

Bergsonian Axioms For Physics

Doug McLellan

Summer 2016 Draft

We show here how a more philosophically satisfying foundation for physics can grow from a century-old insight of Henri Bergson, that *the world is the continual creation of new possibilities, rather than the successive realization of pre-existing possibilities*. This insight invigorated fields both scientific and humanistic, and helped win Bergson a Nobel prize (Literature, 1927), yet it vanished from the scientific scene because the mathematics did not exist to make it precise. Today these mathematics do exist; we will use them to translate Bergson's insight into axioms that a mathematical model of physics ought to satisfy.

The core idea motivating this paper is a mathematical form of Bergson's insight: *new mathematical structures are made possible over time*. We will lay out this idea's philosophical appeal in our first main section, where we take up each of four long-standing puzzles, show why each withstands current lines of attack, and let the core idea suggest a postulate that might resolve it. The puzzles that give rise to the postulates are (1) the arbitrariness of natural laws, (2) the infamous set-of-all-sets paradox, (3) the arbitrariness of space-time structure, and (4) the independence of the continuum hypothesis from all "normal" mathematical considerations. Although this first section presupposes familiarity with light cones and at least the gist of the Cohen/Gödel independence results, it operates at a very general level, with no real physics at all. Our aim here is to help the reader see how a shift to a Bergsonian viewpoint is possible, without invoking too much technical detail.

Our second section translates the four postulates into precise mathematical axioms, and outlines the progress that has been made in the search for a model of them. The axioms themselves are fundamentally set-theoretic. The framework for seeking a model of them is therefore not any particular branch of physics; it turns out to be, more or less necessarily, that of set-theoretic

“forcing extensions.” A forcing extension’s properties are determined by the boolean algebra used to obtain it. The challenge posed by our axioms is thus to find a boolean algebra satisfying properties that they dictate.

The author’s initial plan was to find some well-known boolean algebra forcing a simple model of the axioms, and then commend to smarter folks the work of transferring this result to a physically relevant setting. On one hand we can imagine that this scenario, if it succeeded, would secure some philosophical worth for the Bergsonian postulates; on the other, it would probably impose too few constraints on physical theories to yield much properly scientific insight.

But we will see that none of the best-known boolean algebras yields models of our axioms. In fact the only promising ones we are aware of are derived from systems of von Neumann algebras like those used in algebraic quantum field theory. This is unfortunate in one sense, because these structures are notoriously complex. In another sense, though, it is intriguing: what changes would be required for AQFT to satisfy all the Bergsonian axioms? An interesting answer here may provide not only philosophical gratification but testable scientific hypotheses.

Foreword On The Creation Of Mathematical Possibility

The best way to begin explaining our core idea, that new mathematical structures are made possible over time, may be to insist that we mean it literally, straightforwardly. Otherwise we might be suspected of preparing to sublimate it in a puff of wordplay, with fancy glosses for “possible” and “structure.” We have no such intention. The kind of possibility we have in mind is ordinary, philosophically-naive mathematical possibility, and the structures that attain it are, for instance, real numbers. It will help to regard all these structures as sets. This will let us phrase our core idea more clearly: we augment the usual “realist” or “platonic” position in the foundations of mathematics, that there really is a universe V of possible sets, with the claim that V grows over time. Most importantly, this set-theoretic perspective will equip us to see *how* and *why* new mathematical possibilities might arise.

One could be forgiven for suspecting that our core idea refutes itself immediately. If a mathematical structure — the fraction $5/8$, say — is newly made possible, then it would have been correct in the past to say “ $5/8$ is not yet possible”; but the mere fact that $5/8$ could have been *consistently*

referred to back then should have sufficed to ensure its mathematical possibility. In a word, thinkability ought to imply (mathematical) possibility. This argument is valid as far as it goes: any unambiguously *definable* structure should indeed be *eternally* possible. But the argument only goes as far as definable structures. If there are undefinable structures (say, infinite sequences of digits) then our core idea has room to operate. This paper will make little sense unless it is kept in mind that, whenever we speak of “newly possible structures,” we mean structures that are complex beyond definition.

Just because the notion of possibility-creation enjoys a modicum of self-consistency, it does not follow that it is natural or intuitive. On the contrary, it is deeply odd; its only champion to date has been Henri Bergson, the mathematically gifted French philosopher who became famous in the early twentieth century.¹ Seeds of the possibility-creation notion appear in his *Creative Evolution* of 1907. They reach maturity in the 1920 paper “The Possible And The Real,” where he argues sharply that time and freedom make no sense without it. At any rate, *we* find his arguments sharp. Bergson fell out of fashion by the 1930s and his few latter-day boosters tiptoe around the idea of possibility-creation. That a brilliant Nobel laureate should have failed to secure sustained interest in this idea is a warning to anyone who would relaunch it today; in fact it is a source of suspicion serious enough for us to block right now, and we give some quick reasons why we expect to have more luck than Bergson had.

Three related obstacles kept Bergson’s idea from enduring. The first was his mischosen emphasis: possibility-creation for Bergson took a more biological than mathematical form, and this biological form turned out not to bear much fruit. We will avoid this obstacle ourselves by choosing the other approach. In Bergson’s defense, though, his choice may have been forced by the second obstacle: the non-existence, during his lifetime, of mathematical tools for handling possibility-creation. The early twentieth century was the heyday of the so-called logicians’ quest to reduce all mathematical truths to formal tautologies. “Undefinable structure” would have been a senseless phrase in this context — if you can’t write it down, you can’t reduce it to a tautology — and we have already stressed that undefinable structures are needed if possibility-creation is to work as a mathematical idea.

We get around Bergson’s second obstacle thanks to our set-theoretic ap-

¹A much-repeated anecdote blames the world’s first traffic jam on people thronging to Bergson’s lecture at Columbia University in 1913.

proach. No single philosophy of mathematics has inherited the popularity of logicism, which was cut down by Gödel's incompleteness theorem, but the idea that set theory underlies mathematics has gained wide acceptance. Within set theory one can speak meaningfully of undefinable structures; Gödel first posed the question of whether there are any. More precisely, he asked whether any sets fail to have the related property called *constructibility*. In 1938 he proved that the standard axioms of set theory fail to guarantee unconstructible sets. This might have been the death knell of the possibility-creation idea, had Paul Cohen not proved in 1963 that the same axioms also *fail to preclude* unconstructible sets. For all the standard axioms of set theory tell us, then, there might or might not be the sort of structures that our core idea needs. In any case, the last four decades have seen unconstructible sets become central to higher set theory. We thus have the tools Bergson lacked for making a substantive theory of possibility-creation.

The invention of these tools seems not, however, to have suggested to anyone else the use that we are now proposing. This points to a third and more formidable obstacle behind Bergson's failure: the faith of nearly everyone since Plato in mathematics' *timelessness*. The *Republic* declared geometry — like all of mathematics, presumably — to be “knowledge of the eternal, and not of aught perishing and transient.” Hardly anyone in the twenty-five centuries since has thought otherwise. What can be done about this obstacle? One strategy would be to analyze this faith into its various threads and show that none of them enjoys *a priori* necessity. But Bergson himself tried this; the last quarter of *Creative Evolution* is a far more thorough critique of platonism and neoplatonism than we could ever mount, and it made hardly a dent in the orthodoxy.

For those who would take up his fight against platonist orthodoxy, Bergson's career presents few episodes that are encouraging, but one that is quite instructive: his quarrel with Einstein. At the time of their first face-to-face meeting, in 1922, the two men's theories of space and time were both famous. Bergson's theory was impressionistic and backed up by compelling metaphors; Einstein's was quantitative and backed up by compelling formulas. From where we stand today it is fair to say that the quantitative theory won outright, and we hardly need to examine the arguments' details to know that it won chiefly *because* it was quantitative.² The lesson is clear for anyone

²But the details are nicely examined in *The Physicist And The Philosopher* (Princeton, 2015) by Jimena Canales.

building an alternative approach to foundational questions: *show precisely how it fits into modern scientific theories*. Vaunting its philosophical advantages from a humanistic point of view will win it few lasting friends.

Our strategy for dealing with the third obstacle is therefore to turn possibility-creation into a rigorous physical theory. We now turn to the puzzles that will give rise to postulates for such a theory.

Puzzle 1: The Arbitrariness of Physical Laws

God does not play dice with the universe: You have surely heard this rebuke of arbitrariness, and you surely know that Einstein meant it to deny that quantum laws are random. The arbitrariness that concerns us here, however, is of a broader kind. Grant Einstein what he insists; grant (which may be more, depending on your definitions) that nature’s laws are deterministic. They may still be arbitrary in this sense: there may be no explanation for why *these* deterministic laws govern the universe, rather than *those* deterministic laws. This worry is our starting point.

No sooner have we announced this starting point than certain readers will sigh: “Worries about nature’s *whys* belong in the cloud-cuckoo land of philosophy; real science keeps its feet on the ground, describing phenomena, generalizing from them by induction, and making testable predictions.” We sympathize with these defenders of empiricism. Historically speaking, science’s empirical strain does tend to outperform what we might call its “worrying about the why” strain. But it is not always idle to demand reasons. Relativity, for example, can be considered an answer to the question: “Why and on what grounds might nature choose a privileged rest frame if light’s speed is the same to all observers—and how might nature arrange itself so as to do without such a choice?”

In this spirit, we are going to build a theory of how nature might evade *all* arbitrary choices, even of its own laws. To do so we must first set up our puzzle precisely, and show that existing approaches do not resolve it. This will require a formal and very general framework for space-time theories: see Figure 1 below. Alongside the definitions we will sketch simple illustrations of them. Please keep in mind that the illustrations’ simplifying features, like the finiteness of space-time, do not lessen this framework’s ability to handle more realistic cases.

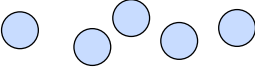
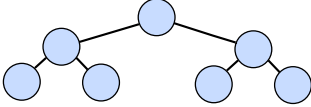
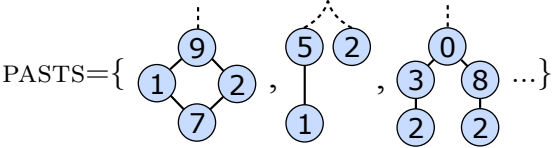
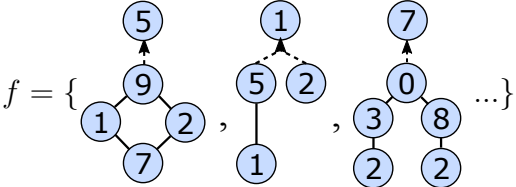
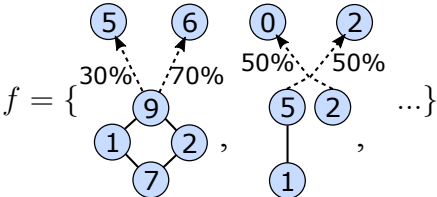
Definitions	Simple Illustrated Examples
Let the word <i>point</i> refer to space-time points (a.k.a. events).	
Stipulate that $<$ will always be an “absolutely earlier than” partial-ordering relation on some set of points, so $p < q$ means that p is in q ’s past light cone.	 <p>(a point p is $<$ another point q if one can reach q from p via lines going upwards)</p>
Assume that the state of each point consists of mathematical data (mass, charge, spin, state vector, and/or something more complex), and let all possible states be collected in a set called STATES.	<p>STATES = {‘mass=0’, ‘mass=1’, ‘mass=2’ ...} = {0, 1, 2, 3...}</p> <p>A point’s state will be drawn inside it, like so: ②</p>
Let PASTS be the set of possible absolute pasts that a point might have; specifically, of all structures $\langle P, <, \phi \rangle$ where P is a set of points, $<$ its ordering as above, and ϕ a function from P to STATES.	<p>PASTS = {</p>  <p>...}</p>
Let a <i>causal natural law</i> be a function from PASTS to STATES, <i>i.e.</i> a function from pasts to present states of the universe.	<p>$f = \{$</p>  <p>...}</p>
Let CNL be the set of all causal natural laws, <i>i.e.</i> $\{f : f \text{ is a function from PASTS to STATES}\}$.	<p>CNL = $\{f, g, h, \dots\}$</p>
If you wish, modify the above definitions so that states can have probabilities at a point, and so that causal natural laws can incorporate randomness. The exact method will not matter for our purposes.	<p>$f = \{$</p>  <p>...}</p>

Figure 1: Our formal framework for space-time theories

With the definitions of Figure 1 established, we can pose our puzzle precisely:

Puzzle 1. The Law of the Universe is some particular function $f \in \text{CNL}$, but why not some other function g or h ?

Now, the choice of a natural law from among CNL is not the only arbitrariness in the above definitions; the order type of $<$ in our real world $\langle P, <, \phi \rangle$ is also unexplained, as is the extension of the set STATES. You are more than welcome to make a note of these gaps, to which we will return, but we will begin by worrying about the arbitrariness of f .

Plan I to Deal With Arbitrariness: Deny, On Skeptical Grounds, That The Puzzle Can Even Be Posed

This puzzle's set-up has put us squarely on the side of that famous philosopher who said, "The world is the totality of facts." This may bother some readers — fans of philosophical "pragmatism," "anti-realism," or "skepticism" — who prefer an earlier philosopher's line, "There are no facts, only interpretations." We know where these readers are coming from. To us too, in certain moods, the world seems an unfathomable mystery; in these moods we can only laugh at anyone claiming to dissect it into discrete "points" or "states."

What we would like to stress, though, is that we indulge these moods largely *because* they soothe our worries about arbitrariness. By letting our world be an unfathomable mystery rather than a set of facts, we spare ourselves the needling question, "Why is *this* set of facts real, rather than *that* set?". We trade in our feeling of arbitrariness for one of mute mystical awe.

This is not an unequivocally good trade. Skepticism/pragmatism/anti-realism has its drawbacks too, the worst being its inability to explain the tremendous success of quantitative science. Current theories of "photons" and "quarks" are accurate to as many decimal points as you care to measure. The simplest way to explain their success is to admit that their terms really do refer to nature's constituent parts. And as soon as we appreciate this, our philosophical pendulum will swing back the other way, towards the world-as-totality-of-facts view. "Scientific realism" and "anti-realism" both have their attractive and repellent sides, then, and what results is a philosophical yo-yo

effect. We are seeking a way to avoid the yo-yo-ing, a more *stable* response to the problem of arbitrariness than shoulder-shrugging mystagogy. So let us put ourselves back in the “world as totality of facts” mood, and see if there is a better direction to go in, once the question of arbitrariness starts nudging us out of it.

Plan II to Deal With Arbitrariness: Just Accept It

The puzzle, again, is this: The Law of the Universe is some particular function $f \in \text{CNL}$, but why not some other function g or h ?

Barring the sort of cleverness we’ll see in Plan IV, it seems that any answer we give could be expressed as “because f is better suited than the other functions in CNL to govern the universe.” This would imply some ordering BS that tells when one causal natural law is Better Suited than another to govern the universe. But then we could go up a level and ask: why is BS best suited to be the ordering that reflects functions’ suitability to govern the universe? Why not the ordering BS’, in which some other function g is maximal? We are in an infinite regress. At no point does any magic hand anoint the ordering that determines the ordering best suited to determine the ordering (...) best suited to determine the function that governs the universe. It looks as if the laws of nature must be arbitrary.

At this point the empiricists will pipe up again: “Look, with all this *why why why why why*, you’re regressing not merely to infinity, but to age three. Obviously explanations stop *somewhere*. It’s a component of maturity to accept that some brute facts are just *given*.”

We certainly respect this position’s forthrightness. And for all we know for sure, this sort of mature acceptance may turn out to be the healthiest attitude available. The problem is that we can’t quite shake our distaste for “brute givens.” Maybe it would be easier if we didn’t have examples of elegant, non-arbitrary truths, but we do have examples, and in abundance: *mathematics*. “ $2 + 2 = 4$ ” is a truth that nobly disdains all issues of justification or explanation. To understand it is to understand it to be *right*. We can’t help hoping that nature’s laws will somehow share this elegant quality.

Plan III to Deal With Arbitrariness: Find The “Least Arbitrary” Law

Some philosophers have in fact claimed to know a law with this elegant quality, or at least to have deduced its existence. Despite the infinite regress just described, they believe in a function that, in virtue of its *quasi*-mathematical elegance, of its *near*-logical-necessity, “really is” best suited to govern the universe. Leibniz is the archetype of this mindset. Here is his own remarkable argument:

In practical affairs one always follows the decision rule in accordance with which one ought to seek the maximum or the minimum: namely, one prefers the maximum effect at the minimum cost, so to speak. And in this [metaphysical] context, ... the receptivity or capacity of the world can be taken for the cost or the plot of ground on which the most pleasing building possible is to be built, and the variety of shapes therein corresponds to the pleasingness of the building and the number and elegance of the rooms. And the situation is like that in certain games, in which all places on the board are supposed to be filled in accordance with certain rules, where at the end, blocked by certain spaces, you will be forced to leave more places empty than you could have or wanted to, unless you used some trick. There is, however, a certain procedure through which one can most easily fill the board. ... And so, assuming that ... something is to pass from possibility to actuality, although nothing beyond this is determined, it follows that there would be as much as there possibly can be, given the capacity of time and space (that is, the capacity of the order of possible existence); in a word, it is just like tiles laid down so as to contain as many as possible in a given area.

From this we can already understand in a wondrous way how a certain Divine Mathematics or Metaphysical Mechanism is used in the very origination of things, and how the determination of a maximum finds a place.³

Our puzzle requires a slightly different “determination of a maximum”; we are seeking the maximal world-governing *function*, rather than directly

³Leibniz, “On the Ultimate Origination of Things,” p. 150-151 in *Philosophical Essays* (Hackett, 1989).




Figure 2: A “least arbitrary” natural law à la Leibniz

seeking the maximal *world* as Leibniz does.

In the context of our simple discrete-space-time illustrations, we nominate as the “maximizing” function something we’ll call Σ : the function that just *sums* the masses of a point’s immediate predecessors. For instance, if p comes immediately after three points in its past light cone having mass 0, 3, and 8 respectively, then Σ will give p mass 11. We argue that there is a nice scalar purity in summation; the sum of a set of numbers is in some sense “the most that can be got from them.” It seems less arbitrary than other laws we might choose.

It’s hard to deny that there is something pleasingly natural about the thought that such a maximizing law should be the real one, even if (as Leibniz admits elsewhere) this naturalness cannot quite amount to logical necessity. But the consensus is that this sort of view is unsatisfactory—even on purely philosophical grounds, to say nothing of empirical ones. Just which variable is the “natural” one to maximize? Why should there be a unique law or unique world that maximizes it (especially since worlds of infinite size seem possible)? Why, instead of our function Σ , did we not make a pitch for a multiplication function Π , or a function that sums *all* a point’s predecessors, not just the immediate ones? And finally, there is this deeper question: Why, given that we have broken free of the theology that constrained Leibniz’s era, should we posit in the first place a mystical power capable of making a choice (let alone one agreeing with us humans on what seems like the “natural” choice)?

Plan IV to Deal With Arbitrariness: The All-Possible-Worlds-Are-Equally-Real Thesis

The above criticisms of the Leibnizian approach motivate today’s clever trick for exempting natural laws from arbitrariness, the *all-possible-worlds-are-equally-real* thesis. Actually, people speak more often of the “anthropic prin-

ciple,” but insofar as it succeeds in eliminating arbitrariness, our term is a better one for it. Let us explain why. The anthropic principle (or one strong version of it) says: “I think, therefore the universe *must* be so arranged as to incorporate thinking beings.” If this is taken to explain why the universe is roughly as it is, then the logic is faulty: It is perfectly possible that a universe *without* thinking beings, obeying some particularly stultifying natural law $g \in \text{CNL}$, could have been the real one. No entities would have existed to wonder about their own existence, but so what? Why was God bound to actualize thinking beings? Maybe He would have been guilty of bad taste had He chosen not to, but He would not have done anything impossible.

An extra idea is needed to make the anthropic principle do the explanatory work it seems intended to do—namely, the idea that all possible worlds are equally real. Scientists know this idea as the many-worlds interpretation of quantum mechanics; in philosophical circles it is mainly associated with the late Princeton professor David Lewis, although something like it seems to have motivated Nietzsche’s theory of eternal recurrence a century earlier. With this new assumption, it becomes possible to explain why *our* world, the world *we* experience, must be roughly as it is: because none of the possible worlds that differ significantly from it embeds any *minds* who might inquire about it.

There is no need for us to delve too deeply into this theory and its variants, which most readers will themselves have mulled over already (it has an undeniable appeal). Let us simply acknowledge that it eliminates arbitrariness more fully than the other views we’ve canvassed, and is every bit as stable and internally consistent. Many of its counterintuitive consequences can be happily accepted once one’s paradigm has shifted. —Many, but not all.

The big problem for this view is the violence it does to our intuitions about time and freedom. Of course, nobody has ever explained time or freedom quite satisfactorily, but this view’s failure is especially bad, especially *immediate*. By handing us the ensemble of all possible worlds as a timelessly existing block, it practically begs us to ask the tough questions: Why do we, who are bits embedded in one of these static worlds, experience it as evolving through time? Why do we seem capable of choosing different sandwiches off our lunch menu if our future is fixed? Adherents of the theory will of course clamber over each other to field these questions. “Time is an illusion!” some will say, “Lunchtime doubly so!” Others will insist, “By *freedom* we must have

meant a particularly nifty form of determinism all along!”⁴ The enthusiasm that the all-worlds-are-equally-real thesis begets in its fans is remarkable. But we are among the majority that ultimately finds their answers too glib.

Choiceless Space-time: Our New Plan to Deal With Arbitrariness

Plan IV is unsatisfactory because it gets time and freedom wrong, but it does purge arbitrariness from the world more fully than Plan III did. We would now like to return to Plan III and suggest a different way to purge its arbitrariness, a way that is friendlier to our intuitions of time and freedom.

We begin by recalling that the natural law f was not the only arbitrariness in our puzzle. f determines which member of STATES should obtain at each space-time point, but we never explained why STATES, rather than some other set, should be the range of data that can obtain at points. Physicists, of course, define STATES according to the needs of their theories. If their theory characterizes points by their non-negative real-valued mass, then STATES = $\{m \in \mathbb{R} : m \geq 0\}$. If their theory characterizes points with a vector in some Hilbert space, then STATES is the set of all such vectors. But we are not interested, here in our armchairs, by the empirically-inspired *what* of STATES; we are interested in the *why*. Why should STATES be any particular set rather than another?

In essence, the answers that have been given to this question are the same ones that Plans I–IV applied to the choice of f . They are no more satisfactory in this context than they were in the other. So, calling up the core Bergsonian idea that we stationed at the back of our mind, we propose an entirely new answer. We dispense with the choice of a special set STATES, dispense with the choice of the function that picks out a member of STATES for each point, and instead *characterize each point just by the totality of possibilities there*. And more specifically, by the totality of *mathematical* possibilities there.

Postulate 1a: Each space-time point is fully characterized just by the totality of mathematical possibilities there.

It is worth reading Postulate 1a a few times, as it contains the core of our “arbitrary-choice-less” space-time theory. The most jarring thing about it is its last word, “there.” Traditional philosophers of mathematics will balk at it:

⁴Daniel Dennett’s *Freedom Evolves* (Viking, 2003) is an entertaining exposition of this point of view. And the “lunchtime” joke is of course Douglas Adams’s.

“Surely there is one unique and timeless totality of mathematical possibilities, and it is perversely otiose to talk of the totality of possibilities *at* this or that point in space-time!” If this objection were right—if the same structures were possible always and everywhere—then our theory would be vacuous. Every point in space-time would be the same. Nothing would happen.

Against this objection, our second puzzle will give us a reason to think that new mathematical structures *must* become possible over time. For now, however, let us just reiterate what we said in our foreword: that this scenario’s oddness does not make it formally contradictory, and that the notion of multiple “totalities of possibilities” has become fundamental to set theory, where they are referred to as different “universes” or “models of ZF” (the axioms of Zermelo-Fraenkel set theory). Because we will formalize our ideas in set theory, we will henceforth usually replace the terms “mathematical structure” and “mathematical possibility” with the term “set.”

Now, we will merely have replaced one arbitrariness with another until we justify each space-time point’s association with its *particular* collection of sets. Having allowed the answer to the question “What is possible?” to be just as time- and space-dependent as the answer to “What is real?”, we may seem to have saddled it with the same intractability that Plans I–IV failed to overcome. We would *like* to say, “There is nothing arbitrary about what holds at p , because what holds there is simply the totality of all possibilities, with no choice made among them.” But why is the totality of possibilities available at p not itself arbitrary?

We do have (the beginnings of) a good answer to this question:

Postulate 1b: At each point, precisely those sets are possible that must be possible, given the sets that were possible at past points. We call the resulting set-totality the synthesis of the prior points.

Postulate 1b should be understood (not yet as a rigorous mathematical statement but) as an application of the belief, nearly universal among people who think about such matters, that the class of all mathematically possible structures is *closed under definable operations*. If an irrational decimal number $x = 0.2224 \dots$ is possible, for instance, then the decimal number $2x = 0.4448 \dots$ is also possible. And this is true regardless of whether x itself is definable in an absolute sense, *i.e.* by some mathematical formula. Postulate 1b notes that this closure idea ought to apply to the *past*: whatever structures happen to have been mathematically possible in the past, they and

any other structures *definable from them* ought to be possible now. Moreover, no *other* structures should be possible, since they would be *unaccountable*, in a sense that we will make clearer in a moment.

Let us compare our new plan side-by-side with the Leibnizian Plan III, whose strategy it largely borrows. In both cases the state of each point is supposed to be a “self-evident,” “eminently reasonable,” “not-at-all-arbitrary” function of its past. Plan III (or our simple cartoon of it) declared that distinguished function to be the summation Σ , and asked, in a haughty voice, “What else could it be? Multiplication? Indeed!” But the haughty voice was not an argument and the Plan fizzled. To succeed, the strategy must be deployed towards a different question. The old question was, “Which of this fixed set of possibilities is *real*, given what was real in the past?”. Our new question is “What is mathematically *possible*, given what was mathematically possible in the past?”. And, modulo some ambiguities we will discuss later, this question has only one answer, at least if we grant that no possible mathematical structure can become *impossible* later.

A comparison between Plan III and our new theory is sketched in Figure 3. The elements of the left-hand sketch have been defined; on the right, we need to explain what $L_{\min}(x)$ means. It is what the phrase “everything that must be possible given that the set x is possible” cashes out to in mathematics. L is the symbol for the *constructible universe*, the barest structure obeying the standard Zermelo-Fraenkel set theory axioms. L is also called the *constructive hierarchy* due to the way it is defined: we start with the empty set at level 0, then take the set of all of level 0’s definable subsets to be level 1, and so on. Now, given a set x , $L(x)$ is defined the same way except that level 0 is “seeded” with x (more precisely, with the union of $\{x\}$ and x ’s transitive closure). In both cases we carry out this iteration of taking previous levels’ definable subsets not just to every finite stage, but to any ordinal stage, taking the union of the previously obtained sets at infinite limits. Familiar mathematical objects like real numbers are identified in a canonical way with certain sets that arise in this hierarchy. There is, however, some ambiguity as to what ordinals are possible, so we write $L_{\min}(x)$ to mean that this process is carried out only up to a stage that is minimal in some sense (essentially, to the first ordinal that is not itself constructible from x , but we will not define this rigorously here). With the technique of forcing that Paul Cohen invented in the 1960s, we can find “mutually generic” real numbers x and y that instantiate the situation shown in the right-hand picture.

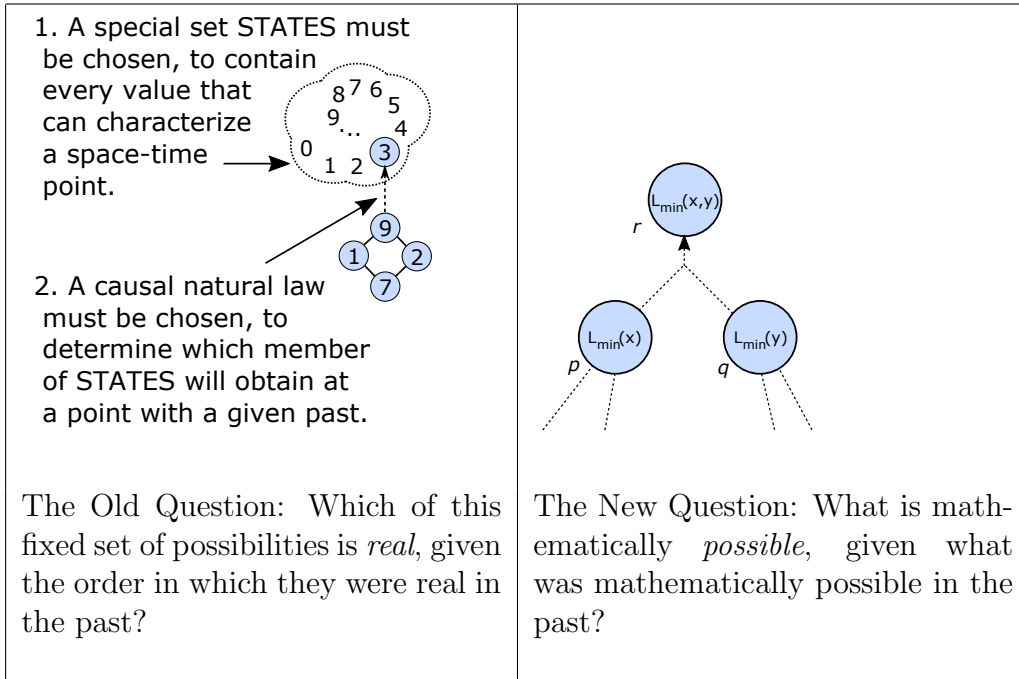


Figure 3: A comparison of the framework of classical space-time (left) with that of our new Bergsonian postulates (right)

The key difference between the pictures in Figure 3 is that the left-hand one involves three arbitrary choices, whereas the right-hand one involves only one. (The order type of the space-time ordering $<$ of points needs to be explained in both pictures and we will address this choice in a later section.) At point r in the right-hand picture, exactly those sets are possible that must be possible, given what was possible earlier. Specifically, what is possible is the *collection* of all previously possible sets, plus the minimum collection of sets needed to close it under definable operations. For instance, consider the real number z , defined by “interlacing” the digits of x and y . If $x = 0.33338 \dots$, and $y = 0.11112 \dots$, then $z = 0.3131313182 \dots$. (Just to be clear, the ellipses do not mean that these are *repeating* decimals!) z must be possible at r , but not at p or q since x and y are mutually unconstructible.

Of course, as we stressed at the outset, many of these structures, including this z , will be undefinable in an absolute sense. What guarantees their possibility is that they are definable *relative* to previously possible structures, so that if one *could* refer to the earlier ones, one could define the later ones

too.

It is worth stressing the notion of “accountability” which uniquely fixes the totality of mathematical possibilities at r . When we say that all possibilities at r must be accounted for, we mean essentially what Leibniz meant when he insisted that all events have a *sufficient reason*. Plan III itself *wanted* to avail itself of this notion; it held that a point’s mass could be “accounted for” if and only if it were the sum of immediately preceding masses. Again, this failed because no logical or philosophical (as opposed to aesthetic) principle made multiplication any less “accountable” than summation. The case is different for choiceless space-time, where we must account for possibility rather than reality. Here, if at point r a decimal number w were possible that were not constructible from the sets possible before r , then w really would be absolutely arbitrary, absolutely unaccountable — it would be the output of some cosmic random-number generator. There is a clear sense in which the interlaced decimal z has a sufficient reason for being possible at r , while the mysterious w would not.

Before going any deeper into the mathematics of Postulate 1, we ought to address its fundamental oddness: few readers will be comfortable with the idea that mathematical possibilities “arise” at all. Therefore we turn now to our second philosophical puzzle, which will lead us to a perspective from which this arising seems not only sane, but natural and necessary. From this perspective it would seem downright *paradoxical* if new sets *failed* to become possible.

Puzzle 2: The Set-Theory Paradoxes

The infamous paradoxes of set theory all begin by having us imagine a “set of all sets.” On some level it seems that this should be legitimate: we know what sets are, and now we are gathering them all up into a hypothetical mega-set. Note that we do more hereby than postulate the universe V of all sets, as we’ve already done; we say that V is itself a *set*, and is by virtue of this a member of the set of all sets, *i.e.* of itself.

The logic hits the fan when we start reasoning about this set of all sets—about the set of its non-self-containing members (the Russell paradox), or about the set of its members that are ordinals (the Burali-Forti paradox). Let’s recall in detail how the latter case leads to a contradiction. The definition of a (von Neumann) ordinal is a set x that is transitive (x contains all its members’ members, so that $z \in y \in x$ entails $z \in x$) and well-ordered by the

membership relation \in (of any pair of x -members, one contains the other, and there is no “descending infinite membership chain” $x \ni y_0 \ni y_1 \ni \dots$). Now suppose there is a set Ord of all ordinals. It would have to be transitive and well-ordered by the membership relation \in , and would thus be an ordinal itself. Therefore $Ord \in Ord$. But this yields an infinite descending chain of membership: $Ord \ni Ord \ni Ord \dots$. So Ord is not well-ordered by \in after all: contradiction.

Puzzle 2: How can we refer to a set of all sets (or all ordinals) without landing in mathematical contradiction?

The standard answer to this puzzle is well known: it is to decree *thou shalt not reason about the totality of all sets as though it were a set itself*. This “works,” but seems shamefully *ad hoc* to a certain philosophical mindset. To wit:

It is an embarrassment in set theory, as it is often understood, that an absolute distinction must be drawn between totalities such as the totality of ‘all ordinals’ or ‘all cardinals’ or ‘all sets’—the totalities which Cantor called ‘inconsistent manifolds’ and we call *proper classes*—on the one hand, and those totalities which form sets. For when we take the former totalities to be well-defined objects, then we must make this absolute distinction: the two kinds of objects must be treated quite differently. But why, if the totality of all sets has a well-defined extension, is it not a set in a more extensive totality?⁵

As with our first puzzle, today’s marketplace of ideas offers this puzzle little but glib, skeptical, and arbitrariness-embracing answers. The glib ones try to convince us *without* invoking the paradoxes directly that we were wrong to expect a universal set to exist. The arbitrariness-embracing ones insist that the paradoxes really do *explain* by themselves the universal set’s nonexistence. And the skeptics shrug off the issue by saying that mathematics is ultimately just a language game whose rules we can make up as we see fit.

Despite the many well-written, good-faith presentations of these answers, there are people like the quoted passage’s author who are not satisfied, who

⁵W. W. Tait, “Constructing Cardinals From Below,” p. 10, on his University of Chicago website at <http://home.uchicago.edu/~wtx/Foundations3.pdf>.

still think there ought to be a “more extensive totality” that could include a set of all sets. What Bergson’s core idea suggests is that such a totality could *become* possible, if the class of all sets *grew over time*. Consider: it takes *time* to think the thought “The set of all sets is a member of the set of all sets,” time over which new (bigger) ordinals could become possible, in which case the “set of all sets” referenced at the *start* of the thought would be *different*, and *smaller*, than the “set of all sets” referenced at the *end*. Rather than the paradoxical $V \in V$, we will have thought the innocuous proposition $V_1 \in V_2$. (Or $Ord_1 \in Ord_2$, in the Burali-Forti case.)

We will call any process that makes a new (higher) *ordinal* possible a *paradox escape*. At the moment of a paradox escape, it becomes possible to “step outside” all previously possible sets, and consider them *as* a totality; the set of all previously possible sets becomes possible.

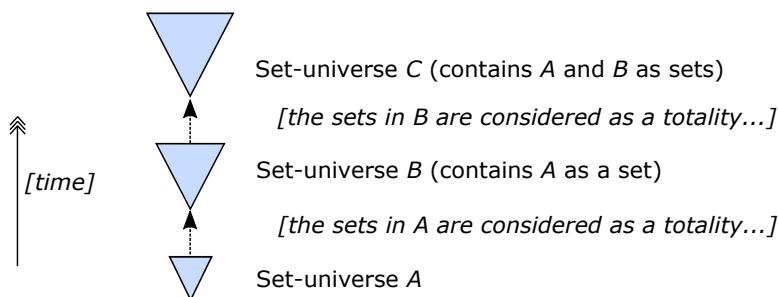


Figure 4: Escaping the paradoxes with a sequence of bigger and bigger set-totalities

This resolution to the set-theory paradoxes might already be celebrated as the most satisfying, were it not for the question: *Who* or *what* is responsible for the creation of new ordinals? We want to say that these ordinals are “automatically” made possible as soon as the totality of all ordinals is considered *as* a totality, but such use of the passive voice cannot ward off the question: *Who* is it that considers the totality of ordinals as a totality? Who is the synthesizing intellect? We hardly aspire to this role ourselves; Burali-Forti’s proof is not a magic spell that we chant to conjure noumenal entities. We might borrow the answer that Berkeley gave to explain the continuation of his own universe, and say that God is the synthesizing intellect. But a Berkeleyan God is the sort of arbitrary, unexplained demiurge that it has been our task to eliminate.

Our solution to this puzzle comes from the “choiceless space-time” model that grew out of our first puzzle. The solution is not to name a being who, by thinking of the totality of possible sets, pushes our world across the gaps in Figure 4. It is instead to eliminate those troublesome gaps by rearranging the set-totalities. If we put the set-totalities into a partial ordering that is not a discrete linear ordering, then each might be a *synthesis* of those prior to it, as per our Postulate 1. Such a synthesis can (depending on the sets it gathers together) also be a paradox escape. And as we already argued in our discussion of choiceless space-time, such a synthesis is philosophically innocuous. It simply establishes what *must* be mathematically possible, given what is *already* possible.

The right-hand picture of Figure 3 above can serve to illustrate a philosophically innocuous paradox escape. In our initial explanation of it we said that x and y , the real numbers that generate all the possibilities at the two earlier points, were “mutually generic.” With the most common kinds of generic real numbers, the point immediately later than these two — where just those sets constructible from $\{x, y\}$ are possible — will have no additional ordinals. Well, let us now suppose instead that x and y are the same kinds of generic real numbers, but not *mutually* generic. It turns out to be consistent that $\{x, y\}$ constructs an ordinal higher than any constructible from x or y alone. We thus have a mathematical green light to advance our second Postulate:

Postulate 2: There are points at which some ordinals are possible that were not possible at any previous point.

Strengthening Postulate 1 To Simplify Our Work

Before developing our ideas further, we propose a slightly stronger version of Postulate 1 that will simplify our work significantly.

Postulate 1a was that the mathematical possibilities at a given space-time point “fully characterize” it, that some model of set theory tells you everything there is to know about it (except maybe its relations with other space-time points). Postulate 1a⁺ will have us *identify* the point with its model of set theory. The collection of sets no longer merely “characterizes” or “gets associated with” the space-time point; it *is* the point. Postulate 1a⁺ rules out the existence of two distinct space-time points whereat exactly the

same sets are possible.

Postulate 1b was that whatever sets were possible in point p 's past, they and any other sets constructible from them (but no others) are possible at p . From this there follows a connection between the “earlier than” relation among points and the inclusion relation among their respective ZF models. Specifically, $p < q$ implies that the collection of sets possible at p is included among those possible at q , which we can write thanks to Postulate 1a⁺ as $p \subseteq q$. Postulate 1b⁺ is that the converse of this should also hold: $p \subseteq q$ should imply $p < q$. In fact we wish to *identify* the earlier-than relation $<$ among space-time points with the (strict) inclusion relation \subset among the ZF models associated (and henceforth identified) with them. The intuition here is that since p 's possibilities can all be seen to “contribute to” the possibilities at q , p ought to be considered *prior to* q in the causal ordering. (Postulate 1b⁺ precludes any p spacelike with respect to q from being associated with a *superset* of q 's ZF model.)

Both of these strengthenings imply, roughly speaking, that you never have exactly the same unconstructible sets popping up independently at two spacelike points in space-time. We don't claim to have an unshakeable metaphysical conviction that this never happens; such a claim would probably have to be classed with Leibniz's, that no two snowflakes can ever be alike. From where we sit now, we can imagine it proving useful or necessary for the development of our theory to renounce one or both of these strengthenings. But for now we observe that they make our work far simpler: when we're speaking of space-time points, p becomes a *synonym* for the sets that are possible at p , and the temporal ordering $<$ becomes a *synonym* for \subset .

Puzzle 3: The Arbitrariness of Space-time Structure

The Postulates presented so far address only two of the three kinds of arbitrariness we identified in the first section, that of the range of possible states for space-time points, and that of the natural law which selects one for each point. There remains the arbitrariness of space-time structure. We phrase this puzzle in terms of the “synthesis” that we defined in the first postulates:

Puzzle 3. What determines which collections of space-time points get “synthesized” into an immediate successor point?

If space-time had only finitely many points, we might appeal to a “maximality principle” like the Leibnizian one cited earlier, and say that *every* collection of space-time points get synthesized. But our first Postulate entails that every point (except for the earliest one, if there is one) must have infinitely many predecessor points. The argument is simple. First, if a point has no predecessors, it must be the *constructible* set-totality (taken up to some ordinal stage), for there is nothing “previously possible” in terms of which an unconstructible set could be defined. Second, if some point p had finitely many predecessors, then some q among them (or p itself) would have as its *sole* predecessor this constructible set-totality. But then the “new” unconstructible sets in q (or p) would again be unaccountable in terms of what was possible beforehand.

The problem with appealing to a maximality principle when there are infinitely many points—with asserting that “every sub-collection of the collection of space-time points should synthesize a new point”—is that there is not, in general, any fact of the matter about what sub-collections an infinite collection has. It would depend on what “background” or “outer” ZF model we took ourselves to be working in—and the whole point of our endeavor is to deny that any such thing is given in advance.

A maximality principle would also be in tension with our arbitrariness-purging goals; there lurks in it an unexplained “demiurge” of the kind we tried to avoid when discussing paradox-escapes. Recall that the issue was how possibilities at earlier points lead to new possibilities at later points, and that we hoped never to need any mysterious force to “act” to obtain one from the other. Our worry here is that whenever *unrelated* set-totalities are brought together and synthesized, there must be some force that reaches out and brings them together.

Is there an alternative principle that could banish these worries from our system? We suggest a *vanishingly-little-work principle*. Intuitively, it should allow a collection X of points to synthesize into a successor point p just if X ’s points are “already infinitely close to attaining all the possibilities in p ” and require “vanishingly little intelligence” to synthesize. We will wait until the next section to formalize this, but it will require that for any pair $q, q' \in X$, there must be $r \in X$ such that $q \subset r$ and $q' \subset r$. Thus q and q' have “already been brought together” in X ; the synthesis of X into p will not be yoking together unrelated models.

Postulate 3. Letting p be any space-time point (and hence, by Postulate

1, a set-totally) and considering any smaller set-totally $q \subseteq p$, q will itself be a space-time point just if “vanishingly little work” is required to synthesize it from some subcollection of p ’s predecessors.

Puzzle 4: How Can We Account For What Is Mathematically Possible?

A perennial puzzle in the foundations of mathematics concerns statements like the continuum hypothesis. Roughly speaking, this hypothesis says that there are no sets “bigger than” the set of integers and “smaller than” the set of real numbers. Such statements are puzzling because, as Gödel and Cohen showed, they cannot be proved or refuted by the standard axioms of mathematics, nor by any axiom that we are likely to admit in the future on classical mathematical grounds. There are two main schools of thought regarding this puzzle. One holds that these statements really are true or false, that the hypothesized sets are or are not possible, in a metaphysical sense of “possible” that is neither the same as physical possibility nor amenable to investigation by classical mathematical methods. The other school holds that there is no truth at all about these statements; they might be shown to lead or not to lead to contradictions, but any talk about whether the non-contradictory ones are “really” true is ethereal nonsense.

Our approach to this puzzle should be easy to guess from what we have said already. It is closer to the first school of thought, agreeing that the various hypothesized sets really are or really aren’t possible. But it overcomes in a novel way the common objection to this position, that their possibility or impossibility cannot be explained. Our Bergsonian view is that every possible set has been *made possible* through the synthesis of less complex sets. There is a web of explanations for the possibility of sets leading back to an earliest point, at which only the constructible sets are possible. *Space-time is this web.*

Some care is necessary if this approach is to account for *all* mathematical possibilities. The preceding postulates may allow cases in which there is a set x and a point p , such that x appears at all points later than p , but not at p or anywhere else. In such cases it is mysterious how x became possible. To exclude such cases we need a “well-foundedness” requirement on the web of explanations. We need to know that *every* mathematical possibility (ex-

cept for the constructible ones, which are possible always and everywhere) is explained in terms of prior possibilities. The following postulate ensures this:

Postulate 4. If a set is possible at some point p , it is possible at some point q prior to or identical with p , such that it is not possible at any point prior to q .

Translating The Bergsonian Postulates Into Axioms

The considerations driving our postulates have been abstract rather than empirical. It would be to a world's credit, we dare say, if it happened to obey our postulates, but we have not yet asked whether our world does obey them, or even whether any structure *could* obey them. We turn to these questions now. But we cannot pose them clearly until we translate our postulates into purely mathematical statements. What is especially in need of precision is the idea that a space-time point is the “synthesis” of previous points, *i.e.* the closure of their sets under all definable operations.

We have already said that standard axioms of set theory capture the notion of closure under all definable operations. Indeed, most of these axioms explicitly take the form “For every set x (and y), the set derived from x (and y) in such-and-such a way exists.” Two of the axioms, however, involve the kind of the inexplicability we have been trying to avoid; these are the “non-constructive” axioms of Power Set and Choice. Like most of the other axioms, these two assert that some set based on x exists—in the case of Power Set, the set of all x 's subsets, and in the case of one popular form of Choice, a set containing one member chosen from each of x 's members—but they provide no way to specify what that set's members are. We thus disown these axioms, and rather than require each point to be a model of ZFC (the Zermelo-Fraenkel set-theory axioms plus the axiom of Choice), we require it only to be a model of ZF^- (the Zermelo-Fraenkel axioms without the Power Set axiom).

We will need, however, to restrict which ZF^- models are candidates to be space-time points. The reason is that Postulate 1 should make each point something like “the ZF^- model generated by the collection of prior ZF^-

models,” and a general theory of what that would mean turns out to be rather gnarly. It would be easier if we could apply the well-known idea of *relative constructibility* to the collection of prior ZF^- models. The only thing keeping us from using it straightforwardly is that it applies to sets, rather than to classes or to families of models. We can avoid this obstacle if we stipulate that all the ZF^- models we deal with are “set-generated” in the following sense:

Definition. For x any set, the ZF^- model generated by x is $L_{\min}(x)$, meaning the constructive hierarchy seeded with $\{x\} \cup \text{TrCl}(x)$ and continued up until the least ordinal stage at which the hierarchy satisfies ZF^- . (Here $\text{TrCl}(x)$ means x ’s transitive closure, and the “min” subscript is meant to denote the least ordinal just described.) A ZF^- model is *set-generated* if some x generates it.

Now, we want each non-least point p to be the synthesis of infinitely many previous points, and we want the “generator” idea to let us translate this as: p is the ZF^- model generated by the union of those previous points’ generators. The hitch is that the union of *all* their generators is a proper class rather than a set. (For if x generates p then so does $\{x, \alpha\}$ for any ordinal $\alpha \in p$.) We can try to avoid that hitch by taking just one generator of minimal rank from each previous point, where “rank” is a standard set-theoretic property. But in doing so we have to be careful not to invoke an uncountable choice function on all the points’ generators.

We can, however, take some particular formula to define the *canonical* generator of a set-generated ZF^- model. Most simply, there is the formula defining the set of *all* a model’s generators of minimal rank. We will write CG to denote the “canonical generator” function defined by this formula.

We can state our axioms now, with the caveat that Axiom 3 will need editing to avoid conflict with other axioms:

Bergsonian axioms that a structure B should satisfy:

Axiom 1a: $p \in B \Rightarrow p$ is a set-generated ZF^- model.

Axiom 1b: $p \in B \Rightarrow p = L_{\min}(\{\text{CG}(q) : q \in B \text{ and } q \subset p\})$.

Axiom 2: There are points $p, q \in B$ such that $p \in q$.

Axiom 3 (draft): For all $x \in \bigcup B$, $L_{\min}(x) \in B$ if and only if it equals $L_{\min}(\{\text{CG}(N) : N \in H\})$ for some directed $H \subseteq B$, downward-closed in B ,

and unbounded in the family of all $L_{\min}(x)$'s set-generated proper submodels.

Axiom 4: $x \in p \in B \Rightarrow B$ contains a minimal member $q \subseteq p$ containing x .

We feel confident that each Axiom here is the simplest and most natural one for resolving the puzzle it was intended for. Sadly, our draft of the third Axiom is in tension with the second, and it is unlikely that all the Axioms can be simultaneously satisfied as written. Resolving this conflict involves more mathematical complexity than we wish to present in this paper; suffice it to say that third Axiom can be edited so as to stop conflicting with the second, while continuing to capture the intuitive meaning of the third Postulate.

We call a structure satisfying these axioms a *self-constructing family of ZF⁻ models*. Our mathematical task is to find one — hopefully, one that bears some structural resemblance to our universe.

Now, the obvious tool for satisfying these axioms is the technique of *forcing* invented by Paul Cohen in the 1960s. This was the first technique to establish the consistency of unconstructible sets, and it remains by far the most popular tool for defining nested models of set theory. The ZF models it yields, called “generic extensions,” can satisfy a wide range of desired properties.

A generic extension of L (the minimal ZFC model) is a model of form $L(G)$, where $G \notin L$ is an ultrafilter on some particular boolean algebra B . The properties of this model are determined by the particular B that is used. Moreover, every ZFC model N that is an inner model of $L(G)$ (that is, $L \subseteq N \subseteq L(G)$) has form $L(G \cap C)$ for some subalgebra $C \subseteq B$. Thus what we are seeking is a suitable system of nested boolean algebras, some or all of which will correspond to space-time points; when point p is $\leq q$, any algebra corresponding to p will be a subalgebra of an algebra corresponding to q .

Axiom 2 requires a boolean algebra that constructs (or, in the lingo, “collapses”) new ordinals when it used for forcing. Now, if spacetime points are densely ordered (i.e., if $p > q$ implies that some point r satisfies $p > r > q$) then not every spacetime point p can correspond to an algebra constructing an ordinal $\alpha(p)$ that is greater than any constructed by p 's predecessors; otherwise there is an infinite descending chain of ordinals $\alpha(p) > \alpha(r_1) > \alpha(r_2) > \alpha(r_3) \dots$, which is impossible. Thus if our nested system of boolean algebras yields a densely ordered set of space-time points, as we expect, those points that collapse higher ordinals will be “sprinkled” intermittently among

those that don't. (This is suggestive of the discontinuous measurement events at the core of the quantum theory.) But we will focus now on points whose algebras do not collapse ordinals.

The simplest non-ordinal-collapsing boolean algebras are called Cohen algebras and measure algebras. Alas, neither one yields models of the Bergsonian axioms. We have concocted a much more complicated algebra that satisfies weakened versions of the Axioms — specifically, it satisfies the Axioms when they are restricted to the new countable sets (such as real numbers) that are generated. But this algebra generates new uncountable sets in a way that violates the full versions, and we conjecture that this issue would rule out any algebra built along similar lines.

Since classical examples from set theory don't work, and since we hope the Axioms apply to the real world, we turn to current theories of physics to see if they supply any boolean algebras that would be good candidates. It turns out that algebraic quantum field theory (AQFT) is, at least superficially, quite like our own set-up, associating to nested regions of spacetime correspondingly nested structures called von Neumann algebras. The most natural way to derive boolean algebras from von Neumann algebras is to define topologies on them, and then consider the “regular open algebras” of those topologies. The trick is in defining the right topology; none of the most commonly used topologies (the norm topology, the weak operator topology, etc.) works for our purposes. A chief aim of this paper is to convince experts in this technically difficult field to join in the effort to define a topology that does work.

When the young Henri Bergson decided to pursue his metaphysical ideas, he foreclosed a future in mathematics, in which he had been a brilliant and indeed nationally recognized student. This prompted his professor's famous quip: “You could have been a mathematician; and you will be a mere philosopher.” It seems to us now that Bergson could not have chosen otherwise. The culture was ready for his intuitive rejection of the time-is-just-like-a-spatial-dimension thesis, whereas it would take another century's worth of geniuses to develop mathematics that could truly put Bergson's ideas to a formal test. We hope we have made a good case that the time for this test is now.

Appendix: Consequences For Mathematics

Since the Bergsonian axioms are meant to be a joint solution to deep problems of physics and mathematics, it is worth looking at it more closely from the latter side, to see where it fits into what is called the philosophy of mathematics. Two of the main questions that occupy this field's thinkers are "Is the subject matter of mathematics a realm of objective mathematical objects?" and "Can what is mathematically true change over time?". This gives a matrix of four possible answer pairs, which we will review with great brevity.

"No"/"No" is best represented by classic logicism. This is the belief (associated mainly with Gottlob Frege, Bertrand Russell, and Early Wittgenstein) that mathematics is "just logic," and that logic itself boils down to formal rules about when the very structure of a proposition makes it true or false, irrespective of what its terms refer to. Thus mathematical truth is a matter of rules about propositions, not of mathematical objects; logicians have by and large understood these rules to be immutable.

"Yes"/"No" is "platonism," a view which gained popularity after Gödel showed that no finite set of logical rules can yield all mathematical truths. A platonist takes the propositions "there are infinitely many primes" or "there is an unconstructible set" to be true or false in much the same way that the proposition "there is a lion at the zoo" is true or false (which is to say *objectively*, and *not as a matter of mere logic*), except that mathematical propositions are not permitted to become true or false over time.

"No"/"Yes" are the answers of the so-called intuitionist school, which emphasizes mathematics as the product of human minds. When, in the foreword, we said that "hardly" anyone has doubted mathematics' timelessness, it was this school that made the qualification necessary. It accepts that a proof about a newly defined structure may establish a genuinely new truth. But such a truth is not really objective because, according to them, there is no pre-existing realm of mathematical objects; mathematicians essentially deduce truths about figments of their imaginations.

These three pairs of answers capture the three main schools of mathematical philosophy, which for a century or so have been in stalemate, or, if you prefer, equilibrium — an equilibrium just dynamic enough to support a handful of scholarly journals.

By answering "yes/yes" we propose a truly new approach to the foundational problems of mathematics. Yes, there really is a realm of mathematical

objects (or possibilities); and yes, it does grow over time; moreover, spacetime is the map of that growth.

Of course, it may happen that our work here just turns a stalemate of three not-quite-satisfactory philosophies into a stalemate of four. The answer depends on how the Bergsonian axioms fare as an approach to physics, and that is why we have spent, and will continue to spend, most of our time on that aspect of it.