Why Does the Universe Exist? Some Perspectives from Our Physics Project

April 28, 2021

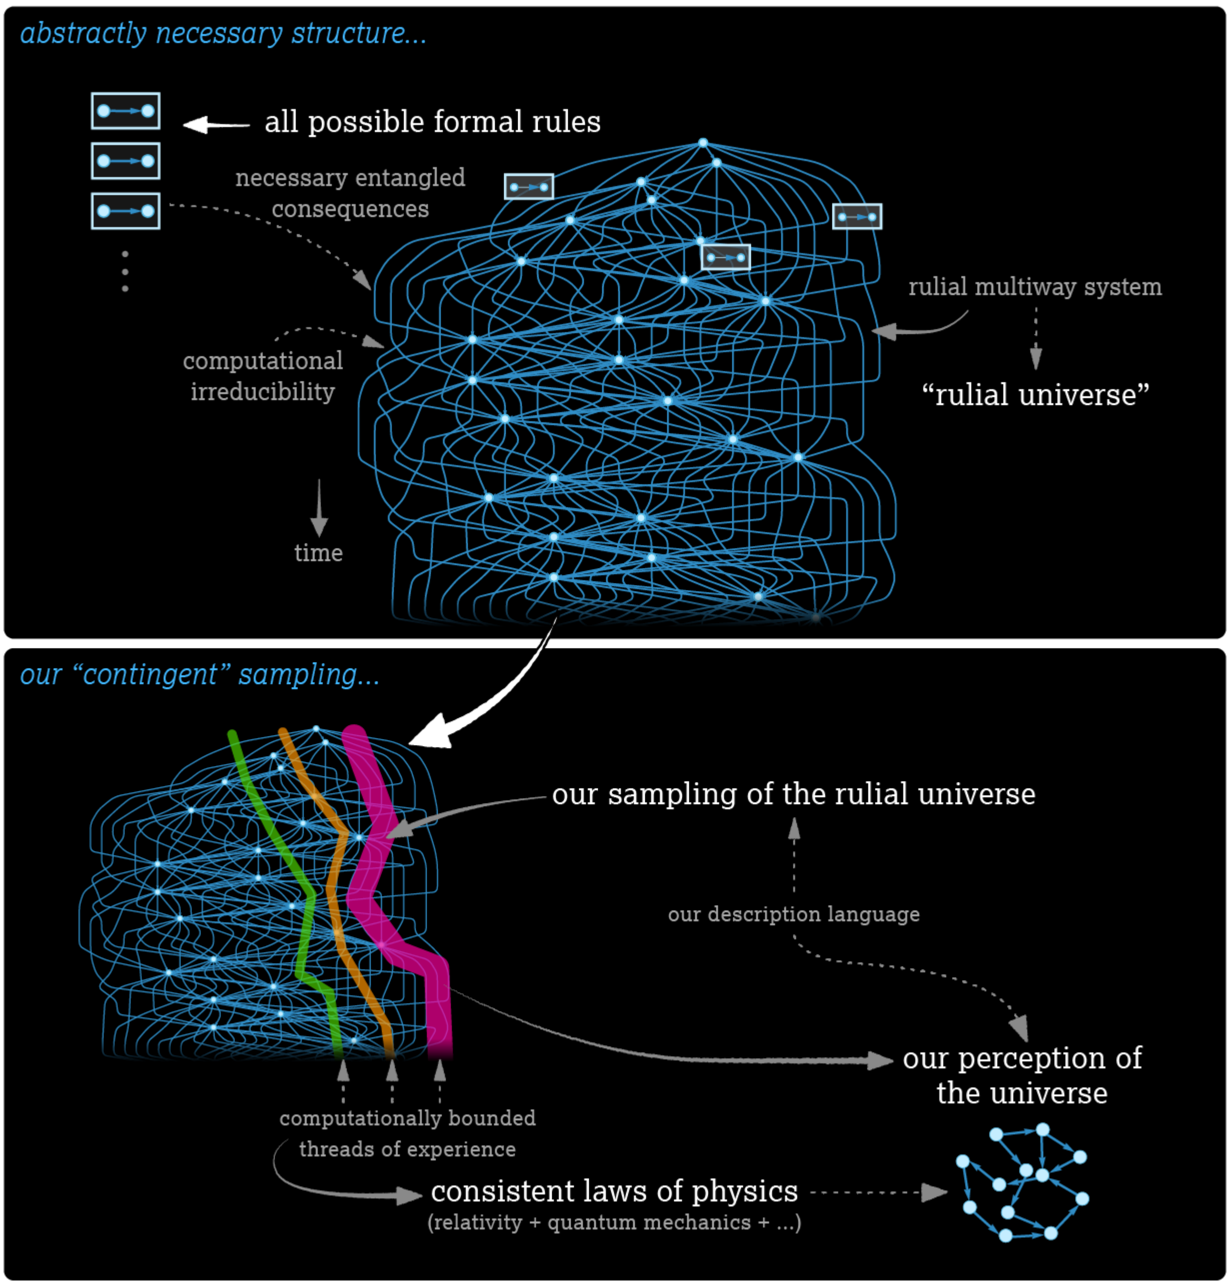

Why does the universe exist? Why is there something rather than nothing? These are old and fundamental questions that one might think would be firmly outside the realm of science. But to my surprise I've recently realized that our Physics Project may shed light on them, and perhaps even show us the way to answers.

We can view the ultimate goal of our Physics Project as being to find an abstract representation of what our universe does. But even if we find such a representation, the question still remains of why that representation is actualized: why what it represents is "actually happening", with the actual stuff our universe is "made of".

It's one thing to say that we have a rule or program that can reproduce a representation of what our universe is doing. But it seems very different to say that the rule or program is "actually being run" and is "actually generating" the "physical reality" of our universe.

As soon as one starts talking about "running programs" some people will immediately ask "On what computer?" But a key intellectual point is that computational processes can ultimately be defined completely abstractly, without reference to anything like a physical computer.

Consider for example one of my favorite cellular automata. Here's its rule:

And here's what it does in a particular case:

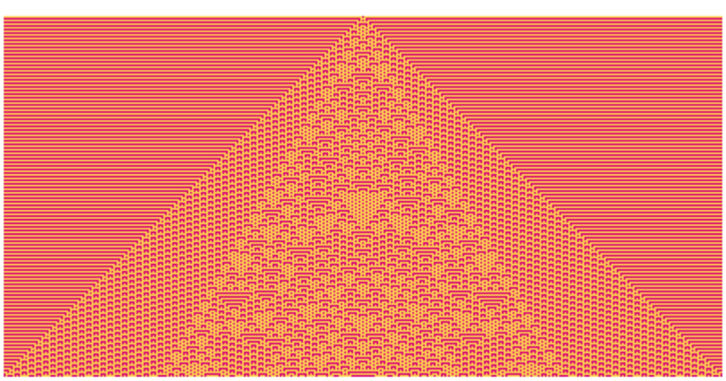

And, yes, I made this picture by running the rule on my actual, physical computer. But we can think of the rule as just giving an abstract definition of a computation to do. It's like the abstract computation 2 + 2 → 4. It's something that necessarily works the way it does, as a consequence of the abstract definitions that specify it. It doesn't depend on having actual

2 + 2 → 4 depends on having actual counters or a person arranging them.

OK, but so we can potentially imagine a purely abstract computational rule that can abstractly reproduce a representation of what the universe does. And since we are part of the universe, the rule must also reproduce whatever processes are involved in our perception of what goes on in the universe.

But the question still remains: why is anything "actually happening"? We can consider a rule. But based on what we've said so far, it's just an abstract possibility—something we can choose to define. There is nothing that says there's anything "real" or "actual" about it, any more than there is about the abstract mathematical statement 2 + 2 = 4.

Is the Universe Inevitable for Us?

There are an infinite number of possible programs that one can abstractly define. But we might assume that when it comes to representing our universe there's just a particular one that gets picked out. In other words, that in the computational universe of all possible programs, there's a specifically selected program that our physical universe follows.

After all, we might argue, we perceive definite physical laws in our universe. And surely, we might assume, there's nothing necessary about these laws; there could perfectly well be a universe with different laws. So there must be something specific about the rule that's being used by our universe, to give its particular laws.

Of course, there's something rather unsatisfying about this. Because it says that in the end there's something about our universe that is arbitrary, and that in a sense has to be "explained from outside". I've long suspected that this isn't quite right, and that once one thinks about observers within a universe observing that universe, the arbitrariness will somehow "cancel out". And, as we'll see, in our new models of fundamental physics something a bit like this happens, though the detailed story is considerably more subtle.

It all has to do with the interplay between computationally bounded observers and the whole computationally irreducible evolution of the universe, that effectively defines the passage of time. The basic point is that while there are all sorts of details of the underlying evolution that depend on the particular rule used, a computationally bounded observer can independent of the rule.

There's a familiar (and closely related) analog of this associated with the Second Law of Thermodynamics. A gas, for example, involves lots of molecules bouncing around. The details of their trajectories will depend on the particulars of their interactions—and will inevitably show computational irreducibility (in the sense that the trajectories cannot be computed much more efficiently than effectively just simulating the whole sequence of interactions). But an observer who uses only bounded computation will inevitably see just a "coarse-grained" version of what's going on—and will "summarize" the behavior of the gas, say in terms of the laws of fluid motion.

And the point is that these laws have the same form regardless of the particular characteristics of the underlying molecules. They're in effect generic—and we can now understand this as an inevitable consequence of the interplay of computational boundedness of observers and computational irreducibility of underlying dynamics.

And in our new models of fundamental physics, fundamentally the same thing happens. Except that now the emerging "generic laws" turn out to be general relativity and quantum mechanics. And instead of "what's underneath" being a gas made of molecules we're now talking about what's underneath being a large number of abstract elements (that we can somewhat whimsically call "atoms of space") with certain abstract relations between them. We can think of these elements and relations as defining a hypergraph:

Then we can say that there's a rule such as

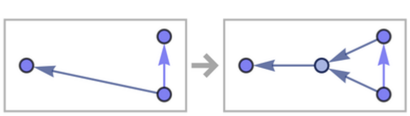

which specifies how updates should be done to this hypergraph. The effect of these updates will typically be a computationally irreducible process whose details depend on the observers in our actual universe must be. Assuming the observer uses only bounded amounts of computation, they can "read" only certain features of the hypergraph, and in so doing they (essentially) inevitably see the universe as having relativistic invariance, and ultimately as following the Einstein equations for the structure of spacetime.

There are many details to this whole picture. When we specify a rule it's effectively saying "here's how to update any piece of the hypergraph where the (left-hand side of the) rule matches". But in general there will be many different ways to do that. And the result is a whole bundle of "possible histories" for the system—which it turns out inevitably work according to the laws of quantum mechanics.

What are we assuming to come to these conclusions? Most fundamentally, there's what I call the Principle of Computational Equivalence. This principle says that rules whose behavior isn't obviously simple will all show the same level of computational sophistication—and they'll all exhibit computational irreducibility. And that's what in effect gives us a notion of time. In essence, the passage of time is just the inexorable and irreducible process of computation.

But there's another piece. Across our hypergraph there's a huge amount of irreducible computation that goes on—with all the various individual "atoms of space" in effect continually getting regenerated. But we humans have a particular way of experiencing all this—that involves sampling just those features that allow us to construct a coherent "thread of consciousness".

It's not obvious that such sampling could maintain its consistency over time, or could make different "consciousnesses" come to the same conclusions about the universe. But the phenomenon of causal invariance (which itself can be an emergent consequence of constructing a thread of consciousness) makes such things inevitable.

And in the end we can conclude that with our particular "sequentialized" way of "reading the universe" as computationally bounded observers, it's inevitable that we'll end up concluding that the universe follows the particular laws of general relativity and quantum mechanics that physics has established.

Underneath, there's all sorts of computational irreducibility that we're not directly sensitive to. But as computationally bounded observers we're picking out a certain slice of the slice we pick out corresponds to our known laws of physics.

But, OK, what this says is that given that we're "reading" the universe in a certain way, it's basically inevitable that we'll conclude that the universe follows certain laws. There could be almost any rule operating underneath; we'll always come to the same conclusion. But in and of itself this doesn't explain why there's a rule operating at all, or in effect why the universe exists, or why there's something rather than nothing.

The Laws of the Rulial Universe

As we work towards the question of why the universe exists, we need to talk about another formal, scientific result. We said above that our underlying rule is used wherever it applies in the hypergraph. And that doing this gives us a structure that we read as following relativity and quantum mechanics.

But so far we're still assuming that there's some particular rule for the universe—that's just getting applied in lots of different places. But what if instead there are lots of different rules that can be applied? In fact, what about the limiting case where any possible rule can be applied?

What if each piece of our hypergraph is updated according to all possible rules—generating many different possible histories? Or, in other words, what if the universe in some sense "simultaneously runs" all possible rules—generating all possible resulting histories?

Our first instinct might be that if all these possibilities are allowed then there could never be anything definite said about what would happen in the universe. But it turns out that this is far from correct. And it all has to do with the entangling of different possibilities associated with the repeated application of rules.

Given a particular state of the universe, applying different rules can lead to different states. But applying rules to different states can also potentially lead to the same state. Or, in other words, the "rulial multiway graph" that represents how one state leads to another can exhibit both branching and merging.

Let's look at a very simple example. Imagine that our rules just take a number x and add a constant n. Starting with 0, and allowing rules with n up to 5, we get after 1 step mostly just a bunch of separate branches:

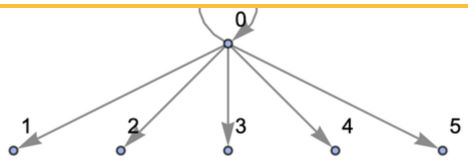

But after 2 steps, some of these branches merge, and we actually get a fairly complicated structure—in effect reflecting equivalences between different sequences of rule applications, associated with facts like 1 + 1 = 2:

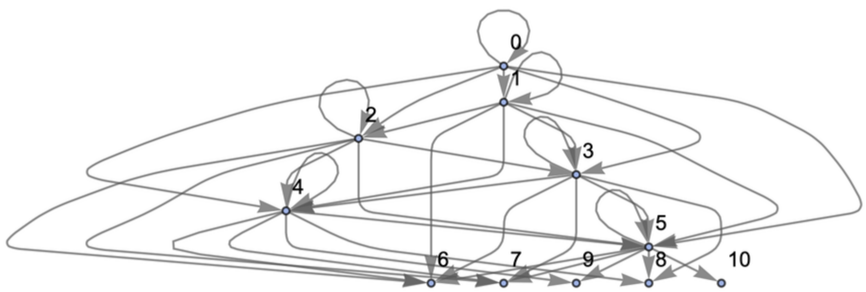

If we use, say, all possible addition rules mod 10 we get a more symmetrical structure:

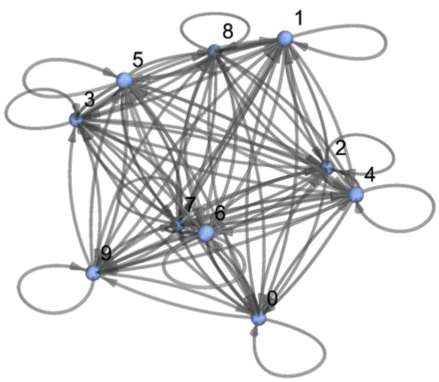

But the important point is that even though all possible rules (at least of a certain type) are being used, there's still a very definite structure that emerges.

So why does this happen? If we just enumerated the results of independently updating some state according to all possible rules, then, yes, in some sense all possible states could be produced. But the crucial point is that we're thinking about repeatedly applying all possible rules—and in doing this, there's inevitably a certain entanglement associated with the fact that different rules (or sequences of rules) can lead to the same states. So in some sense we generate structure from the interplay between the structural relations of states and the computational process of applying rules (or, in effect, in our models, the "passage of time").

rules is something purely formal. Yes, we might choose some particular language or basis for representing these rules (say in terms of Turing machines, or hypergraph rewriting). But in the end the concept of computation universality (or, more tightly, the Principle of Computational Equivalence) tells us that all these different languages or bases will be equivalent. So when we talk about all possible rules, we're not relying on anything about our particular universe. We're just describing something abstract, that is the way it is merely as a result of the definition of terms.

But now what we've seen is that just starting with this "formal concept" of all possible rules, we're getting a definite structure, which, yes, so far we're still also thinking of as something "purely formal".

Remember that we originally started from our models of physics, where we're talking about rules whose application determines the progress of time, the structure of space, and so on. But then we took the "formal limit" of considering all possible rules. And we found that the resulting "rulial multiway graph" has a definite structure. But what is the relation of this to our universe, and to the established laws of physics?

To answer that, we have to think about how this rulial multiway graph would be perceived by observers embedded within it. Looked at "from the outside" the rulial multiway graph involves many different threads of history, sometimes branching, sometimes merging. But an observer within the rulial multiway graph also involves many different threads of history, with the same kinds of branching and merging.

It's very much like the situation for quantum mechanics in our models: we have to work out how a "branching brain" will "read" a "branching universe"—except that now it's a "rulially branching brain" reading a "rulially branching universe". But it turns out that the same basic arguments apply.

There are many different possible "reference frames" that the observer can use to "organize" what they experience in the rulial multiway graph. But if the observer is computationally bounded—and adopts a description of the universe based on a "sequentialized thread of experience"—then it turns out that any reference frame they pick will give as perceived laws of physics general relativity and quantum mechanics.

underlying rules lead to the same perceived laws of physics. But now we're saying that this doesn't just happen for different specific rules; it happens even when the universe is in effect "running all possible rules".

So in a sense, in the end all that matters about the "underlying stuff" of the universe is its formal structure. The laws of physics are then just "formal consequences" of this structure, that we perceive to be the way they are because of the way we choose to "read" the universe, and in particular the fact that we construct a sequentialized thread of experience.

From Formal Inevitability to Actuality

The set of all possible rules is something purely formal—and something that in some sense has no structure. But what we've discovered is that from the formal computational process of applying these rules we inevitably get structure. And if we think about how an observer like us embedded within this system perceives what's going on, we conclude that they inevitably describe the system as following known laws of physics.

So how does this help us understand why the universe exists? We're starting from all possible rules. And basically we're saying that having a universe that operates in the way we perceive ours to operate is an inevitable consequence of there being all these possible rules. Or, in other words, if these rules "exist" then it follows that so will our universe.

But what does it mean for rules to "exist", and in particular for all possible rules to exist? The key point, I believe, is that it's in a sense an abstract necessity. The set of all possible rules is something purely formal. It can be represented in an infinite number of ways. But it's always there, existing as an abstract thing, completely independent of any particular instantiation.

It's crucial that we're talking about all possible rules. If we were talking about particular rules, then we'd need to specify which rules those are, and we'd need a whole language and structure for doing that. But that's not our situation. We're talking about all possible rules. We can construct some explicit symbolic representation for these rules, but the deductions we make ultimately won't depend on this; they would work the same whatever representation we chose to construct.

specific information. But what we're discovering is that our universe is in some sense like a tautology; it's something that has to be the way it is just because of the definition of terms. In effect, it exists because it has to, or in a sense because everything about it is a "logical inevitability", with no choice about anything.

As I mentioned in the previous section, we might have assumed that with all possible rules there'd never be anything definite to say. But the crucial scientific fact is that that isn't the case. Just from the formal idea of all possible rules, lots of specific detail follows.

Is it obvious that this would be the case? Well, no. It's taken the whole stack of ideas and development associated with our Physics Project to figure it out. Although it also actually turns out to be closely related to mathematical concepts, like the so-called infinity groupoid, that arise at what might be thought of as an outer limit of abstraction in modern pure mathematics.

Our Experience of the Universe

I've made an argument for why it's inevitable that the whole rulial universe exists. But what about our particular perception of the universe? Well, that perception is in some sense constructed by us. We are for example picking particular reference frames in which to organize our view of what's happening in the universe.

In the end, the result of our perception is again something abstract—our particular "symbolic description" of what's "actually happening" in the universe. And from this, we might then think that perhaps we're not really saying anything about things "existing". After all, we're starting from abstract rules, and in the end all we have is abstract perception.

But the big question is: why is there consistency in our perception of the universe? It could for example be that each different possible rule would make us get an utterly different impression of the universe. But again, the crucial scientific fact is that the rulial universe has certain inevitable consistent features, that will be perceived in essentially the same way by any "method of perception" that's at all like ours. And what this means is that it makes sense to talk about there being an "objective reality" independent of the details of our methods of perception—that for example we know includes the basic laws of physics.

define their experience of it. In fact, we can in effect imagine a whole rulial space of possible forms of description. It's a bit like saying we experience the familiar universe based on the particular place in ordinary space at which we find ourselves. Well, similarly we experience the whole rulial universe based on the place in rulial space that corresponds to the description of the universe that we use.

The overall structure of the whole rulial universe is inevitable. But the particular place in rulial space at which we find ourselves is a matter of how we choose things. And yes, given our particular history—both biological and intellectual—there are certain constraints on where in rulial space we can readily be. And as I recently discussed in connection with understanding the nature of consciousness those constraints (essentially our computational boundedness, and our notion of a definite thread of experience through time) inevitably lead us to the pillars of known physics: general relativity and quantum mechanics.

OK, so why does the universe exist? What we're saying is that its whole rulial structure is a logical inevitability. The way we choose to describe it, however, has a certain arbitrariness. But across different descriptions there is a fundamental consistency that lets us reasonably think that the universe does indeed exist.

Is there a shorter way to say what's going on? The main scientific result that we're using is that arbitrary abstraction, if played out completely enough, in a sense inevitably reads like physics to observers like us. So the universe—with the physics we perceive in it—exists because it's formally inevitable that it does.

Where Does Everything Come From?

Closely related to the question of why the universe exists is why there is something rather than nothing. And what we can now say is that there is something because—essentially as a matter of definition—all possible formal rules inevitably in some abstract sense exist.

The science says that if one applies these rules then eventually they build up the rulial universe. But why should they be applied? Why aren't they just static, abstract constructs?

In a sense it's confusing to talk about the rules "being applied". The rules just abstractly define what gives what—or how to "build out" the consequences of the rules. It's a bit like with an axiom system in mathematics. Once one has the axioms there's immediately a from the axioms.

But in a sense what we're talking about in terms of arbitrary rules is lower level than the usual treatment of axiom systems in mathematics. The rules define their own "way of being applied". And we are not just looking at "final consequences" of rules, but at the structure generated by the whole "step-by-step" process of every single possible rule application, with all the complex entanglements between them.

At some level, we can just think of rules abstractly being applied. But in our models of physics, there's an interpretation of that: sequences of rules being applied define the passage of time. There are "logically inevitable" chains of rule applications. But for us as observers embedded in the system—particularly with our way of setting up our thread of consciousness—we perceive this as time passing, or in effect, "things happening in the universe".

Talking about things this way might lead one to ask "What got the universe started?" "What makes rules actually get applied?" Well, nothing. Because the rules are just defining how abstract sequences can be constructed. And if you follow a sequence, it can be interpreted as reflecting the passage of time. But there's no "driver" that's saying anything like "now this rule gets applied". The sequences generated by the successive application of rules are somehow just abstract "logical possibilities".

But now remember that we are observers embedded within this whole setup, with everything about our operation also defined by these same abstract rules. So, yes, we can interpret rule applications as being associated with the passage of time. But it's not like time has to pass at a certain rate. Time for us as observers, and for the universe that we observe, in a sense "passes when it passes".

There are rule applications that can occur in the universe as a whole, and there are rule applications that can occur in us. And it's just the abstract interplay of those that makes us perceive the universe the way we do.

Nothing had to "get the rule applications started". They are abstractly defined, and abstract sequences of them correspond to the evolution of the universe and us within it. Like the rules themselves, they are inevitable abstract constructions. And these abstract constructions when played out completely enough yield what we can interpret as the whole structure just in some sense "abstractly exists".

So does this mean that the whole history of the universe can be thought of as an "abstract thing" that is somehow "immediately defined"? Well, no, not in any useful sense. Because the phenomenon of computational irreducibility implies that longer chains of rule applications in a sense "do irreducibly more" or "get irreducibly further". They define a structure that irreducibly needs those rule applications. There's no general way to "jump ahead" and get the outcome without going through all the rule applications.

Computational irreducibility is what in a sense builds up structure in the universe. All possible rules are being applied. But there's an inevitable and irreducible pattern to it, which we interpret as being built up by the passage of time. And while everything about it is in a sense "abstractly defined" from the outset, there's something "fundamentally heavier" about these "consequences" as they are played out.

By the way, in addition to asking how the rules "got started" one might try to ask "what did they start with", or in other words, what was the "initial data structure" of the universe? But once again, we don't need to ask this. Because among "all possible rules" will be ones that effectively create any possible starting condition.

But then, one might ask, why doesn't this mean that absolutely anything could just be inserted into the universe from the beginning? It's all rather subtle, and not as formally worked out as I hope it will be. But the interplay of computational irreducibility, computation universality and the embedding of us as observers using the same rules as the universe does seem to imply that the ultimate perceived structure of the rulial universe can be thought of as being built from "basic rules" that have limited complexity, even though in a sense all possible rules can contribute.

Is This the Only Universe?

If the universe were based on a particular underlying rule there would seem to be no reason why there shouldn't be other universes based on other rules. But an initially quite surprising implication of our approach is that actually the universe is in effect based on all formally possible rules. And a consequence of this is that there can only be one universe—which, as we've argued, in some sense inevitably exists.

experience only a small part of it. In physical space we're used to the fact that we live on a particular planet at the edge of a particular galaxy. But what we now realize is that we are also sampling only a tiny part of the rulial space of all possible descriptions of the universe. Were our sensory apparatus or our intellectual development different, we might describe the universe in very different ways.

As I've argued elsewhere, however, as soon as we imagine that we're operating with something like consciousness, there are constraints, in particular that we must be computationally bounded observers who describe things in terms of a single thread of experience through time. And it's then a crucial fact that these constraints alone make it inevitable that we'll attribute to the universe a physics with familiar core features like general relativity and quantum mechanics.

Even though in the full rulial universe there are lots of other possibilities, the features that we sample, with our consciousness, follow our familiar laws of physics. Or in other words, we have a particular perception of the universe, that operates according to particular laws.

But the crucial point is that those particular laws are not a fundamental feature of the universe; they are merely a feature of our description and our sampling of the universe. And just as we can imagine using spacecraft or telescopes to study distant regions in physical space, we can imagine using different forms of description and analysis to study different parts of rulial space.

I won't discuss this in detail here, but one can for example imagine motion in rulial space: progressively changing with time one's description of the universe. And just like the speed of light limits motion in physical space, there's a constant we call ρ (but whose value we don't yet know) that limits motion in rulial space.

When we first asked the question "Why does the universe exist?" it seemed difficult to understand how one could possibly make an abstract argument that would conclude anything about the existence of the particular universe in which we seem to find ourselves. But the key point is that the full rulial universe involves no particular choices; it is something formally inevitable. But what we perceive as our universe is just a part of that full rulial universe, and a part determined by the particulars of how we—with our consciousness—choose to describe the world.

same rulial universe. But it might be quite distant in rulial space. And so its perception of the universe and its laws might be quite different from ours. It might, for example, have none of our familiar notions of space or time, but instead describe the universe in terms of features of the configurations of atoms of space and their causal connections which we are not set up to access or reason about.

Could we in principle translate from their description to ours? In a sense this involves motion in rulial space—so there are limitations imposed by ρ. But as a practical matter it is something immensely challenging, as our inability to translate even from descriptions presumably close in rulial space, like animals or AIs, suggests.

If we restrict ourselves to descriptions associated with something like consciousness then we've argued that we inevitably attribute to our universe our familiar laws of physics. But without that restriction, is there anything we can say about our universe, and what laws we'll attribute to it? In other words, even beyond the particular slice that we sample, are there general laws of the rulial universe?

The answer is that there is definitely one such law, namely that the universe is fundamentally computational. In other words, whatever description language we use, we can always ultimately represent what the universe is doing in terms of the operation of a standard universal computer, like a universal Turing machine.

We've argued that the rulial universe is in a sense just a representation of the (entangled and computationally irreducible) inevitable consequences of following all possible formal rules. But when we talk about all possible formal rules we're implicitly making the assumption that these rules can be stated in some kind of explicit symbolic form. And this is why we conclude that everything in the rulial universe is computational—in the sense that it can be represented by a standard universal computer—rather than "hypercomputational" (e.g. requiring an oracle that can immediately answer questions that would require infinite computation in a standard universal computer, or involving "hyperrules" that directly represent such operations).

So could there actually be a "hyperrulial universe" that also contains formal rules that involve hypercomputation? The answer is that, yes, there certainly could be. But in some sense it doesn't matter. Because if the way we sample the rulial universe involves at most ordinary universal computation, we will never probe these hypercomputational elements.

shielded from us by a rulial event horizon.

So, yes, a "hyperuniverse" could exist. But it'd never be relevant to us, or to our goal of understanding why the universe—as we perceive it—exists.

To recap then: we start from the idea that all possible formal rules "exist" as abstract constructs that follow directly from abstract definitions. Then we argue that the rulial universe is an inevitable structure derived from these rules. We as observers choose a certain way (consistent for example with our notion of consciousness) to describe this structure. And that form of description is what gives us our perceived laws of physics.

The existence and structure of the rulial universe is in a sense a necessary abstract fact (like 1 + 1 = 2). Our particular sampling of it is a "point of view" we choose. But whatever that "point of view" is—or in other words wherever in rulial space we are—we will still conclude that there is a universe, or that the universe exists. The laws we perceive will depend on our "point of view"—though we know that just adding the constraint of having a "consciousness-based point of view" will already give us core familiar laws of physics.

When it comes to a hyperuniverse, we can say that, yes, if our point of view lay in hyperrulial space then we might want to address why the hyperuniverse exists. But as it is, our point of view is just somewhere in ordinary rulial space, where we know that there's a "rulial universe that exists underneath".

Why This Universe?

Why did we get this universe, with its detailed features, and not another? In the past we might have assumed that this involved arbitrariness in the creation or setup of the universe itself. And indeed endless mythologies describe mechanisms by which choices about the universe might have been made. In recent times it's also been common in some circles to talk about the possibility that the universe might be a "simulation", operating like a videogame that's being run on some kind of "outer computer".

But what I've argued here is that in some sense there's no choice about the fundamental structure of the universe. At the level of the full rulial universe everything is inevitable and necessary. There is no choice being made about the rule—or program—for the universe. All universe has the structure it does.

What is "contingent", however, is how we sample this full rulial universe. We have a certain point of view on the universe, and that is what gives us the particulars we perceive. We could have a different point of view. And maybe at some point in the development of our species and our civilization we will have. And at that point we might describe the universe differently than we do now, invoking different laws of physics to explain what we see in the context of the description we're then using.

But the crucial point is that there is no choice at the level of the fundamental structure of the universe. The choice is about our way of describing the universe, and in effect where we happen to be in rulial space.

If we think about the universe as a simulation, what's being said is that there's no choice of "which videogame is being run". In a sense it's all possible videogames—with all possible underlying rules. And the point is that the resulting rulial universe is just an inevitable formal consequence of the effects of those underlying rules. As a matter of whimsy one could choose to talk about the rules as "running on an outer computer". But it would be like talking about 1 + 1 = 2 in terms of counting physical stones. None of that "instantiation story" makes any difference to the structure of what comes out.

It's perhaps useful to talk more practically about actual "existence inside a videogame". Imagine at some point in the future a digital representation of our brain functions can be uploaded to a computer. And imagine the computer is running some kind of "universe operating system" that implements the videogame—with our "digital minds" inside it. The question is then what perception our digital minds will have about the environment they're in, and for example what its effective laws of physics are.

We might at first assume that we'll just perceive whatever laws of physics were "programmed into the videogame". But it's more complicated than that. Because if we're "part of the videogame" our minds must operate according to the same programming as everything else—which is much closer to the setup for our models of fundamental physics. And—assuming that the "universe operating system" implements universal computation (which presumably it must in order to include our digital minds)—we're basically thrust back into essentially the same formal situation we've been talking about.

experience. But without those, we're again just dealing with a formal, computational system. What about the idea of "running all possible rules"? Well, if we have a universal computational system we can in principle do that. And build up the same formal structure of the rulial universe.

But then the issue is: where in this rulial universe will our digital minds "naturally find themselves"? And the point is that there is no reason to think it will have any correspondence to the "rulial location" for our non-uploaded minds in the ordinary universe. In other words, the description we give of the ordinary universe may be completely different from the description our uploaded minds will give in the "simulation universe".

In both cases (ignoring "resource constraints") there's the same full rulial universe. But we're potentially sampling different parts of it—with views of the universe as incoherent as those that we and putative aliens might have.

Prime Movers and the Path from Abstraction to Abstraction

In a sense the question of why the universe exists is all about why anything is "actualized". It's one thing to say that the universe can be represented by formal rules; it's another to say that those rules are "actualized". And many past analyses have concluded that there has to be some kind of "prime mover" that "breathes reality" into the formal rules. But what I've argued here is that nothing like that is needed.

Instead, the whole elaborate structure of the rulial universe inevitably arises as a necessary consequence of the very definition of all possible formal rules. And the point is that we—as observers embedded in this rulial universe—have a certain perception of what is happening that in a sense "constitutes our reality".

But just where in all of this did the merely formal turn into something "real" or "actual"? Well, at some level, it ultimately didn't. Because in the end what we're talking about is our perception—or in effect our abstract description of what's happening in the universe. So in some sense, everything we're doing is just going from abstraction to abstraction. So where's the "actualization" or the "reality"?

invariance to the final abstract description that we form of the universe. If we started from all formally possible rules and in a sense everyone's perception of what was going on in the universe did something like "picked a different rule" and gave a completely different abstract description of the universe, then it wouldn't be useful to think of anything as being "actualized".

But instead we know that—at least so long as we follow the constraints associated with consciousness—there'll be great alignment between the different abstract descriptions we end up with, and for example they'll all agree about the basic laws of physics.

I suppose we could say that "reality is an illusion" in the sense that in the end we're going from pure abstraction to pure abstraction. But the whole point is that as a matter of science there's necessarily agreement about what "comes in the middle". And that's why we can usefully talk about "reality" as an objective kind of thing.

And the formal inevitability of this is what we can say is "why the universe exists".

The Relation to Mathematics

Most people would view it as self-evident that the physical universe exists. But the question of whether mathematics "intrinsically exists as a definite thing" has been debated since at least Plato.

One view of mathematics is that it involves just writing down whatever axiom system one wants, and then working out its consequences. And at least at first this seems very different from physics, where one might imagine that one starts from our particular world as it is, and then tries to find a formal representation (or "model") of it. And, yes, our effort to find a complete and precise fundamental theory of physics can be thought of as trying to "reduce physics to mathematics"—in the sense that it's trying to give us a particular formal ("axiomatic") system that reproduces what the physical universe does.

At first, we might imagine that this formal system for physics must be something very special, and not something like the "arbitrary axiom systems" that we could choose to write down as foundations for mathematics. But from what we've seen here, the situation is more subtle than this. Because "underneath physics" there are in a sense all possible formal systems.

than physics, because it can operate with whatever formal system we write down. But now it's seeming like the opposite: physics is based on all possible formal systems, but mathematics is based on particular formal systems (geometry, algebra, etc.) that we happen to have written down in the history of human mathematics.

We've argued that the existence of the universe is ultimately a consequence of the fact that all possible formal systems exist as a matter of abstract necessity. So can we use a similar argument for the "existence of mathematics"?

Most likely we can—at least if we subtly change our description of what mathematics is. In the "axiomatic tradition" it's been common to imagine that mathematics could in principle be based on whatever formal axioms we want, although in practice we pick particular ones. But an alternative view is that ultimately mathematics, like physics, is actually based on all possible formal (axiom) systems.

In physics, we describe the construct created from all possible formal systems as the rulial universe. So what is the analog in mathematics? Well, it's the same thing!

In the actual pure mathematics of the past century it's become increasingly common to study spaces of all possible things of certain types. But the analog of the rulial universe is a kind of ultimate limiting example, probably best captured through higher category theory as the infinity groupoid. And, yes, the remarkable fact is that constructs like this—formed "from all possibilities"—actually end up having definite, rich structure.

But, OK, so in the interpretation that mathematics is what's formed from all possible formal systems there's a definite thing underlying it—and it's the very same rulial structure that underlies physics.

At the lowest level the structure is full of computational irreducibility. But in the case of physics we know that we as computationally bounded "single-thread-of-experience" observers sample particular slices of the structure that show computational reducibility—and give us known physical laws.

So how do we sample the rulial structure when we do mathematics? I'm not quite sure what the metamathematical analog of the constraints of consciousness in physics are—though the concept of combining things into reference frames may be related to the "equivalence is equivalent to equality" concept captured by the univalence axiom in homotopy type theory.

language for mathematics (say, in a simple case, algebraic vs. geometrical, etc.). But then it's a fundamental feature of the rulial structure (essentially a consequence of causal invariance) that different reference frames will give equivalent results. In physics this leads to general relativity and quantum mechanics—and a meaningful notion of "objective reality" for spacetime and for quantum phenomena.

And—assuming we can construct something like reference frames—there should be similar "global laws of metamathematics", that in effect describe the "objective reality of mathematics" independent of the particular mode of description we use. I don't yet know quite how to formulate these global laws, though perhaps approaches like category theory give suggestions.

In explaining why the physical universe exists we get part of the way there by talking about the abstract necessity of the rulial universe. And with the idea that the same rulial structure exists "underneath mathematics" we again get part of the way to explaining that mathematics "exists as a definite thing".

But to get further I think we have to consider "perception by mathematical observers" just as we consider "perception by physical observers". And what's crucial about physical observers like us is that they manage to sample computationally reducible slices of the rulial universe—which they perceive as containing "definite things" consistent for example between different observers.

One of the mysteries of mathematics is that it's possible to make progress without continually getting mired in computational irreducibility and phenomena like undecidability. But in some sense this is just the mathematical analog of the fact that we as observers in the physical universe can identify slices of computational reducibility—and not be mired in computational irreducibility and utter unpredictability.

I must say that in the past I wasn't at all sure that mathematics could be thought of as "fundamentally existing as a definite thing". I thought of it more as a human-created artifact, built on specific axiomatic foundations that were in some sense historical accidents. But now that we understand more about why the physical universe exists, it's looking to me more likely that we should think of mathematics also as a "thing that exists".

particular reference frames—into the ultimate structure that "is mathematics".

How Should We Feel about All This?

Can the universe really at some level just be a formal thing? As conceptually elegant—or even beautiful—as this conclusion seems, as I sit here typing these words, some part of me still resists it. After all, the keyboard under my fingers feels like a real thing. The screen in front of me I can reach out and touch. But of course the history of science keeps on showing us how wrong our feelings and intuition can be.

Four centuries ago, for example, we learned that—despite the immediate evidence of our senses—the Earth goes around the Sun. And now what I'm arguing is that again—despite the immediate evidence of our senses— everything we experience is actually just a reflection of an abstract formal structure. And that there is no "special spark of reality" associated with anything.

Our increasingly virtualized modern world might make it seem more plausible that this could be correct. But there is still something shocking about the notion that in a sense everything is fundamentally virtual, or formal. And that there is no ultimate underlying "real substrate"; everything is "formal all the way down".

At first, that might seem disappointing. We might like to think that the things that happen —or that we do—in the universe are somehow "real", not just virtual or formal. After all, we might think, if everything is virtual, nothing can ever fundamentally be "achieved" by, for example, the passage of time.

But there's an important sense in which this is wrong. Because even though an outcome may be "virtual", there may still be something definite and irreducible that's needed to produce it. And in particular the phenomenon of computational irreducibility—together with the Principle of Computational Equivalence—implies that there are many processes defined by formal rules where the only way to find their outcome is effectively to run the rules an irreducible number of times.

If one could always "jump ahead" and immediately determine the outcome one could reasonably say that "nothing is achieved" by "living through" the whole evolution of the effect that it achieves something that couldn't have been achieved without it.

Ever since Copernicus science has been showing us ways in which we are not special. The Earth is not at the center of everything. Life is not made from different stuff than everything else. And indeed my own previous work has argued that the Principle of Computational Equivalence implies that there's nothing fundamentally special about intelligence, beyond "mere computation".

And what I'm arguing here is that there's nothing special even about our whole universe; it's all just determined by abstract necessity. At some level we might see this as the ultimate "Copernican put-down", and the ultimate statement of how "unspecial" and unimportant we are.

But there's a different view one can take—that in a sense instead puts us firmly in the center of everything. Yes, the whole rulial universe is an abstract necessity. But the particular description language—or "rulial position"—that we occupy is something completely special to us.

In the space of all possible rulial positions, the one we occupy can be viewed as arbitrary and "unspecial". But for everything about the universe as we perceive it, it is crucial and central—and in a sense it is what creates for us the universe as we know it.

We might have thought that the ultimate story of science would be about what the universe "gives to us". But what the arguments here imply instead is that it's actually about what we "take from the universe"—and how we are set up to sample the whole rulial universe.

It's notable that our way of describing—or sampling—the rulial universe is not just determined by the "pure physicality" of where we are in physical space, or how we're configured as biological organisms. It also depends on how we "mentally model" the world, and how our formal thinking describes what we perceive is going on.

We might have hoped that intelligence would be something that science would show is fundamentally special about us. But the Principle of Computational Equivalence implies that "abstract intelligence" is actually a common and general phenomenon. So does this mean that there is nothing special about us? Well, no. It just highlights the importance of all those particular details that make human intelligence and the human condition what it details that are actually what's important.

And so it is, I think, when it comes to looking at our whole universe. What is special and significant isn't some general aspect of what underlies the structure of the universe. Instead, it's the details of how we—as humans—describe the universe. (Though, as we've discussed, the general "consciousness features" of that description robustly give us laws of physics such as relativity and quantum mechanics.)

So, in the end, as we reach what we might see as the ultimate foundations of science, we are in some sense back to us humans being at the center of everything. We might have thought that somehow science would long ago have "abstracted us humans" away. But instead, the only way to make anything meaningful is in a sense to put us right in the middle of everything, with features of our perception defining how we describe the universe, and what we consider the reality of the universe.

As I look around at my physical environment it's still a little shocking to think that it's all in a sense just "created by me" by "knitting together" some kind of generic "formal abstraction". Yes, the computational boundedness and sequentialized thread of experience that seem to form the basis for what we view as consciousness force certain properties of the physical environment that correspond to laws of physics we know. But some putative alien—or even future human or human-like entity— could still perceive this physical environment quite differently, even though it's always ultimately based on the same formal structure.

But while it might seem like a different universe, it'll still seem like some universe (albeit perhaps one incoherently different from ours). And however we perceive it, the conclusion will be the same: from pure abstract necessity it will follow that the universe exists, and that there is something rather than nothing.

Notes

I've tried here to sketch my recent thinking about the question defined in the title. I've suspected for some time that the discoveries of our Physics Project (and its precursors in my work from the 1990s) might perhaps have something to say about the question. But in the recent thinking I describe here I've been able to make what at least seems like much more headway than I'd expected.

Contents

And it would certainly be very interesting to know how what I say here—with its foundations in our recent science—might relate to what has been said before. But I expect this will be a challenging project—that will inevitably need to build intellectual bridges between very different kinds of thinking, and I haven't personally tried to do it.

It's worth emphasizing that what I say in this piece is just a sketch. And among other things, I've glossed over many technical details, especially about the structure of rulial space. And while both Jonathan Gorard and I have made studies of at least some aspects of rulial space, there's considerably more to figure out.

Thanks

Thanks particularly to Elise Cawley, Hatem Elshatlawy, Jonathan Gorard, Rowland Stout, Alexander Wolfram and Christopher Wolfram for discussions and suggestions that have helped clarify various aspects of what I talk about here.

Posted in: Big Picture, New Kind of Science, Philosophy, Physics

Join the discussion

+ 10 comments

Related Writings

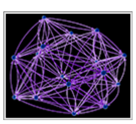

The Concept of the Ruliad
November 10, 2021

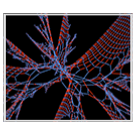

Multicomputation with Numbers: The Case of Simple Multiway Systems
October 7, 2021

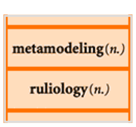

Charting a Course for "Complexity": Metamodeling, Ruliology and More
September 23, 2021

Contents

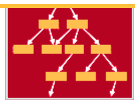

Even beyond Physics: Introducing Multicomputation as a Fourth General Paradigm for Theoretical Science

September 9, 2021

Popular Categories

- Artificial Intelligence
- Big Picture
- Companies and Business
- Computational Science
- Computational Thinking
- Data Science
- Education
- Future Perspectives
- Historical Perspectives
- Language and Communication
- Life and Times
- Life Science
- Mathematica
- Mathematics
- New Kind of Science
- New Technology
- Personal Analytics
- Philosophy
- Physics
- Ruliology
- Software Design
- Wolfram|Alpha
- Wolfram|One
- Wolfram Language
- Other

Writings by Year

2021 | 2020 | 2019 | 2018 | 2017 | 2016 | 2015 | 2014 | 2013 | 2012 | 2011 | 2010 | 2009 | 2008 | 2007 | 2006 | 2004 | 2003 | All