The Anthropic Principle: A Primer for Philosophers[1]

Frank J. Tipler

Tulane University

1. Introduction: The Weak Anthropic Principle

Let's begin with a

Definition: **The Anthropic Principle** is the drawing of scientific inferences from a consideration of Man's Place in Nature.

There are various versions of the Anthropic Principle. The most conservative version of the Anthropic Principle is nothing but a systematic working out of the fact that the astrophysical (and other scientific) data we have is self-selected due to the fact that *Homo sapiens* is a particular type of intelligent being. This conservative version of the Anthropic Principle is called the

> **Weak Anthropic Principle** (WAP): The observed values of all physical quantities are not equally probable, but rather take on values restricted by the fact that these quantities are measured by a carbon-based intelligent life-form which spontaneously evolved on an earthlike planet around a G2 type star.

Again, the Weak Anthropic Principle is just a warning to take into account a grandiose type of selection bias when interpreting data. But selection bias is familiar to scientists. For an astrophysics example, suppose we want to know the fraction of all galaxies that lie in particular ranges of brightness. It is not sufficient to list all the galaxies seen according to their brightness, for the simple reason that many galaxies are too faint to be seen, or are not big enough to be distinguished from stars. Thus our observations are biased toward the very bright galaxies. We must therefore correct for this selection bias. All instruments are subject to some sort of selection bias. The Weak Anthropic Principle merely says that we — *Homo sapiens* — are also a type of measuring instrument, and it is necessary to take into account our special properties (given above) when interpreting data. WAP does *not* claim that our form of life (based on carbon, etc) is the only possible form of life. One can easily imagine non-carbon based forms of life, and indeed, as we shall see, a very speculative form of the Anthropic Principle assumes that non-carbon forms of life can exist. Such forms of life, if they do exist, are obviously not subject to the same selection biases that *Homo sapiens* is. But equally obviously, *Homo sapiens* is subject to the selection biases of *Homo sapiens*. Thus the Weak Anthropic Principle must be accepted, for it is just an application of standard scientific logic.

PSA 1988, Volume 2, pp. 27-48

The Dicke-Wheeler Relation

But the Weak Anthropic Principle is not trivial, for it leads to unexpected relationships between observed quantities that appear to be unrelated! The best example is the Dicke-Wheeler relation, which gives the measured age of the universe in terms of elementary particle time scales and fundamental interaction constants.

Now there is no *a priori* reason for believing in the existence of any definite unique relation between the age of the universe t_U and the various fundamental constants. In fact, if the Universe is open or flat (as the data suggests) t_U will eventually sweep out all values between zero and infinity, so *all* possible relations will hold at some time. However, the Weak Anthropic Principle says we (*Homo sapiens*) must measure t_U to have a value consistent with our carbon-based species having evolved on an earthlike planet. This implies that there must have been sufficient time for the elements heavier than helium to form in stars, and be scattered by supernovae; thus t_U must be greater than lifetime t_S of a typical main sequence star.

Let us make an estimate of main sequence stellar lifetime. The luminosity of a massive star is typically the Eddington Luminosity $L_E = 4\pi GMm_pc/\sigma_T$, where M is the stellar mass, m_p is the proton mass, and σ_T is the Thomson cross-section. Thus if $\eta \leq 10^{-2}$ is the fraction of the star's mass than can be released through nuclear burning, t_S is roughly

$$t_S \approx \eta Mc^2/L_E = (\eta Mc^2\sigma_T)/(4\pi GMm_pc) = \eta(\alpha_e^2/\alpha_G)(m_p/m_e)^2t_p = [\eta(\alpha_e^2(m_p/m_e)^2]\alpha_G^{-1}t_p$$

where $t_p = \hbar/m_pc^2 \approx 10^{-23}$ seconds is a typical strong interaction timescale, $\alpha_e \equiv e^2/\hbar c \approx 1/137$ is the fine structure constant (electromagnetic force strength), $\alpha_G \equiv Gm_p^2/\hbar c \approx 5 \times 10^{-39}$ is the gravitational fine structure constant, and the electron mass is m_e. The collection of constants in brackets is of order unity, so

$$t_S \approx \alpha G^{-1}t_p \approx 10^{10} \text{ years.}$$

Further, t_U can't be too many orders of magnitude larger than this, because otherwise the gas in the interstellar medium will be used up and G2 stars can't form. (This can be made more precise. The amount of gas available for the formation of new stars is, in the standard theories of star formation in the galaxy, decreasing exponentially. The exponential time constant for the exhaustion of the interstellar gas — a few billion years — implies that essentially no G2's will be formed after 10^{12} years, On the other hand, no G2's *can* be formed before 10^8 years. Probability theory tells us (e.g., Alder and Roessler 1964, pp. 33-34) that for processes characterized by a constant ratio between consecutive numbers — exponential decay is such a case, since (Feller 1968, p. 458) it is approximated by a geometric progression— the most likely time is the geometric mean of 10^8 years and 10^{12} years, namely 10^{10} years.) Thus to within a few orders of magnitude, the expected age of the Universe that our particular form of life would measure is 10^{10} years, which should also be the Hubble time H_0^{-1}, which is roughly the age of the universe So we have

$$H_0^{-1} \approx \alpha_G^{-1}t_p \approx 10^{10} \text{ years.} \qquad (1)$$

Thus the Weak Anthropic Principle has allowed us to express Hubble's constant in terms of the fundamental constants, and even to compute it to a few orders of magnitude. Again, it is important to note that (1) need not hold, and indeed won't hold much earlier and later in Universal history, since Hubble's "constant" changes with time. [(1) was derived by Dicke in 1961, and Wheeler later pointed out that it implied the Universe must be at least 10^{10} lightyears across (that is, $H_0^{-1}c$) in order to have a *single* intelligent species like us. This is why I've termed (1) the Dicke-Wheeler relation.]

Dicke *could* have used (1) to rule out the steady state theory, for in this theory, there is no necessary relationship between Hubble's constant and the main sequence lifetime, and the above derivation assumed the Big Bang. (Missed prediction!) This shows that the Weak Anthropic Principle can be used to rule out theories.

The Carter Inequality on the Age of the Earth

The Weak Anthropic Principle can also make predictions. Carter (see Barrow and Tipler, 1986 Section 8.7) has pointed out the additional temporal coincidence:

$$t_e \approx t_{ms} \qquad (2)$$

where t_e is the length of time to evolve *Homo sapiens* (4.6 billion years) and t_{ms} is the length of time the Sun will remain on the main sequence (10^{10} years). There is no reason to expect this approximate equality; indeed quite the reverse. We have seen that t_{ms} is determined by physical constants, while we would think t_e to be biologically based. Thus we would expect *a priori*

$$t_e << t_{ms} \quad \text{or} \quad t_e >> t_{ms} \qquad (3)$$

rather than (2), which is actually observed.

Carter has shown we can explain (2) from $t_{av} >> t_{ms}$ and the Weak Anthropic Principle, where t_{av} is the *average* time needed to evolve an intelligent species on an *immortal* earthlike planet. Now the Earth is not immortal as a life-containing planet, for biological evolution will cease when the Sun leaves the main sequence. If intelligent life doesn't evolve by then, it will never evolve. Thus there is a least upper bound t_{lub} to the time evolution can proceed on an earthlike planet. But the longer life exists, the more likely it is for intelligence to evolve. If t_{av} >> tms, this means intelligence is more likely to appear near t_{lub} than near the beginning of an earthlike planet, so most likely te $\approx t_{lub} \leq t_{ms}$. (Note that we need not have the equality $t_{lub} = t_{ms}$; something other that the death of the Sun may destroy life on Earth.)

This rough argument can be made quantitative. Suppose there are n crucial steps in the evolution of intelligence, each so improbable it is unlikely to occur until long after t_{lub}. Suppose also that the n steps are statistically independent and that the likelihood that the ith step occurs at any given time instant is independent of how long evolution has been going on. The time independence of the ith step is justified as follows. Suppose that the ith step is some property that is coded by a gene or a collection of genes. Each generation the genome changes, by both reassortment and mutation. If the rate of genetic change is roughly constant, and in addition if which genes/codons appear is completely random, then the likelihood that a given collection of codons or genes will arise at any given instant will be time independent. Of course, this time independence will be only an approximation, for a given mutation may be much more likely to give a viable organism if other mutations have occurred first. But I think it is a reasonable approximation. The same argument can be used to justify the statistical independence of the n steps. Now it can be shown (Feller 1968, pp. 328-329, 458-460 of Volume I and p. 8 of Volume II) that the time independence of the ith step implies that the probability it will occur by time t is given by the exponential distribution. Thus the probability that intelligence will evolve on an immortal earthlike planet by time t is given by the product of n exponential probability distributions

$$p(t) = \prod_{i=1}^{n}(1 - \exp[-t/\alpha_i]) \approx \beta t^n$$

where α_i is the expectation time for the occurrence of the ith improbable step. (We have $1 - \exp[-t/\alpha_i] \approx t/\alpha_i$, since $t << \alpha_i$ for all $t \leq t_{lub}$.) From p(t), we can compute the conditional probability that intelligence evolves at time t, given that it *must* evolve on or before t_{lub} . This conditional probability is

$$p(\text{intelligence arises by } t \mid \text{it definitely occurs before } t_{lub}) = \gamma t^n$$

where the normalization $\gamma = t_{lub}^{-n}$ is chosen to make the conditional probability equal to one when $t = t_{lub}$. The expectation $\langle t \rangle$ for the time to evolve intelligence given that it does evolve in $(0, t_{lub})$

$$\langle t \rangle = \int_0^{t_{lub}} t\,dp = \int_0^{t_{lub}} t(t_{lub}^{-n}) n t^{n-1} dt = t_{lub} n/(n + 1)$$

We would expect $\langle t \rangle \approx t_e$, so we finally obtain

$$t_{lub} - t_e \leq t_e/n$$

which is known as Carter's Inequality.

This inequality is testable in principle; we need only know t_{lub} and n. Unfortunately, modern evolutionary biology is unable to give us an estimate for n. But qualitatively, Carter's argument is based on intelligent life being improbable, so SETI (The Search for Extraterrestrial Life) is actually a test of Carter's argument. If the evolution of one-celled organisms happens to be one of the improbable steps (we don't know whether it is or not), then it follows that we are unlikely to find any sort of life elsewhere in the Solar System.

It is extremely important to note that the Dicke-Wheeler relation and the Carter inequality require the *actual* existence of an ensemble out of which we self-select our observations. For the Dicke-Wheeler relation, we are by our very existence selecting one t_U out of (0, t very large). For the Carter inequality, we are by our very existence selecting one earthlike planet out of an actual infinity (in the open or flat case), or out of a near infinity (closed case). (In fact, it is another prediction of the Carter argument that the Universe is huge. If it were small, then the ensemble would not actually exist.)

Domain Universes and the Cosmological Constant

Many modern cosmological models assume that the "total" universe (= everything that exists) is divided up into what we might call "domain universes" of varying sizes in which the various physical constants and initial conditions differ. For example, Linde's Eternal Chaotic Inflation model (Linde 1989) assumes that the total universe is composed of an infinity of domain universes, many of which begin as a Planck-sized (10^{-33} cm) quantum fluctuation inside another domain universe, and inflate to Hubble size (10^{10} lyr). These "child universes" can have different physical constants, and will appear to have different initial conditions. However such domain universes arise, in the ensemble of all of them — i.e., in the total universe — all possibilities are realized. But only in very special subsets is our form of life possible. In this way, WAP self-selection avoids the problem of initial conditions. In fact, any cosmological theory which allows different initial conditions (or constants) must *generally* invoke WAP somewhere. The best example of this is the WAP solution to the Cosmological Constant Problem.

The Cosmological Constant Problem is this (Weinberg 1989). Due to quantum fluctuations, the vacuum has energy, and by Lorentz invariance, the stress-energy tensor of vacuum must have the form $\langle T_{\mu\nu} \rangle = \langle \rho \rangle g_{\mu\nu}$, where $\langle \rho \rangle$ is a constant. This gives an effective cosmological constant $\lambda_{eff} = \lambda + 8\pi G \langle \rho \rangle$, where λ is the classical cosmological constant; equivalently, the total vacuum energy density is

$$\rho_V = <\rho> + \lambda/8\pi G = \lambda_{eff}/8\pi G \qquad (4)$$

This vacuum energy density generates a universal gravational force field: if $\lambda_{eff} > 0$ the force is repulsive and if $\lambda_{eff} < 0$ the force is attractive. Astronomical observations show that $|\rho_V| < 10^{-29}$ gm/cm$^3 \approx 10^{-47}$ GeV^4. On the basis of dimensional arguments, we would expect the energy density of vacuum to be the Planck density $\rho_{Pk} = 5 \times 10^{93}$ gm/cm^3, in which case the two terms in (4) must cancel to 122 decimal places. There are many mechanisms which can give a vacuum energy density, and most require similar cancellations. For instance, if the spontaneous symmetry breaking in electroweak theory is real, then we must have enormous λ ($\geq 10^8$ GeV^4) before the electroweak phase transition in order to have $\lambda_{eff} \approx 0$ today.

The WAP solution to the Cosmological Constant Problem is this. Let us assume the existence of an ensemble of domain universes with λ different in different domain universes. Averaged over all domain universes, the expectation value of λ is huge (the Planck density, say), but only when λ is very small — as in *our* domain universe — can intelligent life develop. We can obtain a lower bound on ρ_V as follows (Barrow and Tipler 1986, section 6.9). This lower bound will be negative. If $\rho_V < 0$, Universe expands and recollapses in time $T \leq 2\pi[8\pi G|\rho_V|/3]^{-1/2}$, independent of whether it is closed, flat, or open (Tipler 1976). As discussed above, we need $T \geq 10^9$ years for our type of intelligence to evolve; i.e., $T \geq 0.1H_0^{-1}$; where $H_0^{-1} = [8\pi G\rho_{M0}/3]^{-1/2}$ is the value of the inverse Hubble constant in an approximately flat universe and ρ_{M0} is the non-vacuum energy density now. This gives a lower bound of

$$10^3\rho_{M_0} > |\rho_V|$$

The upper (positive) bound on ρ_V has been obtained by Weinberg (1987), (1989). If $\rho_V > 0$, then ρ_V generates a repulsive gravitational force which can prevent gravitational condensations (formation of stars and planets). But as long as the non-vacuum energy density is less than ρ_V, repulsion has little effect. We know condensations started in our universe at $z_c \geq 4$, where z is the redshift, when the density was larger than ρ_{M0} by the factor $(1 + z_c)^3$, so if $\rho_V < 100\rho_{M0}$, the evolution of our type of life is definitely possible, since we in fact evolved in our (domain?) universe. On the other hand, if $\rho_V > 104\rho_{M0}$, then galaxies are disrupted very early (several star generations are not allowed), and our type life is impossible. Summarizing, WAP self-selection requires:

$$- 10^3\rho_{M_0} < \rho_V < 10^4\rho_{M_0} \qquad (5)$$

Now the expectation value $<\rho_V>$ computed from the ensemble of domain universes in face of above WAP constraints depends on the distribution of ρ_V's. For instance, if it were a Gaussian distribution with peak and standard deviation equal to *positive* Planck density, then WAP puts us in the far tail of the distribution, and the distribution is very flat over WAP range.Thus any value in range (5) equally likely. On the other hand if Nature gives us an exponential distribution of ρ_V's, with zero probability for negative ρ_V's, then we would expect an ρ_V nearer the upper bound, as happened in Carter's explanation of the age of the Earth.

The crucial point was made by Weinberg (1987, 1989): if ρ_V is actually measured to be non-zero, the *only* plausible explanation for such a state of affairs is WAP self-selection! The reason is that quantum field theory seems most unlikely to give 122 decimal cancellation and leave a residual. It is a corollary that if we actually measure ρ_V to be non-zero, this measurement would be experimental evidence for other domain universes where ρ_V is different. The current opinion in the physics community (an opinion which I share) is that in fact some quantum field mechanism will be found which gives a cosmological constant of exactly zero. The most promising proposal to date is the Hawking-Coleman "baby universe" mechanism (see Weinberg (1989) for a discussion). But if this fails, I would suggest that particle physicists begin to calculate distributions for ρ_V.

2. The Strong Anthropic Principle

In contrast to the self-selection aspects of Man's Place in Nature, consider the possibility that in some way, intelligent life is essential to the Universe. This idea is called **The Strong Anthropic Principle** (SAP). Note that there is no ensemble in SAP! In fact, the existence (or lack of) an ensemble is the basic difference between WAP and SAP. Let me warn the reader that any version of SAP is *VERY* speculative! To emphasize this, Carter (1989) has suggested that we call SAP The Strong Anthropic *Proposal* to distinguish it from WAP, which is a genuine principle of physics that we *must* accept. A *Proposal*, on the other hand, we need not accept; it is merely put forward for consideration. Why should we consider SAP? Let me remind the reader of two important facts:

(1) there is no evidence for intelligence in the Universe today; this strongly suggests it is unimportant to the Universe.

(2) due to WAP selection, we are viewing the Universe very early in its history.

Thus (2) could counter the suggestion of (1) if the true significance of life is made manifest only in the far future. But if life dies out (ultimately wiped out by the Heat Death, say), it is hard to see any essential role for life to play. This leads to

The Final Anthropic Principle (FAP): the Universe is sufficiently benign so that once intelligence first evolves, the laws of physics *permit* its continued existence forever.

I think scientists should take the FAP seriously because we have to have *some* theory for the future of the physical universe — since it unquestionably exists — and the FAP is based on the most beautiful physical postulate: that total death is not inevitable. *All* other theories of the future necessarily postulate the ultimate extinction of everything we could possibly care about. I once visited a nazi death camp; there I was reinforced in my conviction that there is nothing uglier than extermination. We physicists know that a beautiful postulate is more likely to be correct than an ugly one. Why not adopt the Postulate of Eternal Life — FAP, that the extinction of everything we could possibly care about is not inevitable — at least as a working hypothesis? (See Linde (1989), Abramowicz and Ellis (1989), and Barrow (1989) for more discussion of FAP.)

The Omega Point Theory

In order to investigate whether life can continue to exist forever, I shall need to define "life" in physics language. I claim that a "living being" is any entity which codes "information" (in the sense this word is used by physicists) with the information coded being preserved by natural selection. (I justify this definition in section 8.2 of Barrow and Tipler 1986). Thus "life" is a form of information processing, and the human mind — and the human soul — is a very complex computer program. Specifically, a "person" is defined to be a computer program which can pass the Turing Test (See Hofstadter and Dennett 1981, 69-95 (Chapter 5) for a detailed discussion of this test).

There is actually an astonishing similarity between the mind-as-computer-program idea and the medieval Christian idea of the "soul". Both are fundamentally "immaterial": a program is a sequence of integers, and an integer — 2, say — exists "abstractly" as the class of all couples. The symbol "2" written here is a *representation* of the number 2, and not the number 2 itself. In fact, Aquinas (following Aristotle) defined the *soul* to be "the form of activity of the body" . In Aristotelian language, the *formal* cause of an action is the abstract cause, as opposed to the material and efficient causes. For a computer, the program is the formal cause, while the material cause is the properties of the matter out of which the computer is made, and the efficient cause is the opening and closing of electric circuits. For Aquinas, a human soul needed a body to think and feel, just as a computer program needs a physical computer to run.

Aquinas thought the soul had two faculties: the agent intellect (*intellectus agens*) and the receptive intellect (*intellectus possibilis*), the former being the ability to acquire concepts, and the latter being the ability to retain and use the acquired concepts. Similar distinctions are made in computer theory: general rules concerning the processing of information coded in the central processor are analogous to the agent intellect; the programs coded in RAM or on a tape are analogues of the receptive intellect. (In a Turing machine, the analogues are the general rules of symbol manipulation coded in the device which prints or erases symbols on the tape vs. the tape instructions, respectively.) Furthermore, the word "information" comes from the Aristotle-Aquinas' notion of "form": we are "informed" if new forms are added to the receptive intellect. Even semantically, the information theory of the soul is the same as the Aristotle-Aquinas theory.[2]

The "mind as computer program" idea is absolutely central to this paper; indeed, it forms the basis of a revolution now going on in mathematics, physics, and philosophy. The best defense of the idea can be found in Hofstadter and Dennett's *The Mind's I* (1981), particularly pages 69-95; 109-115; 149-201; and 373-382.

In the language of information processing it becomes possible to say precisely what it means for life to continue forever.

Definition: I shall say that life can continue forever if:

(1) information processing can continue indefinitely along at least one world line γ all the way to the future "c-boundary" of the universe; that is, until the end of time.

(2) the amount of information processed between now and this future c-boundary is infinite in the region of spacetime with which the world line γ can communicate.

(3) the amount of information stored at any given time τ within this region can go to infinity as t approaches its future limit (this future limit of τ is finite in a closed universe, but infinite in an open one, if τ is measured in what physicists call "proper time").

The above is a rough outline of the more technical definition given in Section 10.7 of Barrow and Tipler 1986. (See also (Tipler 1986) and (Tipler 1988)) But let me ignore details here. What is important are the physical (and ethical!) reasons for imposing each of the above three conditions. The reason for condition 1 is obvious: it simply states that there must be at least one history in which life (=information processing) never ends. (See below for more on what "c-boundary" means. For now, think of it as meaning "the end of time.").

Condition 2 tells us two things: First, that information processed is "counted" only if it is possible, at least in principle, to communicate the results of the computation to the history γ. This is important in cosmology, because event horizons abound. In the closed Friedmann universe, which is the standard (but over-simplified) model of our actual universe (if it is in fact closed), every comoving observer loses the ability to send light signals to every other comoving observer, no matter how close. Life obviously would be impossible if one side of one's brain became *forever* unable to communicate with the other side. Life is organization, and organization can only be maintained by constant communication among the different parts of the organization. The second thing condition 2 tells us is that the amount of information processed between now and the end of time is potentially infinite. I claim that it is meaningful to say that life exists forever only if the number of thoughts generated between now and the end of time is actually infinite. But we know that each "thought" corresponds to a minimum of one bit being processed. In effect, this part of condition 2 is a claim that time duration is most properly measured by the thinking rate, rather than by proper time as measured by atomic clocks. The length of time it takes an intelligent being to process one bit of information — to think one thought — is a direct measure of "subjective" time, and hence is the most important measure of time from the perspective of life. A person who has thought 10 times as much, or experienced 10 times as much (there is no basic physical difference between these options) as the average person has in a fundamental sense lived 10 times as long as the average person, even if the rapid thinking person's chronological age is shorter than the average.

The distinction between proper and subjective time crucial to condition 2 is strikingly similar to a distinction between two forms of duration in Thomist philosophy. Recall that Aquinas distinguished three types of duration. The first was *tempus*, which is time measured by change in relations (positions, for example) between physical bodies on Earth. *Tempus* is analogous to proper time; change in both human minds and atomic clocks is proportional to proper time, and for Aquinas also, *tempus* controlled change in corporeal minds. But in Thomist philosophy, duration for *incorporeal* sentient beings — angels — is controlled not by matter, but rather is measured by change in the mental states of these beings themselves. This second type of duration, called *aevum* by Aquinas, is clearly analogous to what I have termed "subjective time". *Tempus* becomes *aevum* as sentience escapes the bonds of matter. Analogously, condition 2 requires that thinking rates are controlled less and less by proper time as τ approaches its future limit. *Tempus* gradually becomes *aevum* in the future. (The third type of Thomist duration is *aeternitas*, which can be thought of as "experiencing" all past, present, and future *tempus* and *aevum* events in the universe all at once. But more of *aeternitas* later.)

Condition 3 is imposed because although condition 2 is necessary for life to exist forever, it is not sufficient. If a computer with a finite amount of information storage — such a computer is called a *finite state machine* — were to operate forever, it would start to repeat itself over and over. The psychological cosmos would be that of Nietzsche's Eternal Return. Every thought and every sequence of thoughts, every action and every sequence of actions, would be repeated not once but an infinite number of times. It is generally agreed (by everyone but Nietzsche) that such a universe would be morally repugnant or meaningless. Only if condition 3 holds in addition to condition 2 can a psychological Eternal Return be avoided. Also, it seems reasonable to say that "subjectively," a finite state machine exists for only a finite time even though it may exist for an infinite amount of proper time and process an infinite amount of data. A being (or a sequence of generations) that can be truly said to exist forever ought to be physically able, at least in principle, to have new experiences and to think new thoughts.

Let us now consider whether the laws of physics will permit life/information processing to continue forever. John Von Neumann and others have shown that information processing (more precisely, the irreversible storage of information) is constrained by the first and second laws of thermodynamics. Thus the storage of a bit of information requires the expenditure of a definite minimum amount of available energy, this amount being inversely proportional to the temperature (See Section 10.6 of Barrow and Tipler 1986 for the exact formula). This means it is possible to process and store an infinite amount of information between now and the Final State of the universe only if the time integral of P/T is infinite, where P is the power used in the computation, and T is the temperature. Thus the laws of thermodynamics will permit an infinite amount of information processing in the future, provided there is sufficient available energy at all future times.

What is "sufficient" depends on the temperature. In the open and flat ever-expanding universes, the temperature drops to zero in the limit of infinite time, so less and less energy per bit processed is required with the passage of time. In fact, in the flat universe, only a *finite* total amount of energy suffices to process an infinite number of bits! This finite energy just has to be used sparingly over infinite future time. On the other hand, closed universes end in a final singularity of infinite density, and the temperature diverges to infinity as this final singularity is approached. This means that an ever-increasing amount of energy is required per bit near the final singularity. However, most closed universes undergo "shear" when they recollapse, which means they contract at different rates in different directions (in fact, they spend most of their time *expanding* in one direction while contracting in the other two!). This shearing gives rise to a radiation temperature difference in different directions, and this temperature difference can be shown to provide sufficient free energy for an infinite amount of information processing between now and the final singularity, even though there is only a *finite* amount of proper time between now and the end of time in a closed universe. Thus although a closed universe exists for only a finite proper time, it nevertheless could exist for an infinite subjective time, which is the measure of time that is significant for living beings.

But although the laws of thermodynamics permit conditions 1 through 3 to be satisfied, this does not mean that the other laws of physics will. Its turns out that although the energy is available in open and flat universes, the information processing must be carried out over larger and larger proper volumes. This fact ultimately makes impossible any communication between opposite sides of the "living" region, because the redshift implies that arbitrarily large amounts of energy must be used to signal (This difficulty was first pointed out by Freeman Dyson.). This gives the

First Testable Prediction of the Omega Point Theory: the universe must be closed.

However, as I stated earlier, there is a communication problem in most closed universes — event horizons typically appear, thereby preventing communication. But there is a rare class of closed universes which don't have event horizons, which means by definition that every world line can always send light signals to every other world line. Now Roger Penrose has found a way to define precisely what is meant by the "boundary" of spacetime, where time ends. In his definition of the "c-boundary", world lines are said to end in the same "point" on this boundary if they can remain in causal contact unto the end of time. If they eventually fall out of causal contact then they are said to terminate in different c-boundary points. Thus the c-boundary of these rare closed universes without event horizons consists of a single point. For reasons given in Section 10.6 of Barrow and Tipler 1986 (see also Section 3.7), it turns out that information processing can continue only in closed universes which end in a single c-boundary point, and only if the information processing is ultimately carried out throughout the entire closed universe. Thus we have the

Second Testable (?) Prediction of the Omega Point Theory: the future c-boundary of the universe consists of a single point; call it the *Omega Point.* (Hence the name of the theory.)

It is possible to obtain other predictions. For example, a more detailed analysis of how energy must be used to store information leads to the

Third Testable Prediction of the Omega Point Theory: the density of particle states must diverge to infinity as the energy goes to infinity, but nevertheless this density of states must diverge no faster than the square of the energy .

The Omega Point has an interesting property. Mathematically, the c-boundary is a completion of spacetime: it is not actually in spacetime, but rather just "outside" it. If one looks more closely at the c-boundary definition, one sees that a c-boundary consisting of a single point is formally equivalent to the entire collection of spacetime points. In effect, all the different instants of universal history are collapsed into the Omega Point; "duration" for the Omega Point can be regarded as equivalent to the collection of all experiences of all life that did, does, and will exist in the whole of universal history, together with all non-living instants. This "duration" is very close to the idea of *aeternitas* of Thomist philosophy. We could say that *aeternitas* is equivalent to the union of all *aevum* and *tempus.*

This identification of the Omega Point with the whole of the past, present, and future universal history is more than a mere mathematical artifact. *The identification really does mean that the Omega Point "experiences" the whole of universal history "all at once!"* For consider what it means for us to "experience" an event. It means we think and emote about an event we see, hear, feel, etc. Consider for simplicity just the "seeing" mode of sensing. We see another contemporary person by means of the light rays that left her a fraction of a second ago. But we cannot "see" a person that lived a few centuries before, because the light rays from said person have long ago left the solar system. Conversely, we cannot "see" the Andromeda Galaxy as it now is, but rather we "see" it as it was 2 million years ago. So we experience as "simultaneous" the events on the boundary of our past light cone (for the seeing mode; it is more complicated for all other modes of sensing, for we experience as simultaneous events which reach us at the same instant along certain timelike curves from inside our past light cone).

But all timelike and lightlike curves converge upon the Omega Point. In particular, all the light rays from all the people who died a thousand years go, from all the people now living, and from all the people who will be living a thousand years from now, will intersect there. The light rays from those people who died a thousand years ago are not lost forever; rather these rays will be intercepted by the Omega Point. Or to put it another way, these rays will be intercepted and intercepted again, by the living beings who have engulfed the physical universe near the Omega Point. All the information which can be extracted from these rays will be extracted at the instant of the Omega Point. The beings existing at that last instant of time can experience the whole of time simultaneously just as we experience simultaneously the Andromeda Galaxy and a person in the room with us. (I should warn the reader that I have ignored the problem of opacity and the problem of loss of coherence of the light. Until these are taken into account, I cannot say exactly how much information can in fact be extracted from the past. But at the most basic ontological level, *all* the information from the past [= all of universal history] remains in the physical universe and is available for analysis.)

The preceding analysis was entirely classical. However, since the existence of life for infinite subjective time requires the universe to actually achieve the final singularity — the Omega Point — the Omega Point theory is necessarily a quantum cosmological theory: almost all of subjective experience occurs after the universe has recollapsed past the Planck size! I shall briefly outline how one quantizes the Omega Point Theory.

In classical general relativity, a spacetime is generated from its initial data in the following manner. One is given a *3-dimensional* manifold S, and on S the non-gravitational fields F (and their appropriate derivatives F'), and two tensor fields h and K, with (F, F', h,K) satisfying certain equations called constraint equations. The constraint equations say nothing about the time evolution; rather they are to be regarded as consistency conditions amongst the fields (F, F',h,K) which must be satisfied at every instant of time. The physical interpretation of h is that of a spatial metric of the manifold S, and so S and (F, F',h,K) together represent the entire spatial universe at an instant of universal time. S and (F, F',h,K) are called the initial data. We now try to find a *4-dimensional* manifold M with metric g and spacetime non-gravitational fields F such that (1) M contains S as a submanifold; (2) g restricted to S is the metric h, and (3) K is the "extrinsic curvature" of S in M (roughly speaking, K says how rapidly h is changing in "time"). The manifold M and the fields (g,F) are then the whole of physical reality, including the underlying background spacetime (that is, (M,g)), the gravitational field (represented by the spacetime metric g), and all the non-gravitational fields (given by F). There will be infinitely many such M's and g's, but one can cut down the number by requiring that g satisfies the Einstein field equations everywhere on M, and that the Einstein field equations reduce to the constraint equations on S.

But even requiring the Einstein equations to hold everywhere leaves infinitely many spacetimes (M,g) which are generated from the *same* initial data at the spacetime instant

S. To see this, suppose we have found a spacetime (M,g) which in fact has S and its initial data as the spatial universe at some instant t_0 of universal time. Pick another universal time t_1 to the future of t_0 and cut away all of the spacetime in (M,g) to the future of t_1 (including the spatial instant corresponding to t_1.) This gives a new spacetime (M',g) which coincides with (M,g) to the past of t_1, but which has absolutely nothing — no space, no time, no matter — to the future of t_1. Clearly, both (M,g) and (M',g) are spacetimes which are both generated from S and its initial data. Furthermore, the Einstein equations are satisfied everywhere on both spacetimes. There are infinitely many ways we can cut away (M,g) in this way, so there is an infinity of (M',g)'s we can construct. True, the universe (M',g) ends abruptly at t_1, for no good reason. But what of that? The point is, the field equations themselves cannot tell us that the physical universe should continue past the time t_1. Rather, in classical general relativity one must impose as a separate assumption, over and above the assumption of the field equations and the initial data, that the physical universe must continue in time until the field equations themselves tell us that time has come to an end (at a spacetime singularity, say).

It is possible to prove (Hawking and Ellis 1973, Chapter 7) that there is amongst all the mathematically possible (M',g) —we might call these "possible worlds" — a *unique* "maximal" spacetime (M,g) which is generated by the initial data on S. "Maximal" means that the spacetime (M,g) contains any other (M',g) generated by the initial data on S as a proper subset. In other words, (M,g) is the spacetime we get by continuing the time evolution until the field equations themselves won't allow us to go further. This maximal (M,g) is the natural candidate for the spacetime that is actualized, but it is important to keep in mind that this is a physical assumption: all of the (M',g) are possible worlds, and any one of these possible worlds could have been the one that really exists.

Once we have the maximal (M,g) generated from a given S and its initial data, there is an infinity of other choices of 3-dimensional manifolds in M which we could picture as generating (M,g). For example, we could regard the spatial universe and the fields it contains now as "S with its initial data", or we could regard the universe a thousand years ago as "S with its initial data". Both would give the same (M,g) since the Einstein equations are deterministic. Everything that has happened and will happen is contained implicitly in the initial data on S. There is nothing new under the sun in a deterministic theory like general relativity. One could even wonder why time exists at all since from an information standpoint it is quite superfluous. (I'll suggest an answer to this question in the next Section.) None of the infinity of initial data manifolds in (M,g) can be uniquely regarded as generating the whole of spacetime (M,g). Each contains the same information, and each will generate the same (M,g), including all the other initial data manifolds.

Even in deterministic theories, relationships between physical entities are different at different times. For example, two particles moving under Newtonian gravity are now 2 meters apart (say), and a minute later 4 meters apart. This is true even though given the initial position and velocities when they were 2 meters apart, it is determined then that they will be 4 meters apart a minute later. The question is, will the totality of relationships at one time become the same (or nearly the same) at some later time? If this happens, then we have the horror of the Eternal Return. As is well-known, it is possible to prove that the Eternal Return will occur in a Newtonian universe provided said universe is finite in space and finite in the range of velocities the particles are allowed to have. It is possible to prove that in classical general relativity (Tipler 1979; 1980), the Eternal Return *cannot* occur. That is, the physical relationships existing now between the fields will never be repeated, nor will the relationships ever return to approximately what they now are. What happens is that the Einstein field equations will not permit the gravitational equivalent of the "range of velocities" to be finite: the range simply must eventually become infinite. Thus history, understood as an unrepeatable temporal sequence of relationships between physical entities, is real.

Since I am interested here in discussing quantum cosmology, I am virtually forced into adopting the Many Worlds Interpretation, because only in this interpretation is it meaningful to talk about a quantum universe and its ontology. The Copenhagen Interpretation assumes that a process called "wave function reduction" eliminated quantum effects on cosmological scales an exceedingly short time after the Big Bang, so the universe today is not quantum except on very small scales. The problem with this assumption is that the wave function reduction process is almost entirely mysterious — we have no rules for deciding what material entity can reduce wave functions — so it is impossible to give a sharp analysis of contingency when this process is operating. The Many Worlds Interpretation does not suffer from this drawback: there is no reduction of the wave function, physical reality is completely described by the wave function of the universe, there is an equation (the Wheeler-DeWitt equation) for this wave function, and the universe is just as quantum now as it was in the beginning. Of course, the Many Worlds Interpretation may be wrong; most physicists think it is (most physicists think it's nonsense). But the overwhelming majority of people working on quantum cosmology subscribe to some version of the Many Worlds Interpretation, simply because the mathematics forces one to accept it. The mathematics may be a delusion, with no reference in physical reality. Or the situation may be similar to that of early 17th century physics: astronomers believed the Earth went around the Sun, because the mathematics of the Copernican system forced them to. But few other scholars or ordinary people believed the Earth moved. Their own senses told them it did not. I shall adopt the Many Worlds Interpretation in what follows. For a more detailed defense of this interpretation see Section 7.2 of Barrow and Tipler 1986.

In quantum cosmology the universe is represented by a wave function Ψ(h,F,S), where as in classical general relativity, h and F are respectively the spatial metric and the non-gravitational fields given on a 3-dimensional manifold S. The initial data in quantum cosmology are not (h,F) given on S as was the case in classical general relativity, but rather Ψ(h,F,S) and its first derivatives. From this initial data, the Wheeler-DeWitt equation determines Ψ(h,F,S) for all values of h and F. In other words, the wave function, not the metric or the non-gravitational field, is the basic physical field in quantum cosmology. It is the initial wave function (and appropriate derivatives) that must be given, but once given, it is determined everywhere. What we think of as the most basic fields in classical general relativity, namely h and F, play the role of coordinates in quantum cosmology. But this does not mean h and F are unreal. They are as real as they are in classical theory. But it does mean more than one h and F exist on S at the same time! To appreciate this, recall that the classical metric h(x) is a function of the spatial coordinates on the manifold S. This metric has (non-zero) values at all points on S; that is, for the entire range of the coordinates as they vary over S, which is to say, as we go from one point to another in the universe. Each value of h(x) is equally real, and all of the values of h at all of the points of $_S$ exist simultaneously. Similarly, the points in the domain of the wave function Ψ(h,F,S) are the various possible values of h and F, each set (h,F) corresponding to a complete universe at a given instant of time. The central claim of the Many Worlds Interpretation is that each of these universes actually exists, just as the different h(x) exist at the various points of S: Quantum reality is made up of an infinite number of universes (worlds). Of course, we are not aware of these worlds — we are only aware of one — but the laws of quantum mechanics explain this: we must generally be as unaware of these parallel worlds as we are of our motion with the Earth around the Sun. (In extreme conditions, for instance near singularities, it is possible for the worlds to effect each other in a more obvious way than they do now.)

To fix the classical initial data, we pick a *function* h(x) out of an infinite number of possible metric functions which could have been on S. All of these possible worlds comprise a function space. To fix the quantum initial data, we pick a *wave function* Ψ(h,F,S) out of an infinite number of possible wave functions which could have been on the classical function space (h,F). Remember, however, that all values of the function space (h,F) *really are* on S simultaneously. In quantum cosmology, the collection of all possible wave functions forms the set of the possible worlds; what is contingent is which single unique universal wave function is actualized. But the possible worlds of classical cosmology — the space of all physically possible (h,F) on S — are no longer contingent. All of them are actualized.

In classical deterministic general relativity, we had a philosophical problem with time: since everything that did or will happened was coded in the initial fields on S, time evolution appeared superfluous. What was the point of having time? The problem is solved in quantum cosmology: *there is no time!* The universal wave function Ψ(h,F,S) is all there is, and there is no reference to a 4-dimensional manifold M or a 4-dimensional metric g in the wave function. At the most basic ontological level, time does not exist. Everything is on the 3-dimensional manifold S. How can this be? Of course we see time! Or do we? What we see is relationships between objects — configurations of physical fields —in space. In the discussion of the Eternal Return, I argued that time and history could be truly real only if the spatial relationships between the various fields never returned to a previous state. In quantum cosmology, there is no spacetime in which the spatial relationships between fields can change. Rather, all we have is paths (trajectories) in the collection (h,F) of all possible relationships between the physical fields on S. But this is enough, because each such path defines a history, a complete spacetime.

To understand this, imagine that we are at a point P in (h,F), and have selected a particular path g in (h,F) starting at P. Each point, remember, corresponds to an entire universe (spatially). As we go along g, the relationships between the physical fields vary smoothly from their values at P. *This variation would appear as temporal variation* from inside the path γ, because each point on γ is a complete spatial universe, and thus the sequence of points constitute a sequence of spatial universes. But this is exactly the same as the classical 4-dimensional manifold M with its spacetime metric γ and spacetime fields F, which in the above classical analysis we obtained as an extension of S and its fields! Each path in (h,F) thus is an entire classical universal history, an entire spacetime.

All paths in (h,F) really exist, which necessarily means that all — and I mean all — histories which are consistent with the "stuff" of the universe being (h,F) really exist. In particular, even histories which are grossly inconsistent with the laws of physics really occur! Closed paths in (h,F) obviously exist, so there are histories in which the Eternal Return is true. There are also real histories leading to our presently observed state of the universe (the point P in (h,F)) in which real historical characters — for instance Julius Caesar — never existed. What happens in such a history is that the physical fields rearrange themselves over time (more accurately, over the path corresponding to this strange history) to create false memories, including not only human memories but also the "memories" in a huge number of written records and in massive monuments. Just as there is an infinity of actual pasts which have led to the present state, so there is an infinity of really existing futures which evolve from the present state. So every consistent future is not only possible but it really happens. But not all futures are equally likely to be seen. That is, there is one path in (h,F) leading from a given point P which is overwhelmingly more likely to follow from P than all the others. This path is called *the classical path.* Along this path, the laws of physics hold, and memories are reliable. A classical path in (h,F) very closely resembles a classical spacetime (M,g) obeying the Einstein equations.

So far I have not said what the wave function Ψ itself does. But it must do something physically detectable, something not coded in the fields (h,F) alone. If it did not exert some physical effect, we could just omit it from physics; it would have no real existence. But I claimed above that Ψ was a *real* field, something as real as the fields (h,F).

What Ψ does is determine the set of all classical paths, and also the "probabilities" which are associated with each point and each path in (h,F). A wave function is a complex function, and all complex functions are actually two functions, a "magnitude" and a "phase". The classical paths are by definition those which are perpendicular to the surfaces of constant phase. The square of the magnitude at a point P in (h,F) is the "probability" of that point. The physicists Heisenberg and Mott showed mathematically that if "probability" has its usual meaning, then given the fact that we are (approximately) at P, the conditional probability of going to a nearby point Q is maximum if Q lies along the classical path through P. The relative probability is very close to 1 on the classical path, and it drops rapidly to 0 as one moves away from the classical path connecting P and Q. (See Section 7.2 of Barrow and Tipler 1986 for details about how this works.)

What must be shown is that the square of the magnitude is in fact a "probability" in the usual sense. This is done as follows. We obviously can't get hold of the wave function of the entire universe, but we can prepare in the laboratory a number N of electrons with the same spin wave function. Suppose we measure the vertical component of the electron spin. It turns out that this component can have only two values, spin up and spin down. If the wave function is not in what is called "an eigenstate" of spin up or spin down — in general the electron wave function would not be in an eigenstate, so let's suppose it's not — then each time we measure the vertical component of an electron in our ensemble of N electrons, we will get a different answer. Some of the electrons will be found to have spin up, and the others will have spin down. We can't predict before the measurement what the vertical component of that particular electron will be. But it can be shown that if we compute the relative frequency with which we get spin up, then this number approaches the square of the magnitude of the wave function evaluated at "spin up" as the number N of electrons in the ensemble approaches infinity. And experimentally, this is what we see.

All the physics is contained in the wave function. In fact, the laws of physics themselves are completely superfluous. They are coded in the wave function. The classical laws of physics are just those regularities which are seen to hold along a classical path by observers in that classical path. Along other paths, there would be other regularities, different laws of physics. And these other paths exist and hence these other laws of physics really hold; it is just extremely unlikely we will happen to see them operating. The Wheeler-DeWitt equation for the wave function is itself quite superfluous. It is merely a crutch to help us to find the actual wave function of the universe. If we knew the boundary conditions which the actual universal wave function satisfied, then we could derive the Wheeler-DeWitt equation, which is just a particular equation (among many) which the wave function happens to satisfy. Thus in quantum cosmology, there is no real contingency in the laws of physics. Any law of physics holds in some path, and the law of physics governing the universal wave function can be derived from that wave function. All the contingency in quantum cosmology is in the wave function, or rather, in the boundary conditions which pick out the wave function which actually exists. The well-known Hartle-Hawking boundary condition, which says that "the universal wave function is that wave function for which the Feynman sum over all the paths (classical and otherwise) leading to a given point P is over paths that have no boundaries (more precisely, the 4-dimensional manifold corresponding to a given path is a compact manifold whose only boundary is P)" is one such boundary condition. I should like to propose the

Teilhard Boundary Condition for the universal wave function:

The wave function of the universe is that wave function for which all classical paths terminate in a (future) Omega Point, with life coming into existence along at least one classical path and continuing into the future forever all the way into the Omega Point.

The Teilhard Boundary Condition is enormously restrictive. For example, since classical paths are undefined at zeros of the wave function, we immediately have

Fourth Testable Prediction of the Omega Point Theory: the universal wave function must have no zeros in the spacetime domain.

It turns out (as one might expect) that the Hartle-Hawking boundary condition does not satisfy the Teilhard boundary condition. I have a rough argument that one can construct simple quantized Friedmann cosmological models in which all classical paths terminate in an Omega Point, but I don't know yet what the existence of life requires of a wave function. So at present I can only conjecture, not prove, that a wave function satisfying the Teilhard Boundary Condition in its full generality exists mathematically. (This is not unusual; there is also no general existence proof yet for the Hartle-Hawking boundary condition.)

3. The Universe Necessarily Exists: A Strong Anthropic Ontological Argument

Ever since Kant showed that "existence is not a predicate", most philosophers have felt ontological arguments to be invalid; that is, they have believed it impossible to prove the existence of anything by means of logic alone . I want to claim this is incorrect; I think you can prove that the universe necessarily exists. The proof will be based on an analysis of what the word "existence" means.

Let us begin with some computer metaphysics. Much of computer science is devoted to making *simulations* of phenomena in the physical world. In a simulation, a mathematical model of the physical object under study is coded in a program. The model includes as many attributes of the real physical object as possible (limited of course by the knowledge of these attributes, and also by the capacity of the computer). The running of the program evolves the model in time. If the initial model is accurate, if enough key features of the real object are captured by the model, the time evolution of the model will mimic with fair accuracy the time development of the real object, and so one can predict the most important key aspects which the real object will have in the future.

Suppose we try to simulate a city full of people. Such simulations are being attempted now, but at a ludicrously inaccurate level. But suppose we imagine more and more of the attributes of the city being included in the simulation. In particular, more and more properties of each individual person are included. In principle, we can imagine a simulation being so good that every single *atom* in each person and each object in the city and the properties of each atom having an analogue in the simulation. Let us imagine, in the limit, a simulation that is absolutely perfect: each and every property of the real city, and each and every real property of each real person in the real city is represented precisely in the simulation. Furthermore, let us imagine that when the program is run on some gigantic computer, the temporal evolution of the simulated persons and their city precisely mimics for all time the real temporal evolution of the real people and the real city.

The key question is this: do the simulated people exist? As far as the simulated people can tell, they do. By assumption, any action which the real people can and do carry out to determine if they exist — reflecting on the fact that they think, interacting with the environment - the simulated people also can do, and in fact do do. There is simply no way for the simulated people to tell that they are "really" inside the computer, that they are merely simulated, and not real. They can't get at the real substance, the physical computer, from where they are, inside the program. One can imagine the ultimate simulation, a perfect simulation of the entire physical universe, containing in particular all people which the real universe contains, and which mimics perfectly the actual time evolution of the actual universe. Again, there is no way for the people inside this simulated universe to tell that they are merely simulated, that they are only a sequence of numbers being tossed around inside a computer, and are in fact not real.

How do we know we ourselves are not merely a simulation inside a gigantic computer? Obviously, we can't know. But I think it is clear we ourselves really exist. Therefore, *if* it is in fact possible for the physical universe to be in precise one to one correspondence with a simulation, I think we should invoke the Identity of Indiscernibles and identity the universe and all of its perfect simulations. (For more discussion of whether a simulation must be regarded as real if it copies the real universe sufficiently closely, see (Hofstadter and Dennett 1981), particularly pages 73-78, 94-99, and 287-320.)

But is it possible for the universe to be in precise one-to-one correspondence with some simulation? I think that it is, if we generalize what we mean by simulation. In computer science, a simulation is a program, which is fundamentally a map from the set of integers into itself. That is, the instructions in the program tell the computer how to go from the present state, represented by a sequence of integers, to the subsequent state, also represented by a sequence of integers. But remember, we don't really need the physical computer; the initial sequence of integers and the general rule (instructions or map) for replacing the present sequence by the next is all that is required. But the general rule can itself be represented as a sequence of integers. So, if time were to exist globally, and if the most basic things in the physical universe and the time steps between one instant and the next were discrete, then the whole of spacetime would definitely be in one to one correspondence with some program. But time may not exist globally (it doesn't if standard quantum cosmology is true), and it may be that the substances of the universe are continuous fields and not discrete objects (in *all* current physical theories, the basic substances are continuous fields). So, if the actual universe is described by something resembling current theories, it cannot be in one to one correspondence with a standard computer program, which is based on integer mappings. There is currently no model of a "continuous" computer. Turing even argued that such a thing is meaningless! (There are definitions of "computable continuous functions," but none of the definitions are really satisfactory.)

Let's be more broad minded about what is to count as a simulation. Consider the collection of all mathematical concepts . Let us say that a perfect simulation exists if the physical universe can be put into one to one correspondence with some mutually consistent subcollections of all mathematical concepts. In this sense of "simulation" the universe can certainly be simulated, because "simulation" then amounts to saying that the universe can be exhaustively "described" in a logically consistent way. Note that "described" does not require that we or any other finite (or infinite) intelligent being can actually find the description. It may be that the actual universe expands into an infinite hierarchy of levels whenever one tries to describe it exhaustively. In such a case, it would be impossible to find a Theory of Everything. Nevertheless, it would still be true that a "simulation" in the more general sense existed if each level were in one to one correspondence with some mathematical object, and if all levels were mutually consistent ("consistency" meaning that in the case of disagreement between levels, there is a rule — itself a mathematical object — for deciding which level is correct). The crucial point of this generalization is to establish that the actual physical universe is something in the collection of all mathematical objects. This follows because the universe has a perfect simulation, and we agree to identify the universe with its perfect simulation. Thus at the most basic ontological level, the physical universe is a concept.

But of course not all concepts exist physically . But *some do. Which ones? The answer is provided by our earlier analysis of programs. The simulations which are sufficiently complex to contain observers — thinking, feeling beings — as subsimulations exist physically.* And further, they exist physically by definition: for this is exactly what we mean by existence, namely, that thinking and feeling beings think and feel themselves to exist. Remember, the simulated thinking and feeling of simulated beings are real. Thus the actual physical universe — the one in which we are now experiencing our own simulated thoughts and simulated feelings, exists necessarily, by definition of what is meant by existence. Physical existence is just a particular relationship between concepts. Existence is a predicate, but a predicate of certain very, very complex simulations. It is certainly not a predicate of simple concepts, for instance "100 thalers."

With equal necessity, many different universes will exist physically. In particular, a universe in which we do something slightly different from what we actually do in this one will exist (provided of course that this action does not logically contradict the structure of the rest of the universe). But this is nothing new; it is already present in the ontology of the Many-Worlds Interpretation. Exactly how many universes really exist physically depends on your definition of "thinking and feeling being". If you adopt a narrow definition — such a being must have at least our human complexity — then the range of possible universes appears quite narrow: *The Anthropic Cosmological Principle* (Barrow and Tipler 1986) is devoted to a discussion of how finely tuned our universe must be if it is to contain beings like ourselves. (Although the above discussion of existence comes entirely from physics, namely the physics of computer simulation and the MWI, the conclusions are essentially the same as those of the philosophers Plantinga (1974) and Lewis (1986, p. 73) on the necessary existence of all possible worlds. Plantinga and Lewis are motivated by an analysis of the meaning of modal logic.)

What happens if a universal simulation stops tomorrow? Does the universe collapse into non-existence? Certainly such terminating simulations exist mathematically. But if there is no intrinsic reason visible from inside the simulation for the simulation to stop, it can be embedded inside a larger simulation which does not stop. Since it is the observations of the beings inside the simulation that determines what exists physically, and since nothing happens from their view point at the termination point when the terminating simulation is embedded in the non-stopping simulation, the universe must be said to continue in existence. It is the maximal extension which has existence, for by the Identity of Indiscernibles we must (physically) identify terminating programs with their embedding in the maximal program. (One could use a similar argument for asserting the physical existence of the maximal evolution from given initial data in the classical general relativity evolution problem.) Furthermore, if it is logically possible for life to continue to exist forever in some universe, this universe will exist necessarily for all future time.

4. Philosophical Implications of the Existence (?) of the Omega Point

Suppose the Omega Point really exists. Then even on the most materialistic level, the future existence of the Omega Point would assure our civilization of ever growing total wealth, continually increasing knowledge, and quite literal eternal progress. This perpetual meliorism is built into the definition of "life existing forever" given above. Of course, it is a consequence of physics that although our civilization may continue forever, our species *Homo sapiens* must inevitably become extinct, just as every individual human being must inevitably also die. For as the Omega Point is approached, the temperature will approach infinity everywhere in the universe, and it is impossible for our type of life to survive in this environment. (The non-existence of the Omega Point would not help us. If the universe were open and expanded forever, then the temperature would go to zero as the universe expanded. There is not enough energy in the frigid future of such a universe for *Homo sapiens* to survive. Also, protons probably decay, and we are made up of atoms, which require protons.)

But the death of *Homo sapiens* is an evil (beyond the death of the human individuals) only for a racist value system. What is humanly important is the fact we think and feel, not the particular bodily form which clothes the human personality. If the Omega Point exists, the advance of civilization will continue without limit into the Omega Point. Our species is an intermediate step in the infinitely long temporal Chain of Being (Lovejoy 1936) that comprises the whole of life in spacetime. An essential step, but still only a step. In fact, it is a logically necessary consequence of eternal progress that our species become extinct! For we are finite beings, we have definite limits. Our brains can code only so much information, we can understand only rather simple arguments. If the ascent of Life into the Omega Point is to occur, one day the most advanced minds must be non-*Homo sapiens* . The heirs of our civilization must be another species, and their heirs yet another, *ad infinitum* into the Omega Point. We must die — as individuals, as a species — in order that our civilization might live. But the contributions to civilization which we make as individuals will survive our individual deaths. Judging from the rapid advance of computers at present, I would guess that the next stage of intelligent life would be quite literally information processing machines. At the present rate, computers will reach the human level in information processing and integration ability probably within a century, certainly within a thousand years.

Many find the assurance of the immortality of life as a whole cold comfort for their death as individuals. But recall my discussion of Thomist *aeternitas*. I pointed out that all the information contained in the whole of human history, including every detail of every human life, will be available for analysis by the collectivity of life in the far future. In principle at least (again ignoring the difficulty of extracting the relevant information from the overall background noise), it is possible for life in the far future to construct, using this information, an exceedingly accurate simulation of these past lives: in fact, this simulation is just what a sufficiently close scrutiny of our present lives by the beings of the future would amount to. And I emphasized above that a sufficiently perfect simulation of a living being would *be* alive! Whether the beings of the future would choose to use their power to do this simulation, I cannot say. But it seems the physical capability to carry out the scrutiny would be there. Furthermore, the drive for total knowledge — which life in the future must seek if it is to survive at all, and which will be achieved only at the Omega Point — would seem to require that such an analysis of the past, and hence such a simulation, would be carried out.

I should emphasize that this simulation of people that have lived in the past need not be limited to just repeating the past. Once a simulation of a person and his/her world has been formed in a computer of sufficient capacity, the simulation can be allowed to develop further — to think and feel things that the long dead original person being simulated never felt and thought. It is not even necessary for *any* of the past to be repeated. The beings of the future could simply begin the simulation with the brain memory of the dead person as it was at the instant of death (or 10 years before, or 20 minutes before, or ...) implanted in the simulated body of the dead person, the body being as it was at age 20 (or age 70, or ...). This body and memory collection could be set in any simulated background environment the future beings wished: a simulated world indistinguishable from the long-extinct society and physical universe of the revived dead person, or even a world that never existed, but one as close as logically possible to the ideal *fantasy* world of the resurrected dead person. Furthermore, all possible combinations of resurrected dead can placed in the same simulation and allowed to interact. For example, the reader could be placed in a simulation with *all* of his/her ancestors and descendents, each at whatever age (physical and mental, separately) the future beings please. The intelligent beings of the future could interact — speak, say — with their simulated creatures, who could learn about them, about the world outside the simulation, and about other simulations.

The simulated body could be one that has been vastly improved over the one we currently have: the laws of the simulated world could be modified to prevent a second physical death. We could call the simulated, improved and undying body a "spiritual body", for it will be of the same "stuff" as the human mind now is: a "thought inside a mind" (in Aristotelian language, "a form inside a form."; in computer language, a virtual machine inside a machine). The spiritual body is thus just the present body (with improvements!) at a higher level of implementation.[3] Only as a spiritual body, only as a computer simulation, is resurrection possible without a second death: our current bodies, implemented in matter, could not possibly survive the extreme heat near the final singularity.

Although computer simulation resurrection overcomes the physical barriers to eternal life of individual human beings, there remains a logical problem, namely, the finiteness of the human memory. The human brain can store only about 10^{15} bits (Barrow and Tipler 1986, p. 136) [this corresponds to roughly a thousand subjective years of life], and once this memory space is exhausted, we can grow no more. Thus it is not clear that the undying resurrected life is appropriately regarded as "eternal". There are several options. For example, the beings of the future could guide us to a "perfection" of our finite natures. Whatever "perfection" means! Depending on the definition, there could be many "perfections". With sufficient computer power, it should be possible to calculate what a human action would result in without the simulation actually experiencing the action, so the future beings would be able to advise us on possible perfections without us having to go through the trial and error procedure characteristic of this life. If more than one simulation of the same individual is made, then *all* of these options could be realized simultaneously. Once an individual is "perfected", the memory of this perfect individual could be recorded permanently — preserved all the way into the Omega Point. The errors and evil committed by the imperfect individual could be erased (or also permanently recorded). The perfected individual personality would be truly eternal: she would exist for all future time. Furthermore, when the perfected personality reached the Omega Point, it would become eternal in the sense of being beyond time.

In his *On the Immortality of the Soul*, David Hume raised the following objection to the idea of a general resurrection of the dead: "How to dispose of the infinite number of posthumous existences ought also to embarrass the religious theory." (Hume 1755; reprinted in Flew 1964, p. 187). Hume summarized the argument in a later interview with the famous biographer James Boswell: ".. [Hume] added that it was a most unreasonable fancy that he should exist forever. That immortality, if it were at all, must be general; that a great proportion of the human race has hardly any intellectual qualities; that a great proportion dies in infancy before being possessed of reason; yet all these must be immortal; that a Porter who gets drunk by ten o'clock with gin must be immortal; that the trash of every age must be preserved, and that new Universes must be created to contain such infinite numbers." (Hume 1776 [1977], p. 77).

The ever-growing numbers of people whom Hume regarded as trash nevertheless could be preserved forever in our single finite (classical) universe if computer capacity is created fast enough. By looking more carefully at the calculations summarized above, one sees that they also show it is physically possible to save *forever* a certain constant percentage of the information processed at a given universal time. Thus, the computer capacity will be there to preserve even drunken porters, (and perfected drunken porters) provided only that the beings of the future wait long enough before resurrecting them. Even though the computer capacity required to perfectly simulate is exponentially related to the complexity of entity simulated, it is physically possible to resurrect an actual infinity of individuals between now and the Omega Point, even assuming the complexity of the average individual diverges as the Omega Point is approached, and guide then *all* into perfection.[4] Total perfection of all would be achieved at the instant of the Omega Point.

But this preservation capacity has an even more important implication: *it means that the resurrection is possible even if sufficient information to resurrect cannot be extracted from the past light cone.* Since the universal computer capacity increases without bound as the Omega Point is approached, it follows that if only a bare bones description of our current world is stored permanently, then there will inevitably come a time when there will be sufficient computer capacity to simulate our present-day world by simple brute force: by creating a simulation of *all* logically possible variants of our world. For example, the human genome can code about 10 to the 10^6 power possible humans, and the brain of each could have 2 to the 10^{15} power possible memories. With the computer power that will eventually become available, the beings of the future could simply simulate them all. Just the knowledge of the human genome would be enough for this. And even if the record of human genome is not retained until the computer capacity is sufficient, it would still be possible to resurrect all possible humans, just from the knowledge it was coded in DNA. Merely simulate all possible life forms that could be coded by DNA (for technical reasons, the number is finite) and all logically possible humans will necessarily be included.. Such a brute force method is not very elegant;[5] I discuss it only to demonstrate that resurrection is unquestionably physically possible. And if there is no other way, it almost certainly will be done by brute force in the drive toward total knowledge. In our own drive to understand how life got started on our planet, we are in effect trying to simulate — resurrect — all possible kinds of the simplest life forms which could spontaneously form on the the primitive earth.

Eternal worldly progress and the hope of individual survival beyond the grave, usually pictured as polar opposites, turn out to be the same. An interesting feature of the Omega Point Theory is that it provides a plausible physical mechanism for a universal resurrection[6] (for more discussion of the physics of the resurrection see (Tipler 1989)). As Wolfson (1965) has pointed out, resurrection has been inconsistent with the accepted physics of the day for the past two thousand years. That has now changed.

Notes

[1] I am grateful to Frank Birtel, Michael Heller, Wolfhart Pannenberg, and John Wheeler for their comments on an earlier version of this paper. This work was supported in part by the National Science Foundation under grant number PHY-86-03130.

[2] Unfortunately, Aristotle ruined his own idea of the soul by soiling it with Platonic dualism. This mistake led to Aquinas' contradictory notion of "substantial form". Both ideas suggest that the personality survives death naturally. See Flew (1964, pp. 16-21) and 1987, pp. 71-87). Modern physics tells us however that personality does not survive death. When you're dead, you're dead until the resurrection, if it indeed occurs.

[3] See (Hofstadter and Dennett 1981, pp. 379-381) for a very brief discussion of the extremely important computer concept of "levels of implementation."

[4] This depends in a crucial way on the fact that there will be an actual infinity ($\aleph_0$) of information processed between any finite time and the Omega Point. It is an example of what Bertrand Russell (1931, p. 358) has termed the Tristram Shandy paradox. Tristram Shandy took two years to write the history of the first two days of his life, and complained that at that rate, material would accumulate faster than he could write it down. Russell showed that, if Tristram Shandy lived forever, nevertheless no part of his biography would have remained unwritten. In the case of the Omega Point, which literally does live forever, all beings that have ever lived and will live from now to the end of time can be resurrected and remembered, even though the time needed to do the resurrecting will increase exponentially, a much worse case than Tristram Shandy faced. It is important that at any given time on a classical trajectory, there is only a finite number of possible beings which could exist. If this were not true, then the number of beings that would have to be resurrected between now and the Final State might be the power set of $\aleph_0$, which is higher order of infinity than $\aleph_0$, and thus resurrecting all possible beings via the brute force method might be impossible because only $\aleph_0$ bits can be recorded between now and the Final State.

[5] One could also worry about the morality of such brute force resurrection: not only are the dead being resurrected, but also people who never lived! However, the central claim of the Many-Worlds physics and the Many-Worlds metaphysics discussed above is that all people and all histories who could exist in fact do. They just don't exist on our classical trajectory, and so we have no record of them. So the resurrected dead would presumedly not care which classical trajectory they are resurrected in — their own trajectory or another one — so long as they *are* resurrected.

[6] The version of eternal life discussed here is not attractive to everyone. What is happening is that an exact replica of ourselves is being simulated in the computer minds of the far

future. The philosopher Antony Flew, for example, considers it ridiculous to call this "resurrection", and he puts forward the 'Replica Objection': "No replica however perfect, whether produced by God or man, whether in our Universe or another, could ever be — in that primary, forensic sense — the same person as its original." (Flew 1987, p. 12). "To punish or to reward a replica, reconstituted on Judgement Day, for the sins or the virtues of the old Antony Flew dead and cremated, perhaps long years before, is as inept and as unfair as it would be to reward or to punish one identical twin for what was in fact done by the other." (Flew 1987, p. 9). Flew is wrong about our legal system. It does in fact equate identical computer programs. If I duplicated a word processing program and used it without paying a royalty to the programmer, I would be taken to court. A claim that "the program I used is not the original, it is merely a replica" would not be accepted as a defense. I could also be sued for using without permission an organism whose genome has been patented. Identical twins are *not* identical persons. The programs which are their minds differ enormously: the memories coded in their neurons differ from each other in at least as many ways as they differ from the memories of other human beings. They are correctly regarded as different persons. But two beings who are identical both in their genes *and* in their mind programs *are* the same person, and it is appropriate to regard them as equally responsible legally. I am surprised that an empiricist philosopher like Flew would make the claim that entities which cannot be empirically distinguished, even in principle, are nevertheless to be regarded to be utterly different. Any scientist would think that two physically indistinguishable systems are to be regarded as the same, both physically and legally.

References

Abramowicz, M. and Ellis, G.F.R. (1989), "The Elusive Anthropic Principle", *Nature* 337: 411- 412.

Alder, H.L. and Roessler, E.B. (1964), *Introduction to Probability and Statistics*. San Francisco: Freeman.

Barrow, J. D. (1989), "Anthropic Principle", *Nature* 339: 196.

_ _ _ _ _ _ _. and Tipler, F.J. (1986), The Anthropic Cosmological Principle. Oxford: Oxford University Press.

Carter, B. (1989) in *The Anthropic Principle* , U. Curi (ed.). Cambridge: Cambridge University Press.

Feller, W. (1968), *An Introduction to Probability Theory and Its Applications*, Volumes I & II. New York: John Wiley & Sons.

Flew, A. (ed.) (1964), *Body, Mind, and Death*. New York: Macmillan.

_ _ _ _ _. (1984), *God, Freedom and Immortality: A Critical Analysis*. Buffalo: Prometheus.

_ _ _ _ _. (1987), *The Logic of Mortality*. Oxford: Blackwell.

Hawking, S. W. and Ellis, G.F.R. (1973), *The Large Scale Structure of Space-Time*. Cambridge: Cambridge University Press.

Hume, D. (1977), *Dialogues Concerning Natural Religion*. Norman Kemp Smith (ed.). Indianapolis: Bobbs-Merrill.

Hofstadter, D. R. and Dennett , D.C. (1981), *The Mind's I.* New York: Basic Books.

Lewis, D. (1986), *On the Plurality of Worlds.* Oxford: Blackwell.

Linde, A.D. (1989), "Particle Physics and Cosmology", In *Proceedings of the XXIV International Conference on High Energy Physics*, R. Kotthaus and J. Kühn (ed.). Heidelberg: Springer-Verlag.

Lovejoy, A.O. (1936), *The Great Chain of Being.* Cambridge: Harvard University Press.

Plantinga, A. (1974), *The Nature of Necessity.* Oxford: Clarendon Press.

Russell, B. (1931), *Principles of Mathematics.* New York: Norton.

Tipler, F.J. (1976), "Singularities in Universes with Negative Cosmological Constant", *The Astrophysical Journal* 209: 12-15.

_ _ _ _ _ _ . (1979), "General Relativity, Thermodynamics, and the Poincaré Cycle." *Nature* 280: 203-205.

_ _ _ _ _ _ . (1980), "General Relativity and the Eternal Return." In *Essays in General Relativity*, F.J. Tipler (ed.). New York: Academic Press, pp. 21-37.

_ _ _ _ _ _ . (1986), "Cosmological Limits on Computation." *International Journal of Theoretical Physics* 25: 617-661.

_ _ _ _ _ _ . (1988), "The Omega Point Theory." In *Physics, Philosophy, and Theology: A Common Quest for Understanding*, by R.J. Russell, W. Stoeger, and G. Coyne, pp. 313- 331. Notre Dame: University of Notre Dame Press.

_ _ _ _ _ _ . (1989), "The Omega Point as Eschaton: Answers to Pannenberg's Questions for Scientists", to appear in the July issue of Zygon .

Weinberg, S. (1987), "Anthropic Bounds on the Cosmological Constant", *Physical Review Letters* 59: 2607-2610.

_ _ _ _ _ _ _ _. (1989), "The Cosmological Constant Problem", *Reviews of Modern Physics* 61: 1-23.

Wolfson, H. (1965), "Immortality and Resurrection in the Philosophy of the Church Fathers", in *Immortality and Resurrection*, K. Stendahl (ed.). New York: Macmillan, pp 54-96 .