments. He was all about the data and what can be measured.

C: And let's also clear the terminology: what you call uncomputability is usually called undecidability, and what you call undecidability is usually called incompleteness.

As you noted before, he seemed to thrive on flowery mathematical language and more philosophical ideas. Maybe because he is a theorist, or maybe it's the other way around.

A: Right. Let's start with decidability: the ability to construct an algorithm for a particular problem. The issue here is termination in finite time. If you have ever used a computer, you know it can get stuck in an infinite loop and never end. What one can show is that some problems do not admit an algorithm that always terminates with the correct answer.

C: One does this by assigning a number to all possible Turing machines, and then proving it by diagonalization, which unfortunately is not terribly insightful.

B: An example would be good.

A: Indeed. Suppose you want to find out whether there exists a prime number that has a certain property. I don't know, that it is the sum of the square of two previous prime numbers. A simple approach would be: you take the first, check whether it has that property, if it has it you stop, otherwise you continue. If one of the prime numbers has that property, you will eventually find it and terminate. However, if none of the prime numbers has that property…

B: You do not terminate.

A: Exactly. You see, the statement ‘at least one of the prime numbers has that property’ is the logical OR of infinitely many statements like ‘2 has the property’, ‘3 has the property’ and so on. As long as you know one of the statements is true, then the OR is true, no matter the others. The negation of that, instead, ‘no prime number has that property’ is the logical AND of ‘2 does not have that property’, ‘3 does not have that property’ and so on. Now we have to test all the statements to make sure they are all true. If the statements are infinitely many, we will not terminate. So, if this is the only way you can test the property, you have a problem.

B: So, finite time termination imposes certain logic operations, correct?

C: Exactly. We have finite conjunction, logical AND of finitely many statements, countable disjunction, logical OR of infinitely many statements, but still countably many. And we do not have negation, because if the procedure terminates successfully when the statement is true, it does not mean it will terminate when the statement is false. This is called a Heyting algebra, which is different from the usual Boolean algebra. A locale is a type of…

B: But this is just about computation.

C: It’s not just about computation. In mathematics, for example, you may be able to prove that something exists, such as the well-ordering of the reals using the axiom of choice, but you will not be able to construct that ordering. The constructivists reject any mathematical entity that can only be proven to exist, but can’t be constructed or calculated. In fact…

B: I mean, it has nothing to do with physics… this is math or computer science.

C: Well, you may think the universe as a big computer…

B: It’s not.

The discussion here proceeded as you may have expected. This idea that the universe, or the human mind, is simply a computational device reminds me how, when the first principles of mechanics were discovered, the universe, or the human body, became simply a complicated series of levers. Maybe to a chemist the world is simply a series of chemical reactions, and to an economist is an intricate web of incentives. If you have a hammer…

A: If I may, there is something in science that follows the same rules: experimental verification.

B: That sounds a lot more pertinent. Please, proceed.

A: Suppose you want to verify the statement ‘the mass of the particle is exactly zero’. Can you do it?

B: No. Measurements have uncertainty. So, you can only verify it is less than some value.

A: Very good. But could you verify its negation, ‘the mass of the particle is not zero’?

B: You may be able to exclude the value, yes. If the uncertainty is far enough from it. We did it for neutrinos. Oh, I see you do not have negation… Ah, and measurements also need to take finite time. So, it will follow the same rules. I see, very nice.

C: You see, just like I was saying: the universe can be thought as a big computer...

A: No, no… I am not saying that at all! It's not that experimental verification is the same as computability. It's that computability can be used for experimental verification: if I verified the mass is about 1 Kg and that the volume is 1 liter, I have also verified that the density is 1 Kg/liter, because it can be univocally computed from the other measurements. Same thing for provability: if I experimentally verify the antecedents of a theorem, then I have effectively verified the conclusion. That is where the connection comes from.

B: That is very interesting. But, in practice, what does this give me?

A: At a theoretical level, it severely constrains the formal structures that can be used if one requires the elements of the theory to be experimentally distinguishable. Topologies, for example, serve exactly that role. Pardon my math, but an open interval is in the topology because it corresponds to experimental verification within a finite range. An isolated value is not in the topology because you can't verify the value exactly. So all these structures that mathematicians came up with have actual meaning in physics. Have you ever wondered why functions have to be well behaved in physics? As in, continuous with maybe a few points of discontinuity?

B: Because those are the ones we find useful.

A: Right. But why?

Shrugs. Silence.

A: I'll tell you why. Topologically continuous functions, which provide a more precise characterization of well behaved in this case, can be seen as those that map verifiable statements to verifiable statements. So, if you see the height of the mercury column is between 22 and 23 mm, you also infer that the temperature is between 22 and 23 C. An experimental relationship is that: you infer a measurement from another measurement. So a function that represents an experimental relationship must be well behaved. You don't have a choice.

B: This is admittedly nice. Can you do something more?

A: It is possible to build a theory on just these concepts. Finite time experimental verification gives you already a lot. Suppose you have a set of possible cases, like the states of a system, that must be experimentally distinguishable. You can imagine having a series of verifiable statements to do that. Because of finite time verification, in the limit, the most you can distinguish is a sequence of true and false, of ones and zeros. But a sequence of ones and zeros can be seen as a binary expansion of a number between zero and one. Therefore, you will never be able to experimentally distinguish between more cases than there are on the real line.

B: Again, very nice. But so what?

C: You are ruling out objects. For example, in quantum theory you are ruling out non-separable Hilbert spaces. Any physical theory that uses those, if you believe what she says, is ruled out. Any theorem or property of those spaces would be useless in physics.

B: I understand now. You can separate what is physically relevant from mathematical junk.

C: That's a rather harsh term… Time and time again, what was seen as mathematical junk turned out to be useful to physics later.

B: And a lot of it didn't.

As you and I have often discussed, there is this dichotomy for which physicists are stereotypically, sometimes proudly, sloppy with their math, and on the other side some like to explore very abstract notions (like quasi-projective hypercovering of Zariski varieties, whatever that is) that have very tenuous links to the real world. We badly need definitions and tools that are strictly rigorous both mathematically and conceptually; where each mathematical symbol corresponds precisely to a physical concept and vice-versa. I can't see how we can really make progress at the fundamental level otherwise.

A: Going back to the discussion, undecidability in the form of experimental unverifiability is not a problem: it's a feature. We need it to characterize what can be subject to scientific investigation. It should be embraced as one of the cornerstones in the foundations of science.

C: Fair enough. Since the next session will start soon, I think we should move on to unpredictability, and how it limits our knowledge. But again, we need to be precise: the type of unpredictability coming from chaos theory is different from the one in quantum theory. The first is simply lack of proper knowledge of the initial conditions. Even in chaotic systems, if you knew everything precisely, you would know the precise evolution. The problem is that every finite precision description, however accurate, will eventually worsen in time until, in the end, very little can be said, as the possible states of the system will be scrambled more and more.

B: Shuffling a deck of cards, rolling dice: systems whose outcome bear little correlation with their initial setup are nothing new. It's nice we can study these ideas in more detail, but I don't see it having radically new implications for how we should approach science.

C: Fair enough. Still, quantum mechanics is radically different. We cannot think of the properties of the particles as always having a well-defined value, so we cannot understand the uncertainty as simply coming from not knowing the actual values. In fact, the uncertainty principle is too often mischaracterized as something related to measurement. It is not. It's not about the ability to measure the value, it's the inability to prepare a state in the first place.

He is absolutely right. Many people that work in the foundations of quantum mechanics are really frustrated by this because, even in Heisenberg's original paper, there is a comment to the effect that Bohr disagreed about the interpretation of the result. The name in English, "uncertainty principle", is also very misleading. The name in Italian, for example, is "indetermination principle" which is a bit better.

A: What you say is correct. However, there are classical cases where this happens as well. If we think of thermodynamics variables, such as temperature and pressure, these too are not always well-defined. They are defined only for states at equilibrium. Take the number of particles: if we have a permeable membrane, and particles are allowed to come in and out, the number fluctuates and there is no unique well-defined value for a macrostate. There is an average, though. If we wanted a well-defined value, we would have to change the wall to an impermeable one, which would make the system transition to a new equilibrium. A precise number of particles is not well-defined for all macrostates.

C: I see your point, but this is not a perfect analogy. What would be the conjugate quantity of particle numbers for an uncertainty relationship?

A: I am not saying this is a perfect analogy. I am saying that we already have situations in physics where quantities are not always well-defined, yet the uncertainty we associate with them is still caused by not knowing the exact state and dynamics. So, quantum mechanics does not have to be radically different. In fact, any statistical process can be conceptually turned into a deterministic one where the initial conditions are unknown. Suppose the initial state is not enough to determine some future observable, for example the orientation of spin will not be enough to predict the new orientation after a measurement. For each direction of measurement, you will have a probability distribution. Now, imagine extending the definition of initial state with all the outcomes the spin will have for each possible angle. This is like saying that there are other degrees of freedom that, if included in the description, would render the outcomes determined.

C: This implies that the spin of a particle is not simply one direction, but part of something much more complicated. You are not arguing about a single hidden variable, but many, many hidden variables with their own independent dynamics. Essentially, you are arguing that properties of quantum systems would be emergent properties of something more complex, then. But how would you propose to describe the inside of a fundamental particle?

A: I do not. What I am proposing is to shift the perspective. Let's try to answer this: what makes a system classical and what makes a system quantum?

C: Well, the size. Classical systems are macroscopic.

B: That does not work. We can create entangled systems spanning many kilometers. We can create Bose-Einstein condensates, which are quantum states, with many particles. This idea of a quantum system being small simply does not work. Not all macroscopic systems are classical.

A: Indeed. I think the central theme is reducibility: the ability to reduce the state of the whole to the state of the parts. Say we have a ball; we throw it and study its motion. Or we can take a red marker, make a dot on the ball and study the motion of the dot. We say the ball is reducible meaning that studying the motion of the ball is equivalent to studying the motion of all the possible red dots.

C: And this would be a classical system. All the possible red dots would be points in phase space, and the evolutions would be classical Hamiltonian evolution.

A: Indeed. Now suppose we have an electron. The only way we have to study an electron is to make it interact with a photon. But when the photon interacts with the electron, it interacts with the whole electron. We can't take a red marker and make a dot on an electron. We cannot interact with half an electron.

B: So we can't know what happens at a finer level. Interesting. The inability to interact with parts makes it a quantum system. This makes sense. A proton, for example, in many situations is treated as one particle: it diffracts and interferes. But with deep inelastic scattering we probe the constituents and the quark-gluon structure emerges. And then, we can't treat it as one.

C: So, if that's the case, the uncertainty would simply be coming from lack of knowledge of this internal dynamics. How and why the different parts of the quantum system are interacting with each other.

A: I think so. I may be wrong, but it seems to me that all the causes of decoherence, the transition from superpositions to classical probability distributions, can be seen as coming from there. We have two entangled photons and measure the spin of one separately: we force an interaction with a part. We have a muon which, at some point, decays into an electron and two neutrinos. Confirmation of the decay is interacting with one part, typically the electron. We have a system that emits a photon, a fraction of the energy abandoning the system, which can then be absorbed separately.

C: I still have trouble thinking of parts of a fundamental particle. Also, the fundamental object is the field, not the particle.

A: Well, how do you construct the state space of the field?

C: Mathematically, it's the Fock space: you have the vacuum, the state space for one particle, the one for two particles, and so on…

A: So, why can't you think of a field as just a region of space with varying number of particles?

C: Well, for one, you can have a superposition of those states, so you have a state where the particle number is not well-defined… I see, you are going to say: the number is fluctuating because of all this internal dynamics I cannot describe, and the only things that are well-defined are the averages, etc… Ok. I guess what bothers me the most is the idea that quantum theory is not complete, because it does not tell us what happens at a more fundamental level.

A: No, no, no! That's the gigantic misunderstanding: quantum theory _is_ the fundamental level. It tells us as much as can be told.

C: I don't follow. You are saying there is a level of description that is missing. It's incomplete.

B: But if that level of description is outside the realm of experimental verification, then adding it to the theory would mean adding something not physical. It's philosophy. So, it is complete in the sense that it tells us everything we can access experimentally with a particular set of processes at the best level of description accessible. If we had other processes at our disposal, then we could get to a lower level. If we can do deep inelastic scattering with an electron, we would say it was not fundamental. But so far, we can't.

A: Right, a fundamental system is one that cannot be reduced further. It does not mean it does not have parts or internal dynamics. It means that all the processes at our disposal are not sensitive to that dynamics. They respond to averages and aggregates. What we have to realize is that this is inescapable: given a set of processes, we will always have a limit to what can be done with them. They are going to be sensitive only to a certain level of scale. The dynamics below that scale is unknown by definition. It is physically undefined.

B: What you say it that the foundations of physics should characterize the logic of experimentally verifiable statements and the level of description that these statements can provide. It's all about what we can test experimentally. What we can measure.

A: Right, we should start from there. These, I think, are the two main lessons from computer science's undecidability and from the uncertainty that comes from quantum mechanics and, to a lesser degree, statistical mechanics.

C: I don't find this satisfactory, though. You are basically saying we will not be able to describe the universe at the most fundamental level. We cannot say what these internal processes are doing.

B: All that you can measure _is_ the most fundamental level. There is nothing else. All the rest is philosophical nonsense!

C: The fact that you are not interested does not make it nonsense.

B: Can you measure it? If you can't measure it, it's not physics!

This level of acrimoniousness is not surprising since there are two competing interests. On one side, is the desire to understand how the new advances in the sciences influence the debate on the more open-ended questions about the nature of the universe and our role within; the desire to settle the problem of knowledge once and for all… and this sometimes brings in issues that are not strictly scientific. However, since these philosophical questions are unlikely to be settled, on the other side there is the desire to keep scientific discussion away from these issues; the desire to keep science focused on matters that can be solved and away from philosophical and religious considerations. This sometimes creates the perception that any such aims are a waste of time.

Unfortunately, any system of scientific thought cannot avoid taking its own stance on the relationship between the universe, us and what we can tell about it. If you make a measurement, you are already making an implicit statement of what can or cannot be inferred from it. As you know, between the two extremes, there is a middle ground: humbly characterizing just what can be explored through scientific investigation, and recognizing that everything else, as interesting as it may be, is outside the scope. And do it not to answer philosophical questions, but to do science: develop a formal mathematical theory that can be used in practice to guide the development of new physical theories or a better understanding of the old ones. No matter what scientific theory, there will be verifiable statements with a given logic; there will be a level of description provided by those statements. These will give us the topological, geometrical and algebraic structures we need and the fundamental assumptions taken for their validity.

A: Sorry to interrupt… We haven't yet talked about Gödel's incompleteness and its ramifications. Well, they are very little of course…

B: Yes. Let's change subject. Can you give me an introduction? I never really studied that.

C: There are actually two incompleteness theorems… But let's make it simple: in mathematics there was the desire to create a formal system in which all statements can be proved to be true or false. That would make the system complete. Yet, Gödel proved that, as soon as you have a framework complex enough to describe the natural numbers, you can create statements that cannot be proved either true or false. The proof essentially works by creating a statement that refers to itself in a way that it states its own unprovability. A little bit like "this statement is false" cannot be either true or false.

B: Isn't that just a mathematical trick?

C: Well, yes. But it does tell us that one such statement exists. In practice, there are more interesting ones. For example, "is there a set that has greater cardinality than the integers but smaller than the continuum?" is neither provable nor refutable within standard set theory. A mathematical technique called "forcing" was invented to create a system where such a statement is true. The axiom of choice was also something debated for a long time. So, mathematically, it is very much a problem since it essentially shows the limit of any formal system. But you seem to say this has no ramification for science? I find it surprising since science uses formal systems to describe nature.

A: Well, suppose we have a formal system where F = ma. What does that state in terms of physics?

B: That if I measure the force, to a given level of precision, the product of the mass and the acceleration must be within that level of precision.

A: Did I say that F was a force?

B: No.

A: What if I told you that F is the voltage measured across a conductor, m is the resistance…

B: and a is the current. You got me. I thought it would be about precision. You are arguing you have to define the quantities first.

A: Right. Newton's second law and Ohm's law have, formally, the same structure. The math itself is not enough to characterize the physics. So, it is already incomplete. How to relate the mathematical symbols to experimental verification, that is never going to be in the math. I claim that most of the physics is outside the formal system. If we sent back in time the equations of the standard model and math textbooks, without the schematics of how to build all the equipment needed to build particle accelerators…

B: Agreed! There is no science without experimental techniques!

A: Moreover, if we say that in science the truth is found experimentally, this already implies that our formal system will never be able to prove everything. If it could, there would be no experiment to conduct. We would already know the answer. It would not be science. Defining a scientific question means, in effect, creating a formal system where at least one statement is unprovable. We cannot prove mathematically that electrons exist.

C: Fair enough. What you are saying is that completeness is not as important in science as it is in math. But how exactly? What is contained in the formal system?

A: An intricate web of relationships between different verifiable statements: if "this quantity is between these values", then "this other quantity is between these values"; if "two particles have positive charge" then "there is a repelling force acting on them". These relationships can come from different places. They may be physical assumptions about the system, they may be physical constraints or simply logical constraints. Whatever their origin, they must form a logically sound and well-defined system.

C: Well, but can we create a model for all of the universe, where all relationships between all quantities are well defined? Or we can always find a statement where the relationship between two quantities is not well defined? That would seem a problem to me!

A: Right, but that's not how it works. The idea that we have a single model for everything is not realistic and not even desirable. Suppose you do not know the mass of the electron and want to measure it. Then your starting point is that you do not know the mass, and that there are many possible values in a very broad range.

B: If you do not even have a rough idea, the value may be outside of the range of your experimental setup…

A: Now suppose that you do particle identification in a detector. In that case, you will use the known mass of the electron to decide whether the particle coming in is possibly an electron or not. You see: what is measured and assumed as given changes in different contexts. The logical relationships change depending on what we want to measure and what are the premises.

B: This makes sense to me. I cannot experimentally confirm whether the doppler effect is correct by using measurement equipment that exploits the doppler effect.

C: Ok, what is taken as a premise may change, and the logical relationships I will have in different circumstances will have to reflect that. I still do not understand what this has to do with incompleteness.

A: Because the type of statements one has to use to create incompleteness are meta-statements: statements about statements. These are not things like “does this particular set have more than countably many elements but fewer than the continuum?”. It’s things like “can we ever have a set that has more than countably many elements but fewer than the continuum?” You see, you are going one level up in generalization. It should be quite evident that asking whether “the mass of this particular electron is between these two values” cannot be approached in the same way as asking whether “the mass of all electrons is between these two values”, which to some may simply be a matter of definition. Furthermore, it is yet another matter to ask whether “the mass of all electrons could have been different from our observed value”, whether it is a fixed value like pi or some measured parameter like the mass of the sun. The point is that once you start going up this hierarchy, you are already radically changing the problem. You have to rethink your starting point, how you are going to perform the experimental tests. The type of incompleteness you get from Gödel’s theorems is the least of your problems.

C: I don’t find this satisfying, though. All your points have been about what can be done through experimental testing. You are not telling me whether a theory can, in principle, be constructed.

B: What is the point of constructing the theory if it can’t be tested experimentally?

C: Well, maybe you can test it with other criteria. For example, using internal consistency.

B: But then it’s not science!

They seemed to enjoy going at each other in that manner so I just left to prepare for my presentation.

Probably the biggest misunderstanding about unpredictability, undecidability and incompleteness is that they are not stating something about the nature of the universe. This is like the fisherman who, having pulled his net, concludes there is a minimum size for the fish in the sea. Instead they tell us what we, as finite beings, can tell about the rest of the world through experimentation or logical deduction. These bounds are for us, they are our limitations, and we have to accept them. Either excluding that there may be other things out there, or insisting that we may still be able to study them nonetheless are two sides of the same coin: not accepting the limitations intrinsic to our human condition. That there are more things in heaven and earth than are dreamt in our philosophy.

And you probably know, at this point, what I am going to say: the surprising aspect should not be that we can’t know and prove everything about everything. Having been part of humankind for as long as I can remember, the astonishing fact is that we can actually say something about anything!