Computational Complexity of Cosmology in String Theory

Michael R. Douglas

[1] Simons Center / Stony Brook University

Imperial, April 28, 2017

Abstract

To appear, with Frederik Denef, Brian Greene and Claire Zukowski.

String landscape

In a quantum theory of gravity, the vacuum energy density is an observable.
There is good evidence for inflation in early cosmology, which was driven by a positive vacuum energy at that time.
There is also good evidence that the expansion of the universe is presently accelerating, so that there is positive vacuum energy today, of order $\Lambda_{now} \sim 10^{-120} M_{Pl}^4$.
Conversely, there is no known theory in which zero energy is a distinguished value, and the uncertainty principle implies that this could only be the case when considering large regions of space-time.
Any mechanism to dynamically set the vacuum energy near zero would be non-local.

On the other hand, the anthropic argument in a multiverse looks like a consistent resolution of the problem. This requires $N \sim M^4/\Lambda_{now}$ vacua with uniformly distributed cosmological constants, where M is the natural scale of vacuum energy.

String compactification leads to such vacua, which differ in the topology and geometry of the extra dimensions, branes, fluxes and other structures.

The study of string compactification is still in its early days. Let us list our working assumptions, none of which are indisputable.

Shape of the string landscape

- Long-lived metastable de Sitter vacua exist.
- To be long-lived, supersymmetry must be broken well below the compactification scale.
- A very rough parameterization of vacua is to give the Betti numbers of the compactification manifold M and one or two integral homology classes (wrapped branes, fluxes) satisfying constraints.
- Nearby vacua (in these parameters) are connected by tunneling processes, and all of the vacua are connected by chains of these processes.
- If we place an lower bound on the KK scale, there are a finite number of vacua.
- There are some particularly “simple” M which lead to the vacua which appear as initial conditions in string theory – perhaps T^6, perhaps CY3 with small Betti numbers, or perhaps non-geometric or quantum geometric constructions.

Bousso-Polchinski model

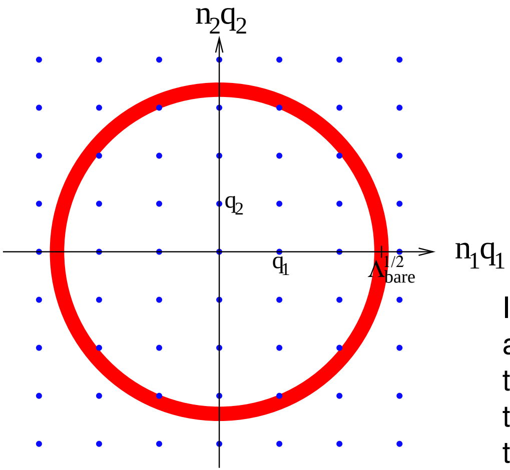

$$E = \sum_{i,j} Q_{i,j} N^i N^j - \Lambda_0$$

In this figure (from Bousso and Polchinski 2000), the lattice points represent quantized values of the flux, while the red circle is the region with $\Lambda \sim 0$.

Denef-Douglas

Suppose we actually had the list of string vacua and the ability to compute observables for each one – cosmological constant, gauge group, matter content, hidden sectors, etc. How hard would it be to find vacua which reproduce the observations to date – Λ_{now} and the Standard Model?

Both in the BP model and in more accurate models of the string landscape, the problem of tuning Λ_{now} is generically hard. If we choose a generic region of the landscape, then DD argued that the problem is NP hard – even if we grant that the cosmological constant can be computed in time polynomial in the size S of the description of the vacuum (say the sum of the Betti numbers), one cannot find a vacuum with $|\Lambda - \Lambda_{now}| < \epsilon$ in time much less than $T \sim 1/\epsilon$.

This is a very general property of landscapes with complicated potentials (terms of different signs, frustration). One does not even need a complicated potential – DD proved it in the BP model (assuming that S could be arbitrarily large).

Denef-Douglas

Suppose we actually had the list of string vacua and the ability to compute observables for each one – cosmological constant, gauge group, matter content, hidden sectors, etc. How hard would it be to find vacua which reproduce the observations to date – Λ_{now} and the Standard Model?

Both in the BP model and in more accurate models of the string landscape, the problem of tuning Λ_{now} is generically hard. If we choose a generic region of the landscape, then DD argued that the problem is NP hard – even if we grant that the cosmological constant can be computed in time polynomial in the size S of the description of the vacuum (say the sum of the Betti numbers), one cannot find a vacuum with $|\Lambda - \Lambda_{now}| < \epsilon$ in time much less than $T \sim 1/\epsilon$.

This is a very general property of landscapes with complicated potentials (terms of different signs, frustration). One does not even need a complicated potential – DD proved it in the BP model (assuming that S could be arbitrarily large).

Denef-Douglas

Suppose we actually had the list of string vacua and the ability to compute observables for each one – cosmological constant, gauge group, matter content, hidden sectors, etc. How hard would it be to find vacua which reproduce the observations to date – Λ_{now} and the Standard Model?

Both in the BP model and in more accurate models of the string landscape, the problem of tuning Λ_{now} is generically hard. If we choose a generic region of the landscape, then DD argued that the problem is NP hard – even if we grant that the cosmological constant can be computed in time polynomial in the size S of the description of the vacuum (say the sum of the Betti numbers), one cannot find a vacuum with $|\Lambda - \Lambda_{now}| < \epsilon$ in time much less than $T \sim 1/\epsilon$.

This is a very general property of landscapes with complicated potentials (terms of different signs, frustration). One does not even need a complicated potential – DD proved it in the BP model (assuming that S could be arbitrarily large).

Computation and cosmology

Given that it is hard to find vacua with small c.c., is there any paradox in the fact that early cosmology found our vacuum? Is this a problem for string cosmology?

To sharpen the question, let us ask: suppose we had a quantum supercomputer which can simulate the fundamental laws of physics – which are, of course, string/M theory!

Just for definiteness, let us postulate that this computer can simulate the entire history of the observable universe since the Big Bang, in a second of our subjective time.

But, before doing so, it has to find a compactification which reproduces our four-dimensional laws, or perhaps any "viable" set of laws (definitely small c.c.). And it has to do this by searching through the possibilities (we will be more precise about this shortly).

How long will it take to find a viable vacuum? Will it take of order a "second," or much less time, or much more time?

Computation and cosmology

Given that it is hard to find vacua with small c.c., is there any paradox in the fact that early cosmology found our vacuum? Is this a problem for string cosmology?

To sharpen the question, let us ask: suppose we had a quantum supercomputer which can simulate the fundamental laws of physics – which are, of course, string/M theory!

Just for definiteness, let us postulate that this computer can simulate the entire history of the observable universe since the Big Bang, in a second of our subjective time.

But, before doing so, it has to find a compactification which reproduces our four-dimensional laws, or perhaps any "viable" set of laws (definitely small c.c.). And it has to do this by searching through the possibilities (we will be more precise about this shortly).

How long will it take to find a viable vacuum? Will it take of order a "second," or much less time, or much more time?

Computation and cosmology

Given that it is hard to find vacua with small c.c., is there any paradox in the fact that early cosmology found our vacuum? Is this a problem for string cosmology?

To sharpen the question, let us ask: suppose we had a quantum supercomputer which can simulate the fundamental laws of physics – which are, of course, string/M theory!

Just for definiteness, let us postulate that this computer can simulate the entire history of the observable universe since the Big Bang, in a second of our subjective time.

But, before doing so, it has to find a compactification which reproduces our four-dimensional laws, or perhaps any "viable" set of laws (definitely small c.c.). And it has to do this by searching through the possibilities (we will be more precise about this shortly).

How long will it take to find a viable vacuum? Will it take of order a "second," or much less time, or much more time?

Computation and cosmology

Given that it is hard to find vacua with small c.c., is there any paradox in the fact that early cosmology found our vacuum? Is this a problem for string cosmology?

To sharpen the question, let us ask: suppose we had a quantum supercomputer which can simulate the fundamental laws of physics – which are, of course, string/M theory!

Just for definiteness, let us postulate that this computer can simulate the entire history of the observable universe since the Big Bang, in a second of our subjective time.

But, before doing so, it has to find a compactification which reproduces our four-dimensional laws, or perhaps any "viable" set of laws (definitely small c.c.). And it has to do this by searching through the possibilities (we will be more precise about this shortly).

How long will it take to find a viable vacuum? Will it take of order a "second," or much less time, or much more time?

Computation and cosmology

Given that it is hard to find vacua with small c.c., is there any paradox in the fact that early cosmology found our vacuum? Is this a problem for string cosmology?

To sharpen the question, let us ask: suppose we had a quantum supercomputer which can simulate the fundamental laws of physics – which are, of course, string/M theory!

Just for definiteness, let us postulate that this computer can simulate the entire history of the observable universe since the Big Bang, in a second of our subjective time.

But, before doing so, it has to find a compactification which reproduces our four-dimensional laws, or perhaps any "viable" set of laws (definitely small c.c.). And it has to do this by searching through the possibilities (we will be more precise about this shortly).

How long will it take to find a viable vacuum? Will it take of order a "second," or much less time, or much more time?

Computational complexity of cosmology

In our joint work (DDGZ), we make this question precise, and show that – if the search for a viable vacuum proceeds by simulating the dynamics of the most popular framework for the study of quantum cosmology, namely eternal inflation and an equilibrium measure factor – the search will take **far, far** longer than the simulation of our universe.

Furthermore we argue that there are other search algorithms – even fairly random ones which do not exploit the detailed structure of string theory – which could find a viable vacuum more quickly, in time polynomial in the number of quantum gate operations $\mathcal{C}_{obs}$ which it would take to simulate the observable universe.

And, if the supercomputer is allowed to take advantage of the details of string theory, there may be "engineered" corners of the landscape in which the search is very efficient. According to DD we would need at least $M^4/\Lambda_{now} \times \mathcal{C}_{avg}$ operations to find a vacuum with small c.c. (where $\mathcal{C}_{avg}$ is the average time to compute the c.c. of a universe), but in these engineered corners it could be far quicker.

Computational complexity of cosmology

In our joint work (DDGZ), we make this question precise, and show that – if the search for a viable vacuum proceeds by simulating the dynamics of the most popular framework for the study of quantum cosmology, namely eternal inflation and an equilibrium measure factor – the search will take **far, far** longer than the simulation of our universe. Furthermore we argue that there are other search algorithms – even fairly random ones which do not exploit the detailed structure of string theory – which could find a viable vacuum more quickly, in time polynomial in the number of quantum gate operations $\mathcal{C}_{obs}$ which it would take to simulate the observable universe.

And, if the supercomputer is allowed to take advantage of the details of string theory, there may be "engineered" corners of the landscape in which the search is very efficient. According to DD we would need at least $M^4/\Lambda_{now} \times \mathcal{C}_{avg}$ operations to find a vacuum with small c.c. (where $\mathcal{C}_{avg}$ is the average time to compute the c.c. of a universe), but in these engineered corners it could be far quicker.

Computational complexity of cosmology

In our joint work (DDGZ), we make this question precise, and show that – if the search for a viable vacuum proceeds by simulating the dynamics of the most popular framework for the study of quantum cosmology, namely eternal inflation and an equilibrium measure factor – the search will take **far, far** longer than the simulation of our universe.
Furthermore we argue that there are other search algorithms – even fairly random ones which do not exploit the detailed structure of string theory – which could find a viable vacuum more quickly, in time polynomial in the number of quantum gate operations $\mathcal{C}_{obs}$ which it would take to simulate the observable universe.
And, if the supercomputer is allowed to take advantage of the details of string theory, there may be "engineered" corners of the landscape in which the search is very efficient. According to DD we would need at least $M^4/\Lambda_{now} \times \mathcal{C}_{avg}$ operations to find a vacuum with small c.c. (where $\mathcal{C}_{avg}$ is the average time to compute the c.c. of a universe), but in these engineered corners it could be far quicker.

The cosmological search

How can we define this search process? We want the computer to be able to exercise choice in where to search – but we want this to be compatible with all the laws of physics, so that an observer within a given universe can not directly detect any sign that there was a search. The concept might still lead to testable consequences – say that certain types of universe which were disfavored in another measure factor, are favored to be found by search.

To begin, let us state our idea in semiclassical quantum gravity. Thus there is a state of the multiverse – namely a geometry and field values on an equal-time surface – which evolves with time, typically by Einstein's equations, but occasionally by tunnelling events. Where is there any choice?

In canonical GR, the equal-time surfaces are a choice, defined by lapse and shift functions. We can give our computer the ability to evolve the equal-time surface in a controlled way, choosing to advance it in some regions of space, and not advancing it in others.

The cosmological search

How can we define this search process? We want the computer to be able to exercise choice in where to search – but we want this to be compatible with all the laws of physics, so that an observer within a given universe can not directly detect any sign that there was a search. The concept might still lead to testable consequences – say that certain types of universe which were disfavored in another measure factor, are favored to be found by search.

To begin, let us state our idea in semiclassical quantum gravity. Thus there is a state of the multiverse – namely a geometry and field values on an equal-time surface – which evolves with time, typically by Einstein's equations, but occasionally by tunnelling events. Where is there any choice?

In canonical GR, the equal-time surfaces are a choice, defined by lapse and shift functions. We can give our computer the ability to evolve the equal-time surface in a controlled way, choosing to advance it in some regions of space, and not advancing it in others.

The cosmological search

How can we define this search process? We want the computer to be able to exercise choice in where to search – but we want this to be compatible with all the laws of physics, so that an observer within a given universe can not directly detect any sign that there was a search.

The concept might still lead to testable consequences – say that certain types of universe which were disfavored in another measure factor, are favored to be found by search.

To begin, let us state our idea in semiclassical quantum gravity. Thus there is a state of the multiverse – namely a geometry and field values on an equal-time surface – which evolves with time, typically by Einstein's equations, but occasionally by tunnelling events. Where is there any choice?

In canonical GR, the equal-time surfaces are a choice, defined by lapse and shift functions. We can give our computer the ability to evolve the equal-time surface in a controlled way, choosing to advance it in some regions of space, and not advancing it in others.

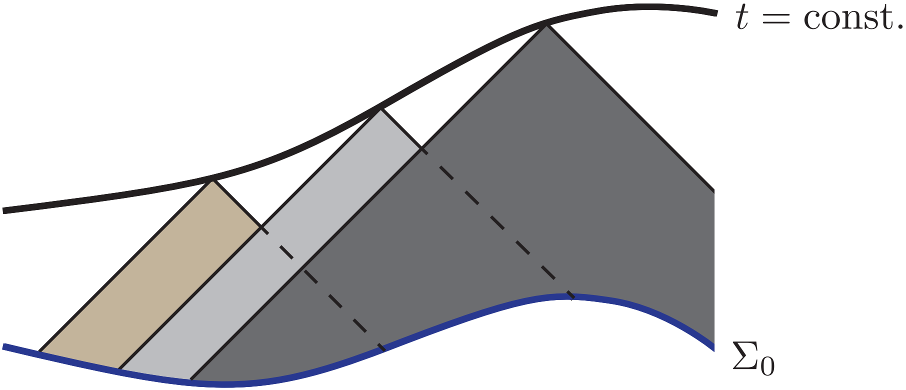
$t = \text{const.}$
Σ_0

Thus we give the computer two basic abilities, besides those of general-purpose computation.

One is to do measurements on and to the past of the equal-time surface. One can discuss what is allowed, but for simplicity let us imagine that each time a tunneling between vacua occurs, a fixed finite number of operations $\mathcal{C}_{meas}$ suffices to notice this and determine the parameters of the vacuum (topology, fluxes, etc.). To measure the cosmological constant to accuracy ϵ, one must be able to observe a region with space-time volume $1/\epsilon$.

The other operation is to move the equal-time surface forward, thus simulating the laws of physics in some chosen subregion of space-time. This has a computational cost $\mathcal{C}$ which depends on properties of the region being simulated.

The idea is that all of the multiverse potentially exists, but only a tiny fraction of it will be simulated. Now one of the main philosophical criticisms of the multiverse is the extravagance of what is being postulated, compared to what we actually observe. Our proposal answers this criticism.

Thus we give the computer two basic abilities, besides those of general-purpose computation.

One is to do measurements on and to the past of the equal-time surface. One can discuss what is allowed, but for simplicity let us imagine that each time a tunneling between vacua occurs, a fixed finite number of operations $\mathcal{C}_{meas}$ suffices to notice this and determine the parameters of the vacuum (topology, fluxes, etc.). To measure the cosmological constant to accuracy ϵ, one must be able to observe a region with space-time volume $1/\epsilon$.

The other operation is to move the equal-time surface forward, thus simulating the laws of physics in some chosen subregion of space-time. This has a computational cost $\mathcal{C}$ which depends on properties of the region being simulated.

The idea is that all of the multiverse potentially exists, but only a tiny fraction of it will be simulated. Now one of the main philosophical criticisms of the multiverse is the extravagance of what is being postulated, compared to what we actually observe. Our proposal answers this criticism.

Thus we give the computer two basic abilities, besides those of general-purpose computation.

One is to do measurements on and to the past of the equal-time surface. One can discuss what is allowed, but for simplicity let us imagine that each time a tunneling between vacua occurs, a fixed finite number of operations $\mathcal{C}_{meas}$ suffices to notice this and determine the parameters of the vacuum (topology, fluxes, etc.). To measure the cosmological constant to accuracy ϵ, one must be able to observe a region with space-time volume $1/\epsilon$.

The other operation is to move the equal-time surface forward, thus simulating the laws of physics in some chosen subregion of space-time. This has a computational cost $\mathcal{C}$ which depends on properties of the region being simulated.

The idea is that all of the multiverse potentially exists, but only a tiny fraction of it will be simulated. Now one of the main philosophical criticisms of the multiverse is the extravagance of what is being postulated, compared to what we actually observe. Our proposal answers this criticism.

Thus we give the computer two basic abilities, besides those of general-purpose computation.

One is to do measurements on and to the past of the equal-time surface. One can discuss what is allowed, but for simplicity let us imagine that each time a tunneling between vacua occurs, a fixed finite number of operations $\mathcal{C}_{meas}$ suffices to notice this and determine the parameters of the vacuum (topology, fluxes, etc.). To measure the cosmological constant to accuracy ϵ, one must be able to observe a region with space-time volume $1/\epsilon$.

The other operation is to move the equal-time surface forward, thus simulating the laws of physics in some chosen subregion of space-time. This has a computational cost $\mathcal{C}$ which depends on properties of the region being simulated.

The idea is that all of the multiverse potentially exists, but only a tiny fraction of it will be simulated. Now one of the main philosophical criticisms of the multiverse is the extravagance of what is being postulated, compared to what we actually observe. Our proposal answers this criticism.

Let us use these ideas to define a measure factor. To start, we choose a large but finite three-manifold as our equal-time surface, and some vacuum as the initial condition. Thus this definition is only well motivated in a theory in which some vacua are "simple" and are preferred candidates for the initial condition. This seems likely in string theory.

The computer then has a rule which, given its observations and calculations, specifies how to evolve the equal-time surface forward. We grant that its observations are powerful enough to decide whether a vacuum is viable – as we discuss later this is probably **not** testing for the existence of observers, but rather small c.c., the existence of structure and chemistry, and other criteria which are necessary but perhaps not sufficient to get observers. The computer then stops with the **first** viable vacuum it finds.

Because this is semiclassical gravity and tunnelling events are probabilistic, this is a probabilistic computation (and a quantum computation in fully quantum gravity). Thus one will get a probability distribution over the viable vacua. This is the measure factor.

Let us use these ideas to define a measure factor. To start, we choose a large but finite three-manifold as our equal-time surface, and some vacuum as the initial condition. Thus this definition is only well motivated in a theory in which some vacua are "simple" and are preferred candidates for the initial condition. This seems likely in string theory.

The computer then has a rule which, given its observations and calculations, specifies how to evolve the equal-time surface forward. We grant that its observations are powerful enough to decide whether a vacuum is viable – as we discuss later this is probably **not** testing for the existence of observers, but rather small c.c., the existence of structure and chemistry, and other criteria which are necessary but perhaps not sufficient to get observers. The computer then stops with the **first** viable vacuum it finds.

Because this is semiclassical gravity and tunnelling events are probabilistic, this is a probabilistic computation (and a quantum computation in fully quantum gravity). Thus one will get a probability distribution over the viable vacua. This is the measure factor.

Let us use these ideas to define a measure factor. To start, we choose a large but finite three-manifold as our equal-time surface, and some vacuum as the initial condition. Thus this definition is only well motivated in a theory in which some vacua are “simple” and are preferred candidates for the initial condition. This seems likely in string theory.

The computer then has a rule which, given its observations and calculations, specifies how to evolve the equal-time surface forward. We grant that its observations are powerful enough to decide whether a vacuum is viable – as we discuss later this is probably **not** testing for the existence of observers, but rather small c.c., the existence of structure and chemistry, and other criteria which are necessary but perhaps not sufficient to get observers. The computer then stops with the **first** viable vacuum it finds.

Because this is semiclassical gravity and tunnelling events are probabilistic, this is a probabilistic computation (and a quantum computation in fully quantum gravity). Thus one will get a probability distribution over the viable vacua. This is the measure factor.

Shortly we will discuss this in more detail assuming that the multiverse is generated by eternal inflation. However, our definitions do not assume this, it could be any dynamics.

The main thing we need to know to make this precise is the computational cost of simulating a given region of space-time. In our early discussions we used the ansatz that a space-time region of volume V would require $\mathcal{C} = M_{Pl}^4 V$ operations to simulate.

In Brown *et al* 1509.07876 it was conjectured that the complexity to produce a state in semiclassical quantum gravity from a reference state is proportional to the action integrated over the region of space-time causally related to the surface where the state is measured – they call on gauge-gravity duality and consider the state of the boundary theory, so it is measured on the boundary.

Although we do not have gauge-gravity duality for semiclassical quantum cosmology, one can make a similar conjecture for the computational cost of simulating a new region of the multiverse.

Shortly we will discuss this in more detail assuming that the multiverse is generated by eternal inflation. However, our definitions do not assume this, it could be any dynamics.
The main thing we need to know to make this precise is the computational cost of simulating a given region of space-time. In our early discussions we used the ansatz that a space-time region of volume V would require $\mathcal{C} = M_{Pl}^4 V$ operations to simulate.
In Brown *et al* 1509.07876 it was conjectured that the complexity to produce a state in semiclassical quantum gravity from a reference state is proportional to the action integrated over the region of space-time causally related to the surface where the state is measured – they call on gauge-gravity duality and consider the state of the boundary theory, so it is measured on the boundary.
Although we do not have gauge-gravity duality for semiclassical quantum cosmology, one can make a similar conjecture for the computational cost of simulating a new region of the multiverse.

Shortly we will discuss this in more detail assuming that the multiverse is generated by eternal inflation. However, our definitions do not assume this, it could be any dynamics.

The main thing we need to know to make this precise is the computational cost of simulating a given region of space-time. In our early discussions we used the ansatz that a space-time region of volume V would require $\mathcal{C} = M_{Pl}^4 V$ operations to simulate.

In Brown *et al* 1509.07876 it was conjectured that the complexity to produce a state in semiclassical quantum gravity from a reference state is proportional to the action integrated over the region of space-time causally related to the surface where the state is measured – they call on gauge-gravity duality and consider the state of the boundary theory, so it is measured on the boundary.

Although we do not have gauge-gravity duality for semiclassical quantum cosmology, one can make a similar conjecture for the computational cost of simulating a new region of the multiverse.

Shortly we will discuss this in more detail assuming that the multiverse is generated by eternal inflation. However, our definitions do not assume this, it could be any dynamics.

The main thing we need to know to make this precise is the computational cost of simulating a given region of space-time. In our early discussions we used the ansatz that a space-time region of volume V would require $\mathcal{C} = M_{Pl}^4 V$ operations to simulate.

In Brown *et al* 1509.07876 it was conjectured that the complexity to produce a state in semiclassical quantum gravity from a reference state is proportional to the action integrated over the region of space-time causally related to the surface where the state is measured – they call on gauge-gravity duality and consider the state of the boundary theory, so it is measured on the boundary.

Although we do not have gauge-gravity duality for semiclassical quantum cosmology, one can make a similar conjecture for the computational cost of simulating a new region of the multiverse.

We are making such a conjecture, that the computational cost of simulating a region R of space-time in semiclassical cosmology is

$$\mathcal{C} = \frac{A}{\pi\hbar} \tag{1}$$

where A is the action integrated over R. In a de Sitter vacuum we have $A \propto \int \Lambda$ so (in line with holography) this is much less than the volume, though still polynomially related.

Now this definition does not always make sense, for example in the Minkowski vacuum the cost would be zero. But the semiclassical cosmologies we want to consider are largely made up of patches of metastable de Sitter. These have positive action and we will argue that in this context, the definition makes sense. It is not obvious because there are AdS bubbles.

According to this definition, the quantum complexity to simulate the observable universe is $\mathcal{C} \sim 10^{120}$. Thus we are asking whether a viable vacuum with c.c. $\sim 10^{-120}$ can be found in roughly this computational time.

We are making such a conjecture, that the computational cost of simulating a region R of space-time in semiclassical cosmology is

$$\mathcal{C} = \frac{A}{\pi\hbar} \tag{1}$$

where A is the action integrated over R. In a de Sitter vacuum we have $A \propto \int \Lambda$ so (in line with holography) this is much less than the volume, though still polynomially related.

Now this definition does not always make sense, for example in the Minkowski vacuum the cost would be zero. But the semiclassical cosmologies we want to consider are largely made up of patches of metastable de Sitter. These have positive action and we will argue that in this context, the definition makes sense. It is not obvious because there are AdS bubbles.

According to this definition, the quantum complexity to simulate the observable universe is $\mathcal{C} \sim 10^{120}$. Thus we are asking whether a viable vacuum with c.c. $\sim 10^{-120}$ can be found in roughly this computational time.

We are making such a conjecture, that the computational cost of simulating a region R of space-time in semiclassical cosmology is

$$\mathcal{C} = \frac{A}{\pi\hbar} \tag{1}$$

where A is the action integrated over R. In a de Sitter vacuum we have $A \propto \int \Lambda$ so (in line with holography) this is much less than the volume, though still polynomially related.

Now this definition does not always make sense, for example in the Minkowski vacuum the cost would be zero. But the semiclassical cosmologies we want to consider are largely made up of patches of metastable de Sitter. These have positive action and we will argue that in this context, the definition makes sense. It is not obvious because there are AdS bubbles.

According to this definition, the quantum complexity to simulate the observable universe is $\mathcal{C} \sim 10^{120}$. Thus we are asking whether a viable vacuum with c.c. $\sim 10^{-120}$ can be found in roughly this computational time.

This definition of cost motivates a new definition of global time, which we call "action time" T_A. We postulate an initial value surface Σ_0. The action time of a point p in the future of Σ_0 is then the integral of the action of the intersection of the past causal domain of dependence of p, with the future of Σ_0.

Consider a 4D dS vacuum with metric (in conformal time)

$$ds^2 = L^2 \frac{-du^2 + dx^2}{u^2}. \quad (2)$$

Let the initial slice be at $u = a$, and consider a point P at $x = 0, u = b > a$. The action of spacetime within this past lightcone is

$$T_A \sim M_P^2 L^2 \log(a/b) \sim M_P^2 L^2 \frac{T_P}{L} \quad (3)$$

where T_P is the elapsed proper time.

Thus the action time equals the elapsed proper time in Hubble units times the number of holographic bits.

This definition of cost motivates a new definition of global time, which we call “action time” T_A. We postulate an initial value surface Σ_0. The action time of a point p in the future of Σ_0 is then the integral of the action of the intersection of the past causal domain of dependence of p, with the future of Σ_0.

Consider a 4D dS vacuum with metric (in conformal time)

$$ds^2 = L^2 \frac{-du^2 + dx^2}{u^2}. \tag{2}$$

Let the initial slice be at $u = a$, and consider a point P at $x = 0, u = b > a$. The action of spacetime within this past lightcone is

$$T_A \sim M_P^2 L^2 \log(a/b) \sim M_P^2 L^2 \frac{T_P}{L} \tag{3}$$

where T_P is the elapsed proper time.

Thus the action time equals the elapsed proper time in Hubble units times the number of holographic bits.

As the dynamics proceeds, there will be a chain of tunneling events and a sequence of dS ancestor vacua leading up to a specified point p. Adding up the succession of action times, one finds

$$T_A \sim M_P^2 \sum_i L_i^2 \frac{C_i T_{P,i}}{L_i} \tag{4}$$

where T_i is the proper time spent in vacuum V_i. Thus the total action will be the total proper time along the path in (varying) Hubble units, weighted by the (varying) number of accessible holographic bits.

In general there will also be tunnelings to AdS bubbles in which the action time does not make sense. These bubbles will crunch and nobody knows quite what they mean in the landscape. Because dS regions can have AdS bubbles in their past, one needs to check that the action time is continuous and monotonically increasing. We have shown that this is the case in the dS regions.

As the dynamics proceeds, there will be a chain of tunneling events and a sequence of dS ancestor vacua leading up to a specified point p. Adding up the succession of action times, one finds

$$T_A \sim M_P^2 \sum_i L_i^2 \frac{C_i T_{P,i}}{L_i} \tag{4}$$

where T_i is the proper time spent in vacuum V_i. Thus the total action will be the total proper time along the path in (varying) Hubble units, weighted by the (varying) number of accessible holographic bits.
In general there will also be tunnelings to AdS bubbles in which the action time does not make sense. These bubbles will crunch and nobody knows quite what they mean in the landscape. Because dS regions can have AdS bubbles in their past, one needs to check that the action time is continuous and monotonically increasing. We have shown that this is the case in the dS regions.

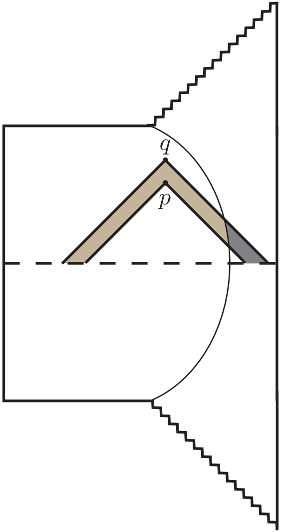
q
p

The computational meaning of action time is that it is the minimal computational time at which a point could be generated by simulation. One can also say that it is the computational difficulty of the non-deterministic version of the search problem.

If the dynamics were deterministic, we could also say that the action time of p is the minimal time needed to **verify** that a proposed cosmology including p satisfies the laws of physics. This leads us into the definition of a **complexity class** of a class of vacua in cosmology.

What is the maximal cost C_{search} (or expected cost) of finding a viable vacuum? We can generalize the problem a bit by looking, not just for $\Lambda \sim 10^{-120} M_{Pl}^4$, but to ask for the cost as a function of Λ. Now we do not know what Λ are attainable in string theory and there are arguments that the list of possibilities is finite (Acharya and Douglas). But these arguments assume a lower bound on the Kaluza-Klein scale M_{KK} – if we consider decompactification limits we can get arbitrarily small AdS $|\Lambda|$, and plausibly metastable dS as well.

The computational meaning of action time is that it is the minimal computational time at which a point could be generated by simulation. One can also say that it is the computational difficulty of the non-deterministic version of the search problem.

If the dynamics were deterministic, we could also say that the action time of p is the minimal time needed to **verify** that a proposed cosmology including p satisfies the laws of physics. This leads us into the definition of a **complexity class** of a class of vacua in cosmology.

What is the maximal cost $\mathcal{C}_{search}$ (or expected cost) of finding a viable vacuum? We can generalize the problem a bit by looking, not just for $\Lambda \sim 10^{-120} M_{Pl}^4$, but to ask for the cost as a function of Λ. Now we do not know what Λ are attainable in string theory and there are arguments that the list of possibilities is finite (Acharya and Douglas). But these arguments assume a lower bound on the Kaluza-Klein scale M_{KK} – if we consider decompactification limits we can get arbitrarily small AdS $|\Lambda|$, and plausibly metastable dS as well.

The computational meaning of action time is that it is the minimal computational time at which a point could be generated by simulation. One can also say that it is the computational difficulty of the non-deterministic version of the search problem.

If the dynamics were deterministic, we could also say that the action time of p is the minimal time needed to **verify** that a proposed cosmology including p satisfies the laws of physics. This leads us into the definition of a **complexity class** of a class of vacua in cosmology.

What is the maximal cost $\mathcal{C}_{search}$ (or expected cost) of finding a viable vacuum? We can generalize the problem a bit by looking, not just for $\Lambda \sim 10^{-120} M_{Pl}^4$, but to ask for the cost as a function of Λ. Now we do not know what Λ are attainable in string theory and there are arguments that the list of possibilities is finite (Acharya and Douglas).

But these arguments assume a lower bound on the Kaluza-Klein scale M_{KK} – if we consider decompactification limits we can get arbitrarily small AdS $|\Lambda|$, and plausibly metastable dS as well.

The computational meaning of action time is that it is the minimal computational time at which a point could be generated by simulation. One can also say that it is the computational difficulty of the non-deterministic version of the search problem.

If the dynamics were deterministic, we could also say that the action time of p is the minimal time needed to **verify** that a proposed cosmology including p satisfies the laws of physics. This leads us into the definition of a **complexity class** of a class of vacua in cosmology.

What is the maximal cost $\mathcal{C}_{search}$ (or expected cost) of finding a viable vacuum? We can generalize the problem a bit by looking, not just for $\Lambda \sim 10^{-120} M_{Pl}^4$, but to ask for the cost as a function of Λ. Now we do not know what Λ are attainable in string theory and there are arguments that the list of possibilities is finite (Acharya and Douglas). But these arguments assume a lower bound on the Kaluza-Klein scale M_{KK} – if we consider decompactification limits we can get arbitrarily small AdS $|\Lambda|$, and plausibly metastable dS as well.

We might postulate a lower bound on M_{KK} depending on Λ to get an infinite family of problems, so we can ask:
What is the asymptotic behavior of $\mathcal{C}_{search}(\Lambda)$?
This way of phrasing the question singles out Λ so it is not so general. A better way is to say, suppose the cost of simulating a single known universe is $\mathcal{C}_{univ}(\Lambda)$, then what is the relation between $\mathcal{C}_{univ}(\Lambda)$ and $\mathcal{C}_{search}(\Lambda)$? If it is polynomial, then we could say that the problem of finding such a vacuum in string cosmology is in P.
We can also define whether the problem of finding a vacuum from a given set (say, viable) is in NP. It will be if $\mathcal{C}_{search}(\Lambda)$ grows polynomially in $\mathcal{C}_{univ}(\Lambda)$, where we have an **oracle** that always makes the best choices for the search (out of polynomially many). Equivalently, we require that the problem of **verifying** that a cosmology creates the vacuum satisfy the laws of physics be doable in polynomial time.
If we advance the space-like surface Σ_0 everywhere, then if a viable vacuum appears in action time polynomial in $\mathcal{C}_{univ}(\Lambda)$, the problem of finding it would be in NP. This question has the advantage that we don't need to say much about how the search is guided.

We might postulate a lower bound on M_{KK} depending on Λ to get an infinite family of problems, so we can ask:
What is the asymptotic behavior of $\mathcal{C}_{search}(\Lambda)$?
This way of phrasing the question singles out Λ so it is not so general.
A better way is to say, suppose the cost of simulating a single known universe is $\mathcal{C}_{univ}(\Lambda)$, then what is the relation between $\mathcal{C}_{univ}(\Lambda)$ and $\mathcal{C}_{search}(\Lambda)$? If it is polynomial, then we could say that the problem of finding such a vacuum in string cosmology is in P.
We can also define whether the problem of finding a vacuum from a given set (say, viable) is in NP. It will be if $\mathcal{C}_{search}(\Lambda)$ grows polynomially in $\mathcal{C}_{univ}(\Lambda)$, where we have an **oracle** that always makes the best choices for the search (out of polynomially many). Equivalently, we require that the problem of **verifying** that a cosmology creates the vacuum satisfy the laws of physics be doable in polynomial time.
If we advance the space-like surface Σ_0 everywhere, then if a viable vacuum appears in action time polynomial in $\mathcal{C}_{univ}(\Lambda)$, the problem of finding it would be in NP. This question has the advantage that we don't need to say much about how the search is guided.

We might postulate a lower bound on M_{KK} depending on Λ to get an infinite family of problems, so we can ask:
What is the asymptotic behavior of $\mathcal{C}_{search}(\Lambda)$?
This way of phrasing the question singles out Λ so it is not so general. A better way is to say, suppose the cost of simulating a single known universe is $\mathcal{C}_{univ}(\Lambda)$, then what is the relation between $\mathcal{C}_{univ}(\Lambda)$ and $\mathcal{C}_{search}(\Lambda)$? If it is polynomial, then we could say that the problem of finding such a vacuum in string cosmology is in P.
We can also define whether the problem of finding a vacuum from a given set (say, viable) is in NP. It will be if $\mathcal{C}_{search}(\Lambda)$ grows polynomially in $\mathcal{C}_{univ}(\Lambda)$, where we have an **oracle** that always makes the best choices for the search (out of polynomially many). Equivalently, we require that the problem of **verifying** that a cosmology creates the vacuum satisfy the laws of physics be doable in polynomial time.
If we advance the space-like surface Σ_0 everywhere, then if a viable vacuum appears in action time polynomial in $\mathcal{C}_{univ}(\Lambda)$, the problem of finding it would be in NP. This question has the advantage that we don't need to say much about how the search is guided.

However as stated this would only make sense if the dynamics were deterministic. Of course the dynamics is probabilistic or quantum – so it is better to ask whether the problem of finding a viable vacuum is in BPP or in BQP. These are more or less defined by asking that the probability of finding the vacuum in polynomial time is bounded below by a number greater than $1/2$.

The nondeterministic (or verification) analog of this is the protocol classes MA (Merlin-Arthur) and QMA. Arthur is a computer with a random number generator which can solve polynomial time problems (in BPP) and Merlin is an oracle with infinite computational power. Arthur is allowed to ask Merlin questions about the problem (so, does this candidate cosmology satisfy the laws of physics), and Merlin will answer, but Arthur cannot blindly trust Merlin's answers. If there is a protocol by which Merlin can convince Arthur of the correct answer to a question with high probability, then the problem is in MA.

However as stated this would only make sense if the dynamics were deterministic. Of course the dynamics is probabilistic or quantum – so it is better to ask whether the problem of finding a viable vacuum is in BPP or in BQP. These are more or less defined by asking that the probability of finding the vacuum in polynomial time is bounded below by a number greater than $1/2$.

The nondeterministic (or verification) analog of this is the protocol classes MA (Merlin-Arthur) and QMA. Arthur is a computer with a random number generator which can solve polynomial time problems (in BPP) and Merlin is an oracle with infinite computational power. Arthur is allowed to ask Merlin questions about the problem (so, does this candidate cosmology satisfy the laws of physics), and Merlin will answer, but Arthur cannot blindly trust Merlin's answers. If there is a protocol by which Merlin can convince Arthur of the correct answer to a question with high probability, then the problem is in MA.

To apply this to cosmology, the idea (we think) is that Merlin proposes a cosmological history in which a viable vacuum is created in polynomial time, and then Arthur checks both the equations of motion and whether any random tunneling events which took place were likely or rare (by computing the amplitude using the laws of string theory), thus verifying the proposed cosmology.

Using this definition, we can check whether a class of vacua V_i are in MA by following the time evolution along a sequence of space-like surfaces of increasing action time, and defining a probability distribution over spatial geometries where the probabilities reflect the probabilities of tunneling events between vacua. We define $\mathcal{C}$ to be the time T_A after which probability that a vacuum in the class is created is greater than $2/3$. If $\mathcal{C}$ grows polynomially in $\max_i \mathcal{C}_{univ}(V_i)$, then the class is in MA.

To apply this to cosmology, the idea (we think) is that Merlin proposes a cosmological history in which a viable vacuum is created in polynomial time, and then Arthur checks both the equations of motion and whether any random tunneling events which took place were likely or rare (by computing the amplitude using the laws of string theory), thus verifying the proposed cosmology.

Using this definition, we can check whether a class of vacua V_i are in MA by following the time evolution along a sequence of space-like surfaces of increasing action time, and defining a probability distribution over spatial geometries where the probabilities reflect the probabilities of tunneling events between vacua. We define $\mathcal{C}$ to be the time T_A after which probability that a vacuum in the class is created is greater than $2/3$. If $\mathcal{C}$ grows polynomially in $\max_i \mathcal{C}_{univ}(V_i)$, then the class is in MA.

So, we can ask whether the problem of finding a given class of vacua (say dS with c.c. at most Λ) is in MA or QMA. Even if it is, we can ask whether a particular way to solve the problem attains this theoretical possibility. There are many problems for which a naive algorithm is exponential, and it takes some cleverness to find a polynomial-time algorithm – famous examples are linear programming and testing primality.

So, to summarize the questions we formulated,

1. Is it possible to find a viable vacuum in time polynomial in C_{univ} ?
2. Is it possible to verify the cosmology which finds such a vacuum in polynomial time?
3. Does the usual discussion of eternal inflation find a viable vacuum in polynomial time?
4. Can one at least verify such a cosmology in polynomial time?

We are pretty sure the answer to 4, and thus 3, is NO. We believe the answer to 2 is YES. We don't know the answer to 1.

So, we can ask whether the problem of finding a given class of vacua (say dS with c.c. at most Λ) is in MA or QMA. Even if it is, we can ask whether a particular way to solve the problem attains this theoretical possibility. There are many problems for which a naive algorithm is exponential, and it takes some cleverness to find a polynomial-time algorithm – famous examples are linear programming and testing primality.

So, to summarize the questions we formulated,

1. Is it possible to find a viable vacuum in time polynomial in $\mathcal{C}_{univ}$?
2. Is it possible to verify the cosmology which finds such a vacuum in polynomial time?
3. Does the usual discussion of eternal inflation find a viable vacuum in polynomial time?
4. Can one at least verify such a cosmology in polynomial time?

We are pretty sure the answer to 4, and thus 3, is NO. We believe the answer to 2 is YES. We don't know the answer to 1.

Eternal inflation and the multiverse

Let us now grant that the dynamics of the multiverse is eternal inflation combined with Coleman-DeLuccia tunneling. Thus, there are regions in the inflationary potential with $\delta\rho/\rho \sim 1$, so quantum fluctuations of the inflaton can compensate for the classical rolling, and some part of the multiverse is always undergoing inflation.

This has been studied since the 80's and simple arguments have been given which lead to a measure factor, given some assumptions. In the original discussions, one chooses some definition t of global time, and defines a dynamics for the evolution of $N_i(t)$, the number of universes of type i at time t. This is a Markov process,

$$\frac{d}{dt}N_i = \alpha_i N_i + \sum_j M_{i\leftarrow j}N_j - M_{j\leftarrow i}N_i, \tag{5}$$

where $M_{i\leftarrow j}$ is the matrix of transition rates from j to i, and α_i represents the effects of inflation adjusted to reflect the definition of global time. (There are also "local time" arguments.)

Eternal inflation and the multiverse

Let us now grant that the dynamics of the multiverse is eternal inflation combined with Coleman-DeLuccia tunneling. Thus, there are regions in the inflationary potential with $\delta\rho/\rho \sim 1$, so quantum fluctuations of the inflaton can compensate for the classical rolling, and some part of the multiverse is always undergoing inflation.

This has been studied since the 80's and simple arguments have been given which lead to a measure factor, given some assumptions. In the original discussions, one chooses some definition t of global time, and defines a dynamics for the evolution of $N_i(t)$, the number of universes of type i at time t. This is a Markov process,

$$\frac{d}{dt}N_i = \alpha_i N_i + \sum_j M_{i\leftarrow j}N_j - M_{j\leftarrow i}N_i, \quad (5)$$

where $M_{i\leftarrow j}$ is the matrix of transition rates from j to i, and α_i represents the effects of inflation adjusted to reflect the definition of global time. (There are also "local time" arguments.)

Define

$$P_i(t) = \frac{N_i}{\sum_j N_j} \tag{6}$$

to be the fraction at time t of vacua of type i out of all vacua. In terms of P_i, the transition matrix satisfying

$$\frac{d}{dt}P_i = \sum_j \kappa_{i\leftarrow j}P_j - \kappa_{j\leftarrow i}P_i \tag{7}$$

is a stochastic matrix, and we can interpret P_i as a probability distribution. This was worked out in Garriga *et al* 2005, and

$$\kappa_{i\leftarrow j} = \frac{4\pi}{3}M_{i\leftarrow j}H_j^{\beta-4} \tag{8}$$

where time t is related to proper time τ as

$$dt = H(\phi(\tau))^{1-\beta}d\tau. \tag{9}$$

So $\beta = 1$ is proper time, $\beta = 0$ is scale factor time, and $\beta = 2$ is action time.

$$\frac{d}{dt}P_i = \sum_j \kappa_{i\leftarrow j}P_j - \kappa_{j\leftarrow i}P_i \tag{10}$$

The most important point about this equation, which can be seen right from the start, is that it is **linear**. This is because the dynamics in any given universe is independent of all the others. Mathematically it means that we can analyze it by diagonalizing the matrix which appears on the RHS. Using detailed balance, one can find a rescaling of the N_i so that $d/dt \sum_i N_i = 0$ and there is a probability interpretation.

To go further we need to know something about $\kappa_{i\leftarrow j}$, whether there are terminal vacua which cannot tunnel back into the others. Physically this would include supersymmetric Minkowski vacua and some AdS vacua. However, even before we know this, we expect that there is a universal long time or "equilibrium" limit. It is given by N_i proportional to the **dominant** eigenvector of the matrix on the RHS, the one with the largest eigenvalue.

$$\frac{d}{dt}P_i = \sum_j \kappa_{i\leftarrow j}P_j - \kappa_{j\leftarrow i}P_i \tag{10}$$

The most important point about this equation, which can be seen right from the start, is that it is **linear**. This is because the dynamics in any given universe is independent of all the others. Mathematically it means that we can analyze it by diagonalizing the matrix which appears on the RHS. Using detailed balance, one can find a rescaling of the N_i so that $d/dt \sum_i N_i = 0$ and there is a probability interpretation.

To go further we need to know something about $\kappa_{i\leftarrow j}$, whether there are terminal vacua which cannot tunnel back into the others. Physically this would include supersymmetric Minkowski vacua and some AdS vacua. However, even before we know this, we expect that there is a universal long time or "equilibrium" limit. It is given by N_i proportional to the **dominant** eigenvector of the matrix on the RHS, the one with the largest eigenvalue.

Taking N_i to be the dominant eigenvector is attractive at first, because it removes the dependence on the initial conditions. The resulting measure factor still depends on the definition of global time, but there are probably not too many well motivated choices, and many of the choices lead to absurd results.

For example, taking t to be proper time along geodesics leads to the "youngness paradox." Let us follow the usual anthropic discussion according to which the probability of finding vacuum type i is proportional to the number of observers in such vacua. With this choice of t, there is an exponentially large advantage to delaying the end of inflation. Thus, we predict that observers will appear as early as possible – 13.7 billion years seems very late!

Taking N_i to be the dominant eigenvector is attractive at first, because it removes the dependence on the initial conditions. The resulting measure factor still depends on the definition of global time, but there are probably not too many well motivated choices, and many of the choices lead to absurd results.

For example, taking t to be proper time along geodesics leads to the "youngness paradox." Let us follow the usual anthropic discussion according to which the probability of finding vacuum type i is proportional to the number of observers in such vacua. With this choice of t, there is an exponentially large advantage to delaying the end of inflation. Thus, we predict that observers will appear as early as possible – 13.7 billion years seems very late!

This problem can be solved by instead using scale factor time, $T_S \equiv T_P/L_i$. (And in other definitions of time and other cutoffs we won't have time to discuss.) This takes out the advantage of inflation and leads to the following result (Garriga et al 0509184):
The dominant eigenvector is concentrated on the longest lived metastable dS vacuum (or "master vacuum"). Its value for a vacuum i is approximately the rate for the fastest tunneling chain from the master vacuum to i.
This prescription was studied in detail in the BP model in Schwartz-Perlov and Vilenkin 0601162. It appears sensible but they found that the measure factor is wildly varying for nearby vacua. This leads to a potential problem with the anthropic solution to the c.c. problem: if the number of viable vacua near the master vacuum is small, then it becomes likely that we live in an unusual region of a larger Λ universe. This is not a problem if this number is large.

This problem can be solved by instead using scale factor time, $T_S \equiv T_P/L_i$. (And in other definitions of time and other cutoffs we won't have time to discuss.) This takes out the advantage of inflation and leads to the following result (Garriga et al 0509184):
The dominant eigenvector is concentrated on the longest lived metastable dS vacuum (or "master vacuum"). Its value for a vacuum i is approximately the rate for the fastest tunneling chain from the master vacuum to i.
This prescription was studied in detail in the BP model in Schwartz-Perlov and Vilenkin 0601162. It appears sensible but they found that the measure factor is wildly varying for nearby vacua. This leads to a potential problem with the anthropic solution to the c.c. problem: if the number of viable vacua near the master vacuum is small, then it becomes likely that we live in an unusual region of a larger Λ universe. This is not a problem if this number is large.

So what is the longest-lived metastable vacuum in the string landscape ? Here is a reasonable guess.
A general way to suppress tunneling rates is to lower the scale $m_{3/2}$ of supersymmetry breaking. One can show (Dine et al 0901.1169) that

$$\Gamma \leq exp - \frac{2\pi^2 M_P^2}{m_{3/2}^2}. \qquad (11)$$

Thus the longest-lived vacuum will be the dS vacuum with the smallest $m_{3/2}$.
In dynamical supersymmetry breaking (and in its flux dual) we have $m_{3/2} \sim \exp -1/g^2 N$ for some effective gauge coupling g. This coupling will be set by moduli stabilization to a ratio of quantized fluxes. Thus the question is what sets this ratio.
In general the fluxes satisfy a bound depending on topological invariants of the compactification manifold M. For example in F theory we have $\sum N_{NS} \cdot N_{RR} \leq \chi/24$ where χ is the Euler character of the fourfold.

Thus the longest lived vacua come from complicated compactification manifolds with many homology cycles. For example in F theory there is a fourfold with $\chi/24 = 75852$.

This is also the compactification expected to lead to the largest number of flux vacua (Taylor and Wang 1511.03209). Thus “entropic” arguments lead to the same conclusion, that the master vacuum is a very complicated vacuum. Thus, the measure factors derived by running eternal inflation to equilibrium are concentrated on vacua which are easy to reach from a special, complicated master vacuum.

This is not inconsistent so far as I know. But it is strange. Suppose we found a relatively simple string compactification which contained the Standard Model and just enough vacua to realize the anthropic solution to the c.c. problem. Considerations such as Occam's razor would favor this vacuum.

But according to the prescription we just described, this vacuum is irrelevant. One needs to know how easy it is to produce from the master vacuum. Since it is so different, and there are surely complicated vacua containing the Standard Model, it is disfavored.

Thus the longest lived vacua come from complicated compactification manifolds with many homology cycles. For example in F theory there is a fourfold with $\chi/24 = 75852$.

This is also the compactification expected to lead to the largest number of flux vacua (Taylor and Wang 1511.03209). Thus "entropic" arguments lead to the same conclusion, that the master vacuum is a very complicated vacuum. Thus, the measure factors derived by running eternal inflation to equilibrium are concentrated on vacua which are easy to reach from a special, complicated master vacuum.

This is not inconsistent so far as I know. But it is strange. Suppose we found a relatively simple string compactification which contained the Standard Model and just enough vacua to realize the anthropic solution to the c.c. problem. Considerations such as Occam's razor would favor this vacuum.

But according to the prescription we just described, this vacuum is irrelevant. One needs to know how easy it is to produce from the master vacuum. Since it is so different, and there are surely complicated vacua containing the Standard Model, it is disfavored.

Thus the longest lived vacua come from complicated compactification manifolds with many homology cycles. For example in F theory there is a fourfold with $\chi/24 = 75852$.
This is also the compactification expected to lead to the largest number of flux vacua (Taylor and Wang 1511.03209). Thus "entropic" arguments lead to the same conclusion, that the master vacuum is a very complicated vacuum. Thus, the measure factors derived by running eternal inflation to equilibrium are concentrated on vacua which are easy to reach from a special, complicated master vacuum.
This is not inconsistent so far as I know. But it is strange. Suppose we found a relatively simple string compactification which contained the Standard Model and just enough vacua to realize the anthropic solution to the c.c. problem. Considerations such as Occam's razor would favor this vacuum.
But according to the prescription we just described, this vacuum is irrelevant. One needs to know how easy it is to produce from the master vacuum. Since it is so different, and there are surely complicated vacua containing the Standard Model, it is disfavored.

How does our proposal change these results? While the basic derivations we just discussed carry over, there are two main differences.

First, we postulate that a hospitable vacuum is found by search, and that the search stops with the **first** such vacuum to be found. Thus, every hospitable vacuum is to be considered as terminal in our prescription.

Second, if the search process produces a very long-lived metastable vacuum, then rather than wait for a tunneling event out, it can switch to simulating a different region of space-time. The long-lived vacua include the hospitable vacua but are a much larger subset.

Where it switches to depends on the search algorithm and what has already been seen, but two simple options are:

- Go back to the most recent shorter-lived vacuum. In this case it is almost as if the tunneling event did not happen (except for the time lapse needed to decide the vacuum was long-lived).
- Go back to the simple initial condition. In this case the vacuum counts as terminal.

How does our proposal change these results? While the basic derivations we just discussed carry over, there are two main differences.

First, we postulate that a hospitable vacuum is found by search, and that the search stops with the **first** such vacuum to be found. Thus, every hospitable vacuum is to be considered as terminal in our prescription.

Second, if the search process produces a very long-lived metastable vacuum, then rather than wait for a tunneling event out, it can switch to simulating a different region of space-time. The long-lived vacua include the hospitable vacua but are a much larger subset.

Where it switches to depends on the search algorithm and what has already been seen, but two simple options are:

- Go back to the most recent shorter-lived vacuum. In this case it is almost as if the tunneling event did not happen (except for the time lapse needed to decide the vacuum was long-lived).
- Go back to the simple initial condition. In this case the vacuum counts as terminal.

How does our proposal change these results? While the basic derivations we just discussed carry over, there are two main differences.

First, we postulate that a hospitable vacuum is found by search, and that the search stops with the **first** such vacuum to be found. Thus, every hospitable vacuum is to be considered as terminal in our prescription.

Second, if the search process produces a very long-lived metastable vacuum, then rather than wait for a tunneling event out, it can switch to simulating a different region of space-time. The long-lived vacua include the hospitable vacua but are a much larger subset.

Where it switches to depends on the search algorithm and what has already been seen, but two simple options are:

- Go back to the most recent shorter-lived vacuum. In this case it is almost as if the tunneling event did not happen (except for the time lapse needed to decide the vacuum was long-lived).
- Go back to the simple initial condition. In this case the vacuum counts as terminal.

How does our proposal change these results? While the basic derivations we just discussed carry over, there are two main differences.

First, we postulate that a hospitable vacuum is found by search, and that the search stops with the **first** such vacuum to be found. Thus, every hospitable vacuum is to be considered as terminal in our prescription.

Second, if the search process produces a very long-lived metastable vacuum, then rather than wait for a tunneling event out, it can switch to simulating a different region of space-time. The long-lived vacua include the hospitable vacua but are a much larger subset.

Where it switches to depends on the search algorithm and what has already been seen, but two simple options are:

- Go back to the most recent shorter-lived vacuum. In this case it is almost as if the tunneling event did not happen (except for the time lapse needed to decide the vacuum was long-lived).
- Go back to the simple initial condition. In this case the vacuum counts as terminal.

Let us consider the question, is the search more likely to find a "simple" hospitable vacuum or a "complicated" one, measured by topological invariants such as Betti numbers.

A toy model for this is to consider a landscape which is a sum of BP models. It has one component for each Betti number from 1 to N_{max}, and each component is a flux lattice $\mathbb{Z}^N$. We grant that the flux vacua only exist if $\Lambda \leq \Lambda_{max}$ so the total number is finite, and idealize this as $\sum n_i^2 \leq L^2$, so the number of vacua in the N'th component is $B_N L^N$ where B_N is the volume of an N-dimensional unit ball.

The cosmological constant is a uniformly distributed random variable. While the rates for tunnelings which change fluxes can be worked out as in Schwartz-Perlov and Vilenkin, we will idealize the tunnelings as connecting arbitrary pairs of vacua with the same N with some constant rate ϵ.

There are also tunnelings between certain vacua with N differing by 1 with average rates α_N. This average rate may be small in itself, or it may be small because only a small fraction of the vacua admit such tunneling events.

Let us consider the question, is the search more likely to find a "simple" hospitable vacuum or a "complicated" one, measured by topological invariants such as Betti numbers.

A toy model for this is to consider a landscape which is a sum of BP models. It has one component for each Betti number from 1 to N_{max}, and each component is a flux lattice $\mathbb{Z}^N$. We grant that the flux vacua only exist if $\Lambda \leq \Lambda_{max}$ so the total number is finite, and idealize this as $\sum n_i^2 \leq L^2$, so the number of vacua in the N'th component is $B_N L^N$ where B_N is the volume of an N-dimensional unit ball.

The cosmological constant is a uniformly distributed random variable. While the rates for tunnelings which change fluxes can be worked out as in Schwartz-Perlov and Vilenkin, we will idealize the tunnelings as connecting arbitrary pairs of vacua with the same N with some constant rate ϵ.

There are also tunnelings between certain vacua with N differing by 1 with average rates α_N. This average rate may be small in itself, or it may be small because only a small fraction of the vacua admit such tunneling events.

Let us consider the question, is the search more likely to find a "simple" hospitable vacuum or a "complicated" one, measured by topological invariants such as Betti numbers.

A toy model for this is to consider a landscape which is a sum of BP models. It has one component for each Betti number from 1 to N_{max}, and each component is a flux lattice $\mathbb{Z}^N$. We grant that the flux vacua only exist if $\Lambda \leq \Lambda_{max}$ so the total number is finite, and idealize this as $\sum n_i^2 \leq L^2$, so the number of vacua in the N'th component is $B_N L^N$ where B_N is the volume of an N-dimensional unit ball.

The cosmological constant is a uniformly distributed random variable. While the rates for tunnelings which change fluxes can be worked out as in Schwartz-Perlov and Vilenkin, we will idealize the tunnelings as connecting arbitrary pairs of vacua with the same N with some constant rate ϵ.

There are also tunnelings between certain vacua with N differing by 1 with average rates α_N. This average rate may be small in itself, or it may be small because only a small fraction of the vacua admit such tunneling events.

Let us take $P_H \sim 10^{-120}$ to be the fraction of hospitable vacua, independent of N once it is large enough. In the toy model, the search for such vacua is essentially a random walk. The question is, what is the distribution of N when such a vacuum is found.

This depends on the choice of how the search algorithm treats long-lived vacua and on the α_N. If long-lived vacua are treated by going back to the most recent vacuum, then there is nothing stopping the search from drifting in the direction of increasing N, and entropy favors this. Still, if $\alpha_N \lesssim P_H \epsilon$, simple vacua will be favored.

On the other hand, if long-lived vacua are treated as terminal, then if it is more likely to meet any long-lived vacuum than to tunnel in N, then simple vacua are favored. Since almost all vacua with $\Lambda \ll M_P^4$ are long lived, this seems likely.

In more formal terms, we modify the Markov process by either removing the long-lived vacua or by counting them as terminal. We then solve for the time evolution of the resulting process. If it runs for at least the mixing time (the inverse of the gap between the largest and second largest eigenvalues of the Markov matrix) then we can approximate the solution as the dominant eigenvector. Thus the question is how the mixing time compares to $1/P_H$.

Within a given N sector, and if transitions vary fluxes by ± 1, the gap is roughly ϵ/L^2. Similarly, we can expect the gap in the α transition matrix to be roughly α/N_{max}^2. If α is small, it will determine the smallest gap. So the criterion for mixing is $\alpha \gtrsim P_H L^2$ which is more or less the same as the one to reach complicated vacua.

The upshot is that if the search process does not tend to jump back to simple vacua, and if transitions between CY's of different topology are as likely or more likely than finding hospitable vacua (average rate $\alpha \gtrsim P_H L^2$), then our prescription favors complicated vacua (for intuitively sensible reasons), much as does the equilibrium prescription.

On the other hand, the actual vacua that are favored are very different – the longest lived metastable vacuum plays no role, because the search will abandon long-lived nonhospitable vacua. In the toy model, the search looks random. One needs much more detailed information about the landscape to say more.

What about the youngness paradox? One might think with $\beta = 2$, action time would be even worse than proper time from this point of view. We take from this that the measure factor cannot have an exponential inflationary factor in the P_i for hospitable vacua. In fact no such factor is there, because we are treating all the hospitable vacua as terminal. Once the search procedure generates such a vacuum, it simulates it long enough to decide whether there are observers (or perhaps a simpler criterion). If inflation produces N copies of the vacuum, it will cost time N to simulate them all, and it is up to the search algorithm whether to do this or to switch and consider other vacua. Thus multiplicity leads to no particular advantage in being found.

Finally, if the search is not random, one can imagine taking advantage of "engineered" features of the landscape. For example, consider a Calabi-Yau manifold with 100 conifold throats stabilized at a succession of energy scales $M_F, 10^{-1} M_F, 10^{-2} M_F, \ldots, 10^{-100} M_F$. By adjusting the fluxes in each throat, one could tune the c.c. easily, avoiding the need to solve an NP hard problem.

What about the youngness paradox? One might think with $\beta = 2$, action time would be even worse than proper time from this point of view. We take from this that the measure factor cannot have an exponential inflationary factor in the P_i for hospitable vacua. In fact no such factor is there, because we are treating all the hospitable vacua as terminal. Once the search procedure generates such a vacuum, it simulates it long enough to decide whether there are observers (or perhaps a simpler criterion). If inflation produces N copies of the vacuum, it will cost time N to simulate them all, and it is up to the search algorithm whether to do this or to switch and consider other vacua. Thus multiplicity leads to no particular advantage in being found.

Finally, if the search is not random, one can imagine taking advantage of "engineered" features of the landscape. For example, consider a Calabi-Yau manifold with 100 conifold throats stabilized at a succession of energy scales $M_F, 10^{-1}M_F, 10^{-2}M_F, \ldots, 10^{-100}M_F$. By adjusting the fluxes in each throat, one could tune the c.c. easily, avoiding the need to solve an NP hard problem.

Let us compare the computational complexity of the two proposals. In both cases we have argued that the time required to find the preferred vacua is of order the mixing time. But whereas in our proposal this is basically the time required to do a random walk through the landscape, in the usual definitions of measure factor, the transition rates are small, and the important ones will be $\sim \exp -M_P^2/m_{3/2}^2$. Thus, simulating the approach to equilibrium takes time exponential in $\mathcal{C}(\Lambda_{obs}) \sim 10^{120}$. Even verifying a cosmology which produces the equilibrium distribution will take exponential time.

This answers questions 3 and 4 above. As for question 2, we believe that a good deal of the search can be done using tunnelings at large Λ, which go fairly quickly. Furthermore the definition of action time $T_A \sim M_P^2 T L_i$ tells us that large Λ vacua are easier to simulate. Using this to propose a real search procedure is work in progress.

Let us compare the computational complexity of the two proposals. In both cases we have argued that the time required to find the preferred vacua is of order the mixing time. But whereas in our proposal this is basically the time required to do a random walk through the landscape, in the usual definitions of measure factor, the transition rates are small, and the important ones will be $\sim \exp -M_P^2/m_{3/2}^2$. Thus, simulating the approach to equilibrium takes time exponential in $\mathcal{C}(\Lambda_{obs}) \sim 10^{120}$. Even verifying a cosmology which produces the equilibrium distribution will take exponential time.

This answers questions 3 and 4 above. As for question 2, we believe that a good deal of the search can be done using tunnelings at large Λ, which go fairly quickly. Furthermore the definition of action time $T_A \sim M_P^2 T L_i$ tells us that large Λ vacua are easier to simulate. Using this to propose a real search procedure is work in progress.