
Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

By Paul Almond

13 December 2010

Website:

<http://www.paul-almond.com>

E-mail:

info@paul-almond.com

There is a view that, by performing repeated quantum suicide experiments, an observer can gain evidence that the many-worlds interpretation of quantum mechanics is true. The idea is that, if the observer has survived enough quantum suicide experiments, this would seem implausible if many-worlds were not true and the observer were relying just on “luck”, while if many-worlds were true, the observer would know that there would have to be some futures in which he/she survived – providing an explanation of his/her current situation. An observer, therefore, should be able to demonstrate to his own satisfaction that many-worlds is true, with any desired degree of confidence, even though onlookers will mainly experience seeing the observer die and the observer can never “come back” to share his knowledge that many-worlds is almost certainly true. This view is incorrect. In fact, repeated quantum suicide experiments will indicate nothing about the likelihood of the many-worlds interpretation being correct, even to the observer who is undergoing them. The justification for this is based on the likelihood of finding yourself experiencing different kinds of observer moments. The observer moments after a quantum suicide experiment would have a lower combined measure than those immediately before it. If many-worlds is true you should consider it just as “unlikely” that you would find yourself in such an observer moment, when you could have found yourself experiencing an observer moment before the experiment, as it would have been to be “lucky” and survive the experiment if many-worlds were not true. An important idea here is that you should not consider your memories as a priori knowledge when considering the likelihood of finding yourself in a particular location in a many-worlds multiverse. The issue of whether subjective survival should be expected in quantum suicide experiments is a separate one.

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

Table of Contents

| | |
|--|----|
| 1 Introduction | 4 |
| 2 The Scenario and the Calculations | 6 |
| 2.1 Description of the Scenario - and its question..... | 6 |
| 2.2 Non-Many Worlds Calculation of the Probability..... | 6 |
| 2.3 Many-Worlds Calculation of the Probability | 7 |
| 2.3.1 Simple Version | 7 |
| 2.3.2 Observer Moment Justified Version | 7 |
| 2.3.3 What about proliferation of observer moments as worlds split? | 9 |
| 2.4 A Further Question and Its Answer..... | 11 |
| 3 Discussion..... | 13 |
| 3.1 Surviving a sequence of quantum suicide experiments does <i>not</i> tell you that many-worlds is likely to be correct..... | 13 |
| 3.2 Does this mean that we can never have reason to think that many-worlds is correct? | 14 |
| 3.3 From where does the surprise come in the MWI case? | 14 |
| 3.3.1 Being Surprised at What Has Happened vs. Being Surprised at Where You Are | 14 |
| 3.3.2 Sequences of Quantum Suicide Experiments..... | 15 |
| 3.4 Does this mean that quantum suicide is just conventional suicide?..... | 17 |
| 3.5 Low Measure Marginalization | 20 |
| 3.6 A Possible Answer to the Doomsday Argument | 21 |
| 3.7 “Silly” Theories | 22 |
| 4 Conclusion..... | 23 |
| 5 References | 26 |

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

List of Abbreviations

MWI many-worlds interpretation (of quantum mechanics)

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

1 Introduction

Quantum suicide is a thought experiment originally proposed by Hans Moravec (Moravec, 1987) and Bruno Marchal (Marchal, 1988), and further developed by Max Tegmark (Tegmark, 1997).

A quantum suicide experiment involves some mechanism that will, depending on some quantum event, either kill the observer or not kill him: we might assume a 0.5 probability of each. The observer can go through a long sequence of such quantum suicide experiments. In the “everyday” view of reality and probability, a long enough sequence of experiments is almost certain to kill the observer: the observer’s survival relies on being “lucky” in each experiment, and the more chances that the observer takes, the more likely it is that his luck will eventually run out.

There is, however, another way of looking at quantum suicide. The many-worlds interpretation of quantum mechanics (MWI), otherwise known as the relative state formulation of quantum mechanics and proposed by Hugh Everett, states that when quantum events occur, all possible outcomes are realized. The quantum wavefunction is viewed as having physical reality, and decoherence causes splitting into separate “worlds” (Everett, 1957; Price, 1995). In this view of reality, the observer in the quantum suicide experiment has multiple futures, and there will always be a future available in which the observer survives.

A common idea is that, if a scientist wants to know whether MWI is true or not, she could use herself as a subject in repeated quantum suicide experiments to find out. The idea is that, after a sufficiently long sequence of experiments, if the scientist is still alive, she should consider it unlikely that this is due to being “lucky” in all the experiments and that, instead, it should be more likely that she is still alive because, in MWI, there is always some future in which she survives. The scientist can supposedly establish that MWI is true with any confidence that she wishes by performing enough quantum suicide experiments: as she survives each latest experiment, the plausibility of a non-MWI explanation decreases and the likelihood of MWI being true increases. In this way, after surviving enough quantum suicide experiments, the scientist can be effectively convinced that MWI is correct. She can never report this back to her colleagues, however: in most of the futures of the scientist’s colleagues, the scientist is simply killed early in the process. Quantum suicide is therefore supposed to be a way by which you can establish that MWI is true *just to yourself*.

This article will show that this idea is incorrect. It will be shown, instead, that if you have survived any number of repeated quantum suicide experiments, despite what we may intuitively think, this gives you no evidence that MWI is true that you did not already have. Quantum suicide experiments are of no use in establishing the likelihood that MWI is correct – not even to the person undergoing them.

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

The argument will be based on considerations of probability, observer moments, measure and locating yourself within the MWI multiverse. To start, we will consider a scenario in which you are faced with uncertainty about what day it is. This scenario will have some similarity with the Sleeping Beauty scenario (Elga, 2000). This scenario is intended to make a point about the likelihood of finding yourself on the other side of a quantum suicide experiment.

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

2 The Scenario and the Calculations

2.1 Description of the Scenario - and its question

It is Sunday, 10PM. You just took a sleeping pill. Before 11PM, you will fall asleep, and you will remain asleep until woken. There is nothing available to wake you up, beyond what is described in this scenario.

A “quantum death alarm clock” is by your bedside. Its behaviour depends on some initial quantum event which occurs at 11:59PM on Sunday and which sets its future behaviour in one of three ways.

- *There is a 0.5 probability that the quantum death alarm clock will ring and wake you on Monday.*
- *There is a 0.49 probability that the quantum death alarm clock will kill you, silently, in your sleep, on Monday.*
- *There is a 0.01 probability that the quantum death alarm clock will ring and wake you on Tuesday.*

(and $0.5+0.49+0.01=1$)

You fall asleep and the next thing you are aware of is waking up as you hear the quantum death alarm clock ringing.

Without looking at a calendar, what is your estimate of the probability that it is Tuesday?

Two approaches to estimating this probability will be considered. The first will be assuming conventional probability, and the second will be assuming that MWI is correct.

2.2 Non-Many Worlds Calculation of the Probability

You know that the quantum death alarm clock may be ringing on Monday, which has a probability of 0.5. You also know that it may be ringing on Tuesday, which has a probability of 0.01. There is a probability of 0.49 that it does neither – that it kills you before you awake – but the fact that you have now been woken by the clock’s ringing means that you can discount this possibility. The probability that it is Tuesday is the probability that the clock rings on Tuesday as a proportion of all the total probability of you remaining alive

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

$$\begin{aligned} P(\text{It is Tuesday}) &= \frac{P(\text{Clock rings on Tuesday})}{P(\text{Clock rings on Monday}) + P(\text{Clock rings on Tuesday})} \\ &= \frac{0.01}{0.5 + 0.01} \\ &= \frac{1}{51} \end{aligned}$$

So, the probability that it is Tuesday is $\frac{1}{51}$.

2.3 Many-Worlds Calculation of the Probability

2.3.1 Simple Version

You know that in some paths, the quantum death alarm clock will have killed you without ringing, but as the machine is ringing and you are now waking, you know that this has not happened: you know that you are still alive.

In your past, either the machine must have been set by the quantum event to ring and wake you on Monday, which has a probability of 0.5, or it must have been set by the quantum event to ring and wake you on Tuesday, which has a probability of 0.01. You can, of course, discount the situation in which the machine was set to kill you, as you know that, from your perspective, this did not happen.

Therefore, as with the non-MWI situation

$$\begin{aligned} P(\text{It is Tue}) &= \frac{P(\text{Clock rings on Tue})}{P(\text{Clock rings on Mon}) + P(\text{Clock rings on Tue})} \\ &= \frac{0.01}{0.5 + 0.01} \\ &= \frac{1}{51} \end{aligned}$$

So, the probability that it is Tuesday, *for you*, is $\frac{1}{51}$.

Assuming that MWI is correct has done nothing to make the probability that it is Tuesday any greater.

2.3.2 Observer Moment Justified Version

We need first to ensure that we can approach probabilities in MWI correctly.

Let us consider the structure of the multiverse described by MWI, in particular looking at the branches extending out from the instant when the quantum death alarm clock is

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

set by the quantum event. Because we are assuming that MWI applies, we are no longer thinking that one future or another happens: they all happen. This means that the question “What is the probability that it is Tuesday?” is a bit vague. There will be multiple versions of you, some of which meet death from the quantum death alarm clock, some of which wake on Monday and some of which wake on Tuesday: there will definitely be someone who has the experience of waking on Tuesday. What the question means, when we ask you after you have just woken, is this:

Where are you, right now, in the Everett multiverse? What is the probability that you are in a part of the multiverse corresponding to waking on Tuesday?

The way to approach this is to look at *how much multiverse there is* for each possibility. We can do this by considering the branches extending out from the instant at which the quantum event occurs and the quantum death alarm clock was set, and the proportions of those branches that correspond to different futures.

- The probability that the quantum death alarm clock was set to kill you was 0.49. In MWI terms, this means that the proportion of the total measure it receives was 0.49: the proportion of the branches extending out from the instant when the quantum death alarm clock was set is 0.49.
- The probability that the quantum death alarm clock was set to wake you on Monday was 0.5. In MWI terms, this means that the proportion of the total measure it received was 0.5: the proportion of the branches extending out from the instant when the quantum death alarm clock was set to wake you on Monday is 0.5. We can think of this as meaning that the proportion of futures in which you wake on Monday is 0.5.
- The probability that the quantum death alarm clock was set to wake you on Tuesday was 0.01, and this similarly corresponds, in MWI, to a proportion of the total measure of 0.01, meaning the proportion of branches from the instant at which the clock was set is 0.01, and the proportion of the futures in which you wake on Tuesday is 0.01.

Out of all the “places” where you might be in the multiverse, the proportion of those places in which you are dead is 0.49, the proportion of those places where you are waking on Monday is 0.5 and the proportion of places in which you are waking on Tuesday is 0.01. You can discount the possibility that you are in one of the places in the multiverse where you are dead: you know that you are alive. This means that you must be in one of the places in the multiverse where you are waking on Monday or one of the places in the multiverse where you are waking on Tuesday. The probability that you are in one of the places in the multiverse corresponding to waking on Tuesday is therefore obtained using the proportions of places in each case.

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

$$\begin{aligned} P(\text{You are in a Tue waking}) &= \frac{\text{Proportion of Tue wakings}}{\text{Proportion of Mon wakings} + \text{Proportion of Tue wakings}} \\ &= \frac{0.01}{0.5 + 0.01} \\ &= \frac{1}{51} \end{aligned}$$

So, as previously calculated, the probability that it is Tuesday, for you, is $\frac{1}{51}$.

Some readers will recognize the sort of reasoning that we are doing here as being about *observer moments*. This is reasoning similar to that in the Doomsday argument, a controversial argument suggesting that, if human civilization will survive for a long time into the future, we are now part of an extraordinarily small proportion of people living very early in history, and that as this is unlikely it is more plausible that our civilization will not survive long (Carter, 1983; Gott, 1993). Nick Bostrom has described how the Doomsday argument can be viewed in terms of observer moments – the sort of thing with which we are dealing with here (Bostrom, 2005).

Again:

Assuming that MWI is correct has done nothing to make the probability that it is Tuesday any greater.

2.3.3 What about proliferation of observer moments as worlds split?

The above justification, in 2.3.2, only gave consideration to the branches coming immediately from the instant at which the quantum death alarm clock is set. It was assumed that, for example, because the Monday awakening gets 0.5 of these branches, this means that the Monday awakening gets 0.5 of the total measure (including worlds where you end up dead). However, splitting of worlds happens after this instant. In fact, all through Monday worlds will be splitting and proliferating. By Tuesday, a huge amount of splitting should have occurred. Even if only $\frac{1}{51}$ of the initial branches in which you are going to survive were for futures in which you wake on Tuesday, should it not be the case that the extra time that elapses before the time for waking on Tuesday should cause so much splitting and proliferation of worlds that almost all of the observer moments in which you wake are Tuesday ones?

We can approach this issue by assuming that it is the case in everyday life and seeing where it gets us. Suppose that when worlds split, new observer moments are created with the same statistical status, so that you should expect it to be as likely to find yourself experiencing one observer moment as any other, regardless of where it is in the multiverse.

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

This means that if you are experiencing Observer Moment A, and your world splits into two worlds – one world in which you are experiencing Observer Moment B₁ and another world where you are experiencing Observer Moment B₂ – finding yourself “in” Observer Moment B₁ should seem no more unusual than finding yourself “in” the earlier Observer Moment A.

If MWI is true, all through your life so far, worlds have been splitting and, if the above view is correct, observer moments have been proliferating *with no change in the statistical status of a single observer moment*. If you find yourself experiencing an observer moment before the age of ten years old, you should be very surprised because (presumably) the continued proliferation of worlds and observer moments after your tenth birthday party will mean that there are vastly more observer moments in your life after your tenth birthday party than before it. In fact, the proliferation of worlds within a fraction of a second after your tenth birthday should in itself create more observer moments than you had experienced so far.

Let us now take this further. In the last fraction of a second of your life, on any branch, enough proliferation of worlds will occur to generate more observer moments than you have experienced so far in your life, so you should consider it natural to find yourself in one them, and very unlikely to find yourself anywhere else: you should expect to find yourself, right now, within a fraction of a second of death – and if you are not, you should consider yourself to be experiencing a hyper-unusual observer moment. We can take this further still. After your life has ended, worlds will continued to proliferate and other people’s observer moments will exist in those worlds. There should be many more observer moments, say, for people in general, one hundred years after your death than there are at the time of your life, because the proliferation of worlds should have multiplied everything up. You should not merely expect to find yourself a fraction of a second from death, but a fraction of a second from the end of your civilization – or, maybe even someone else’s later civilization. Really, you should expect to find yourself in the last fraction of a second of existence of the last person in the last civilization, because the proliferation of worlds means that that is where all the observer moments should be.

The scenario just described, of course, does not happen. I am not within a fraction of a second of the end of my life or someone else’s later life, the end of this civilization, or someone else’s life at the end of a later civilization. Instead, I am experiencing what seems to be a relatively normal observer moment – and you almost certainly are too. This should tell us that the above view is flawed.

The view must be flawed in its assumption that, when worlds split, statistical measure proliferates too – that when a world splits into two worlds, a single observer moment in one of these two worlds has the same statistical measure – in that you should expect it to be as likely to find yourself experiencing it – as an observer moment in the original world, just before the splitting. Assuming this leads to the nonsense conclusion that we

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

just obtained. To be experiencing anything like a “normal” observer moment at all, without things being statistically skewed one way or another, we need some view in which you are no more or less likely to find yourself at any particular time in your life than at any other time: the splitting of worlds must not count statistically. For this to be the case, statistical measure must not proliferate when worlds split: instead, when a world splits, the worlds that split off from it get the same amount of measure in *total*, so that each individual split-off world has less measure than the world from which it split off.

Previously, we considered the situation in which you are experiencing Observer Moment A, and your world splits into worlds – one world in which you are experiencing Observer Moment B₁ and another world in which you are experiencing Observer Moment B₂. With this correct view, the total measure of B₁ and B₂ is the same as that of A. Finding yourself “in” Observer Moment B₁ or Observer Moment B₂ should seem as likely as finding yourself in Observer Moment A, but finding yourself in Observer Moment B₁ should seem less likely than finding yourself in Observer Moment A, because it has less measure.¹

Viewing the measure of observers as being divided up like this when worlds split should also fit better with our view of what is *physically* going on.

2.4 A Further Question and Its Answer

The question previously asked, in 2.1, was:

Without looking at a calendar, what is your estimate of the probability that it is Tuesday?

It has been shown that the answer should be $1/51$, regardless of whether MWI is correct.

We can now ask a further question.

If you wake up, and look at the calendar and see that it is Tuesday, how does this affect your estimate of the probability that MWI is correct?

The answer should be that *it does not affect it at all*. If the probabilities of it being Tuesday that you calculated before you looked at the calendar were very different in the MWI and non-MWI cases, then finding out that it is Tuesday could be very significant.

¹ In this view, all observer moments are not the same. There is, however, an alternative way of thinking about it, and that is that all observer moments *are* the same, and that what we think of as an observer moment at some instant actually corresponds to many observer moments, and when worlds split, the number of observer moments for each instant in the split-off worlds is reduced, so that the total number of observer moments across all worlds remains constant. This would actually be complicated by the need to deal with things more in a relative, than in an absolute way, but in any event it is really just a matter of semantics. The idea of observer moments having different measure will be used here.

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

For example, if you calculated that the probability of it being Tuesday if MWI were *not* true was 10^{-10} and the probability of it being Tuesday if MWI were true was 0.9 then finding out that it is Tuesday would very strongly suggest that MWI is true: the alternative would be to think that MWI was false and, instead an outcome with a probability of 10^{-10} had just occurred. This, however, is not the case: before you looked at the calendar, your probability of it being Tuesday should have been the same in either case, and when you find out that it is Tuesday the news should be no more or less surprising if MWI is true than if it is not true. *Finding out that it is Tuesday tells you nothing about the probability that MWI is true.*

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

3 Discussion

3.1 Surviving a sequence of quantum suicide experiments does *not* tell you that many-worlds is likely to be correct.

If you start performing quantum suicide experiments on yourself, it is generally accepted that people around you will tend to see you getting killed. This means that nothing is proved to these people. Even if one of them sees you survive, he has no particular reason to think that is because of MWI.

There is a common view, however, that things are different for you: that if you have survived a sufficiently long sequence of quantum suicide experiments, you will be experiencing very uncommon observer moments that are supposed, almost certainly, not to exist without MWI and that this tells you that MWI is almost certainly true. You are supposed to be able to increase your confidence that MWI is true to as high a level as you wish merely by performing more quantum suicides. While you can never demonstrate this to the people you “left behind”, who are left looking at your corpse, the idea is that at least *you* will have some satisfaction from knowing.

The thought experiment just discussed has shown that *this view is wrong*. In the thought experiment, finding yourself in an observer moment that is part of a category of observer moments with a combined low measure, following a quantum suicide event told you nothing. It follows, unless there is any reason to think otherwise, that this generalizes, and that finding yourself in a low measure set of observer moments after any quantum suicide event, or any sequence of quantum suicide attempts, would tell you nothing – beyond the obvious facts that if MWI is not true, you have survived by being lucky, and if MWI is true you happen, right now, to be on one of the branches where you survived.

There is a common idea of the lone scientist, wanting to know if MWI is true, embarking on a series of quantum suicide experiments. The idea is that the scientist can confirm that MWI is true with any degree of confidence required, but can never bring this knowledge back: to anyone else living in the world, in almost all of their futures, the scientist merely gets killed.

The argument in this article suggests that this idea is flawed: Not only will onlookers in the laboratory be denied evidence that MWI is true: even the scientist himself is given no rational reason for preferring MWI as a result of the experiments.

All this means that a quantum suicide experiment tells you nothing, really, beyond the specifics of the experiment itself. (If MWI is not true it tells you what happened in the experiment and if MWI is true it tells you where you after the experiment.)

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

3.2 Does this mean that we can never have reason to think that many-worlds is correct?

No, it merely means that surviving repeated quantum suicide experiments is going to tell you nothing about the truth or otherwise of MWI. There may be other reasons for thinking that MWI is true or false, and these are not affected by the inability of quantum suicide experiments to tell you anything: if you have a good reason to think that MWI is true *before* attempting a quantum suicide experiment, then this reason is just as good afterwards. Max Tegmark has discussed various reasons for preferring an MWI view or a non-MWI view (Tegmark, 1997).

3.3 From where does the surprise come in the MWI case?

3.3.1 Being Surprised at What Has Happened vs. Being Surprised at Where You Are

The results of the thought experiment may seem counter-intuitive to some people, because of the way that your prior calculation of the probability that it is Tuesday, before finding out that it is Tuesday, is the same in both the MWI and the non-MWI cases: you are supposed to have found out that an outcome with the same probability has occurred in each case, so you are supposed to be equally surprised in each case. The issue, here, is one of where the surprise is supposed to come from in the case where MWI is true.

In the case where MWI is *not* true, it is obvious where the surprise comes from. The experience of waking on Tuesday is unlikely to happen. It is much more likely ($^{99}/_{100}$ probability) that you will be killed or woken on Monday instead. Your surprise therefore conforms to our common idea of what surprise is supposed to be: you are surprised that an unlikely event happened. We can also think of this in terms of observer moments. When you wake on Tuesday, you are not experiencing an observer moment that was unlikely ever to exist: when you find out that the observer moment did exist, and you experienced it, you should be surprised that it existed. The important point here is that the surprise is about what *reality* turns out to be like.

In the case where MWI is true, the surprise cannot come from the same source. You should not be surprised because you wake on Tuesday: MWI tells you that different versions of you wake on different days, so some version of you *must* wake on Tuesday. Looking at this in observer moment terms, an observer moment exists in which you wake on Tuesday, so you cannot reasonably be surprised to find out that such an observer moment exists. There can, therefore, be no surprise about what reality turns out to be like: MWI is a completely deterministic model and you already knew what reality was going to turn out to be like. At a casual glance, it may seem that MWI takes away all of our valid causes for surprise, suggesting that, maybe, when you wake on

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

Tuesday it should seem quite natural. *The mistake, here, would be thinking that the surprise in the MWI case has to come from the same source as the surprise in the non-MWI case.*

In the MWI case, the low probability of finding yourself experiencing waking on Tuesday did not come from the unlikelihood of such observer moments existing, nor did it come from the probability of being dead: that was ignored in the calculation, as you can readily eliminate it on waking by observing that you are not dead. The low probability came from the low combined measure of observer moments corresponding to waking on Tuesday when compared with the higher combined measure of all the observer moments corresponding to waking on Monday. You should have expected to find yourself in the higher measure set of Monday waking observer moments instead – or even in an observer moment at some completely different point in your life, before all this messing around with quantum death alarm clocks started. Your surprise, on finding out that you were experiencing an observer moment corresponding to waking on Tuesday, does not come from finding out that such an observer moment exists, but rather *from finding out that you are in it*. It is the same kind of surprise you would experience if you were driving and got lost and found yourself in an unusual place.

3.3.2 Sequences of Quantum Suicide Experiments

In the thought experiment that we have been considering, there is a reasonable chance of survival: the machine may simply wake you on Monday, rather than kill you. The way in which we have dealt with this scenario, however, tells us how to deal with the more common quantum suicide scenario in which someone is supposed to demonstrate to herself that MWI is true by performing a long sequence of quantum suicide experiments such that, if MWI is not true there is hardly any chance of survival. Here, we need to be careful, because cognitive errors can cause people to jump to the wrong conclusions. Let us consider a scenario.

Suppose that Mary, a physicist, wants to know if MWI is correct, and she is prepared to put her life on the line to find out. She sets up a machine which runs a long sequence of quantum suicide experiments, each of which kills her or does not kill her, based on some quantum event, with 0.5 probability. She survives the first few experiments. At this stage, she could just be lucky, so she continues. Eventually, she has been through 1,000 quantum suicide experiments and is still alive. If MWI is not correct, the probability of this happening was $(\frac{1}{2})^{1,000}$, or about 10^{-301} , and this is clearly an extremely low probability: Mary should be surprised, as it was much more likely that she would be dead by now. It may seem obvious that Mary should find it implausible that such a low probability sequence of events has occurred, and that she should instead conclude that MWI is true.

The problem with such reasoning is that it ignores the fact that, if MWI is correct, when Mary finds herself in this observer moment after surviving 1,000 experiments, it is part

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

of a category of observer moments with an extremely low combined measure: the measure has been “thinned” time after time by the repeated quantum suicide experiments. If MWI is correct, Mary should be just as surprised to find herself in such an observer moment, out of all the observer moments she could be in, as she would be to find that she had survived a long sequence of experiments if MWI were incorrect. The situation is no more or less implausible with MWI being true than it was without MWI being true. Any implausibility associated with an unlikely sequence of events happening has merely been replaced by implausibility associated with being located in a low measure part of the Everett multiverse.

Some people will object to this by saying that the MWI situation is not the same as the non-MWI one, because if the sequence of experiments has happened, the only place that Mary *could* be is in an observer moment in which she has survived, as she can only experience these observer moments. The fallacy here is that this requires that Mary regard her knowledge of her place in the multiverse as *a priori* knowledge. When Mary has survived 1,000 quantum suicide experiments, she will have memories of the 1,000 experiments happening and the passage of time while they happened. It may seem tempting to say that, given that Mary’s memories tell her that 1,000 experiments have happened, it should seem obvious to her that the only place she could be is in an observer moment where she *survived* the 1,000 experiments, but the mistake here is in assuming that Mary’s *memory* of the 1,000 experiments counts as *a priori* knowledge. Her memory is specific to where she is and is just a feature of the observer moment in which she happens to be. When considering other observer moments in which she *might* have been, in order to determine the likelihood of finding herself in her current one, Mary should not assume that her memory of past events must be the same in all the other observer moments: she could have found herself in an observer moment from earlier in time, with a memory of only taking part in ten quantum suicide experiments, or she could have found herself in an observer moment from before the quantum suicide experiments started, with no memory of performing any at all.

This deals with the objection that someone who has found herself in an observer moment corresponding to having just survived a long series of quantum suicide experiments, with a supposedly very small chance of survival, could *only* find herself in such an observer moment: she could have found herself in a much *earlier* observer moment, and the fact that such observer moments will have higher measure means that she should have expected to find herself in such an observer moment if she was unaware of the date or what had been going on in her past. In the non-MWI version of a situation like this, the surprise, and any implausibility, would come from not being dead when you should have been. When MWI is correct, the surprise, and any implausibility, would come from looking at the clock, the calendar *or your own memory* and seeing that you here *now*.

The kind of reasoning used here has some similarity with the kind of reasoning used to support the Doomsday argument.

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

3.4 Does this mean that quantum suicide is just conventional suicide?

Some people think that MWI may provide a “loophole” in some situations in which your death is *almost* certain. The idea is that, if you perform a sequence of quantum suicide experiments, like Mary in the above example in 3.3.2, the fact that there will be some version of you that survives – that there will always be observer moments in your future in which you survive – means that you can discount the futures in which you do not survive, as you are not there to observe them, and view yourself as being certain to survive. If this is true, going through a sequence of quantum suicide experiments should hold no fears for you. This, of course, is a controversial idea. Some people think it is nonsense, even if MWI is correct, while others think that the way in which your life ends matters. A related idea is that of quantum immortality – the idea that MWI would guarantee our survival in all circumstances. Max Tegmark is sceptical of the idea that MWI would imply quantum immortality, suggesting in an e-mail in 1998 (available on his personal website) that dying is likely to involve a period of classical deterioration (Tegmark, 1998). In the same e-mail Tegmark mentions a possible objection to these ideas, stating

“I agree that if the argument were flawless, I should expect to be the oldest guy on the planet, severely discrediting the Everett hypothesis.”

The argument made in this article should provide a simple answer to this. Regardless of what we think of the your reasonable expectations of “subjective survival” in these kinds of situations, any observer moments associated with what we would traditionally think of as low probability outcomes, in MWI become members of a low measure collection of observer moments, and therefore, if you had somehow lost track of how old you were – say you had severe amnesia – you should find it unlikely that you were in one of the observer moments in a category with a low combined measure, if these exist, corresponding to being, for example, millions of years old. If you found out that you were in such a situation, and had arrived in it not through some advance in technology or other “conventional” means, but simply through reliance on MWI always continuing you in some branches, you should be extremely surprised.

So, does the argument in this article imply that it is flawed to think of MWI as providing you with guaranteed, subjective survival in quantum suicide experiments? It may seem at first glance that it is saying that – that it is saying that the future observer moments, after you have performed quantum suicide, are somehow “inaccessible” – in the sense that you should never consider it likely to end up in them.

This, however, is not really what the argument is saying. The idea that you should consider it unlikely to find yourself in a particular observer moment is relevant purely in the contexts of not knowing the time and/or date and considering the likelihood that, when you do find out the time and/or date, it will be after some sequence of quantum

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

suicide experiments, or of finding yourself in a given kind of observer moment and deciding how plausible it is that you are there given some theory of physics. The idea that you should consider it unlikely to find yourself in a given observer moment does not necessarily imply that you should not expect your future to be made up of such observer moments. The reason for this is that, when we talk about future expectations, the matter is complicated with various issues of *continuity*. Any observer moment that you expect to be in five minutes from now has some kind of relationship with the observer moment that you are in now: it is not randomly selected from the set of all the observer moments in your life. To be “an observer moment in which you could be in five minutes” an observer moment has to meet the standard of being five minutes from now at the very least, and most observer moments in your life will not meet such a standard. This does not mean that measure is irrelevant: measure will still be an issue when assigning subjective probabilities to the various futures that are valid continuations from your present, but this is very different from saying that a future observer moment is unlikely to be experienced due to the combined low measure of observer moments of that kind in comparison with the measure of observer moments in your past.

This leads to a situation which may seem paradoxical, but would only really seem that way due to our unfamiliarity with it. Let us consider the scenario from 3.3.2 of Mary, the physicist taking part in a sequence of quantum suicide experiments. If MWI does provide guaranteed subjective survival in situations like this it would mean that the following two statements could be made that may seem contradictory about Mary’s expectations as she starts the sequence of experiments.

1. Mary should expect, after some time, to be in a future in which she has survived the sequence of quantum suicide experiments, as a consequence of MWI.
2. When Mary finds herself in this future, she should view it as a very unlikely situation, and not one which is “predicted” any more by MWI than it is by a non-MWI view.

If we assume that quantum suicide does allow guaranteed subjective survival, MWI seems to be telling Mary to expect to be in a future situation which, when she is in it, she should regard as unexpected! There is no conflict here. Our expectations for future observer moments are statistically skewed because we will only regard certain observer moments as being “the future”. However, when we get into an observer moment, what we thought of this observer moment from the point of view of an “earlier” observer moment in the past has nothing to do with how unusual it should seem to be there. (Of course, another possibility is that Mary should not expect subjective survival at all: that what she is doing is little better than *conventional* suicide – and that issue is not going to get resolved in this article.)

As an analogy, suppose there is a man who lives to be a billion years old, and who knows the day of his death in advance. Suppose he loses track of the date completely, as well

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

as losing his memory, so that he does not know his age at all. He then looks at a calendar and sees that it happens to be the day before he is due to die. He should be surprised at this: it locates his current experiences in a very unusual part of his life – right at the end. This does not mean that he is experiencing a low measure or unlikely observer moment: rather, the observer moment he is experiencing is in an unusual part of his life. Now, the man knows his situation is unusual, but he will also expect his experiences five minutes from now to be those of a man *still* near the end of his life, and just five minutes nearer to death. The man knows that his current situation is “unlikely” – in the sense that he did not expect to be in it before he checked the date. He will also know that the situation in which he will be in five minutes will be unlikely for the same reason; however, this does not mean that he will expect to be magically “teleported” to some more plausible time, earlier in his life. There is a requirement for what can constitute “his future” and he expects to experience it, from the point of view of his present, no matter how statistically unlikely it is that anyone would experience it.

This may still seem paradoxical to some readers, who would wonder how your expectation that something is going to happen could rationally turn into surprise that it has happened. This brings us back to the point that you should not regard your memories as a priori information when deciding how likely your situation is. Before the sequence of quantum suicide experiments, Mary may think that MWI provides subjective survival in quantum suicide experiments, and she may have an expectation that she will find herself in the situation of having survived, but that says nothing about the likelihood of this kind of future observer moment in any general sense: it is purely how it looks from Mary’s particular position, pre-quantum suicide. The idea that the kind of observer moment corresponding to having survived the sequence of quantum suicide experiments is “expected” is more a characteristic of the observer moment that Mary is in pre-quantum suicide, combined with a particular view about what constitutes valid continuity, than it is about the later observer moment itself – and, of course, Mary could be very wrong to expect this. It may be that she should really expect simply to die. When Mary finds herself in the observer moment corresponding to having survived, she will remember that she expected to survive before she did the experiments, but that information will be of no use to her in deciding how unlikely her current observer moment is, because she cannot regard her memories as a priori: the fact that she remembers having that expectation in the past is no longer an expectation, but is just an arbitrary feature of her current observer moment. Mary should now reflect that the observer moment corresponding to being in this particular situation with these particular memories exists as part of a category of observer moments with extremely low measure, and she should be surprised to find herself in it. Her memory of expecting this observer moment to happen will be of no use in deciding whether MWI is likely to be true, because if MWI is true the issue now becomes one of why she has the memories that she does.

I want to be clear, here, that none of this means that MWI does or does not provide subjective survival in quantum suicide scenarios: the issue is too involved to deal with

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

fully in this article. I am merely pointing out that the argument given here does not resolve this issue of subjective survival one way or another.²

3.5 Low Measure Marginalization

It has been shown how, if you find yourself in the situation of having the memory of going through a series of quantum suicide experiments, such that survival would be expected to be extremely unlikely, MWI does not offer an easy way out of explaining the apparent low probability of you being in this situation: both the non-MWI view and the MWI view seem to require that you accept a high degree of specificity in your situation.

Survival of repeated quantum suicide experiments will not help you choose between MWI and non-MWI views, but you may already have a view on the subject, which does not change as the number of survived experiments increases. Alternatively, you may not think you can decide between MWI and non-MWI views. What happens, however, as the number of quantum suicide experiments that you have survived increases, and the probability that you would normally assign to surviving this far, or to finding yourself in this situation eventually decreases to an ultra-low value? This low probability will be the same whether MWI is true or whether the alternative explanation – that the quantum suicide experiments simply turned out the right way in a non-MWI world – is true. It will not, therefore support either position and, without any competing view, we should still think that MWI with the observer being appropriately located in the multiverse by chance, or non-MWI with the right sequence of events just happening by chance, is true. However, what if some competing hypothesis is proposed to explain your situation? One hypothesis might be that you could be delusional. Another hypothesis could be that the experiments are somehow rigged by other humans or aliens or even that your current experiences are really the product of a computer simulation by some entity which could be human or nonhuman. Such ideas may seem extreme, but the problem with a long sequence of quantum suicide experiments is that if you think of a weird hypothesis to explain your situation, and you can put a non-zero probability on it, if the probability you would normally assign to being in your situation has become low enough, this “weird” hypothesis will actually appear to be a more sensible explanation than the more “conventional” explanations that you just got lucky in a non-MWI reality or just happen to be in the right observer moment in an MWI reality: it will actually be more rational to believe it!

² There is, actually a very small contribution that the argument here may make to this issue: earlier, I mentioned Tegmark’s comment about expecting to be the oldest person in the universe. The fact that you are not such a person might be considered to be an argument against MWI. Because this issue is resolved by the consideration we have given to the matter here, a possible reason for some suspicion of MWI is removed. While MWI being true does not necessarily imply subjective survival in quantum suicide experiments, it is a prerequisite for it, so some people may consider dealing with this issue as slightly increasing the chance that MWI does provide subjective survival.

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

I discussed this issue in a previous article (Almond, 2008) and referred to it as “marginalization”. The term “marginalization” is not really very descriptive, so it may be better to refer to it as “low measure marginalization.

Low measure marginalization would seem to be something that anyone, or any entity, contemplating quantum suicide should take seriously – and it is a threat even if MWI ensures subjective survival. It does not matter that when you start, with a high measure, you might have no reason to take various weird hypotheses seriously: in the future, when you are in an outrageously low measure group of observer moments, you will have every reason to doubt the correct explanation of your position. Your memories of surviving the sequence of quantum suicide experiments should count for nothing, because they could be more rationally explained in some other way – and the fact that such an explanation would be wrong will not help you in this situation, when *you do not know*. From this perspective, repeated quantum suicide might be considered as “burning the evidential bridge” between you and your previous situation, leaving you uncertain of your status. It might be considered a threat to the sanity of an entity performing it.

The issue of whether you could do anything about this is an entirely separate one, that will not be discussed in detail here – and of course it will only matter anyway if quantum suicide does involve subjective survival. Any defence against low measure marginalization would probably need to be information theoretic in nature somehow, and involve taking something with you – some package of information – which somehow cannot easily be accounted for by outlandish explanations. That in itself raises issues of what you can know and what your basis for knowledge is, and before we end up in the territory of Hume and Kant I will leave this issue for now.

3.6 A Possible Answer to the Doomsday Argument

The Doomsday argument, if correct, suggests that our civilization will end soon, based on the idea that our own position in time should not be too unusual, but in an MWI context, the Doomsday argument could be interpreted as meaning that our civilization will have a huge and ongoing decrease of *measure* in the future: that our civilization will continue in some futures, but that its measure will decrease all the time as something like the quantum suicide experiment – whether contrived by our descendants or imposed on them – occurs on a continual basis. In such a situation, observers should expect to find themselves in relatively high measure parts of the civilization’s history – maybe where we are now.

How useful this is should depend on your views of MWI and subjective survival. If you find MWI implausible, or if you do not think that MWI would provide subjective survival in quantum suicide experiments, it will be of no help. On the other hand, if you think that MWI is true and that you should view yourself as certain to survive quantum suicide experiments, it should possibly be encouraging.

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

3.7 “Silly” Theories

A number of “non-standard” (for that you might read “weird stuff on the Internet”) theories about our status have in common the idea that the end of the world should be expected to happen on a regular basis, and that we survive it purely because MWI demands that some future always exists in which it does not happen. Two examples are as follows.

- Repeated technical failures in physics experiments, most notably at CERN, combined with claims that the experiments, if they actually proceed, will cause the end of the world, have caused speculation that has already happened and continues happening – that CERN regularly causes the end of the world in a substantial number of futures, and we do not know about it because we have to be in a branch where this did not happen – which means we see technical failures instead of experiments.
- We live in a false vacuum which collapses into a low energy state on a weekly basis – or we could think of other similar “cosmic catastrophe” ideas in which the world is supposed to be unstable and we regularly get “quantum suicided” without asking for it.

The argument given in this article should have shown that theories like this are untenable. The theories may *seem* like good explanations of things like technical failures at particle accelerators. (That is being generous actually: the idea that CERN can end the world is hardly plausible.) However, it should be apparent that the theories are not explaining anything. The idea that CERN would kill us when it is turned on, and that our continued existence is explained by MWI is no more plausible than the idea that it is explained by a lucky series of technical problems – and it should be clear that the latter is extremely implausible.

It should be noted that this does *not* apply to anthropic arguments that suggest that it should have been unlikely for the universe to develop in a form that allowed life, or that the universe should have been too unstable to allow life, and that MWI has dealt with this in the past by ensuring that a path was available through the multiverse that allowed our universe to become a stable one, suitable for life, before conscious observers existed. Without observers, or observer moments in the universe before the measure reduction occurs, there is no reason not to think that observer moments after the measure reduction are places in which observers should think it quite reasonable to find themselves. To put this another way, the thought experiment about waking on Monday or Tuesday, in 2.1, is irrelevant if you do not have any observer moments appearing before Tuesday.

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

4 Conclusion

There is a view that, by performing repeated quantum suicide experiments, an observer can gain evidence that the many-worlds interpretation of quantum mechanics (MWI) is true. The idea is that if the observer has survived enough quantum suicide experiments, this would seem implausible if MWI were not true and the observer were relying just on “luck”, while if MWI were true, the observer would know that there would have to be some futures in which she survived – providing an explanation of her current situation.

This view has been shown to be incorrect. Surviving any number of quantum suicide experiments does *not* tell an observer anything about the likelihood that MWI is correct.

The argument to show this started with a thought experiment involving a “quantum death alarm clock”. This would wake you on Monday, kill you or (much less likely) wake you on Tuesday, depending on the outcome of some quantum event. The problem is to determine the probability that it is Tuesday when you have just woken and have not yet seen a calendar. This probability should not depend on whether MWI is true or not. An argument based on observer moments was given to justify this. If MWI is true, the number of branches from the initial quantum event that sets the quantum death alarm clock that lead to Tuesday awakenings is much lower than the number of branches that lead to Monday awakenings. The combined measure of all the observer moments corresponding to a Tuesday awakening is therefore lower than it is for a Monday awakening. On waking you should therefore think it is much more likely that you are experiencing a Monday awakening.

An objection to this could be that worlds proliferate in MWI as time passes. Are we correct just to consider the branches leading from the initial quantum event that sets the quantum death alarm clock? As time passes, these branches will divide into further branches many times. As Tuesday is further in the future from Monday, and more splitting will have occurred by the time for the awakening on Tuesday, will this not mean that there will actually be more observer moments for a Tuesday awakening? The answer to this is that as worlds split, the measure of observer moments is reduced so that, when branching occurs, assuming that the observer survives in all branches, the combined measure of the observer moments in all branches immediately after the branching is the same as the measure of an observer moment immediately before the branching: observer moments proliferate, but the total measure of observer moments does not.

Some people may say that this is untrue, or cannot be known with anything like certainty, but your experience of the world, and your place in it, should tell you that this is the case. If observer moments did not have reduced measure after splitting, the observer moments very near the end of your life would outnumber, and statistically outweigh, all the previous observer moments and you should therefore be very surprised to find yourself living at any time other than in the last fraction of a second of

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

your life – or in fact, the last fraction of a second of the life of someone else who lives later than you. The only way of resolving this issue is for the measure of future observer moments to be reduced as worlds split such that it exactly balances the proliferation of observer moments.

When we are not at risk of death, the likelihood of us being in one moment or the next remains the same, as the proliferation of worlds exactly balances the reduction in measure as worlds split, but in the thought experiment, on some branches you cease to exist and have no further observer moments, meaning that the total measure of observer moments in which you wake on Tuesday is much less than total measure of observer moments in which you wake on Monday. Assuming that MWI is true makes no difference to the estimate that you should make of the probability of a Tuesday awakening. This also works in reverse: the day on which you wake tells you nothing about whether MWI is true. This is an important result, because it is a specific quantum suicide experiment in which the outcome tells you nothing about the likelihood of MWI being true.

We can view repeated quantum suicide as reducing the overall measure in the same way. Quantum suicide reduces your total measure of observer moments across worlds that correspond to being alive at some particular time following a sequence of experiments. Finding yourself alive after a sequence of quantum suicide experiments may *seem* to suggest that MWI is correct, due to the improbable sequence of events that have to be assumed to have occurred, but if you assume that MWI is true, you also have to assume that you are experiencing an extremely low measure kind of observer moment – and this should appear no more or less implausible than the idea that an extremely unusual sequence of coincidences happened in a non-MWI reality. One way of considering this is in terms of “surprise”. If you survive repeated quantum suicide in a non-MWI reality, you should be surprised at *what happened*. If you survive repeated quantum suicide in an MWI reality you should be surprised at *where you are*. The degree of surprise should be the same in each case, and there will be no way to distinguish between the two views.

There is a similarity between this argument and Bostrom’s “observer moment” way of viewing the Doomsday argument, which suggests that we should view ourselves as typical observers, rather than observers who are experiencing unusual observer moments.

Although repeated quantum suicide can tell you nothing about whether MWI is true, this does not mean that you can never know whether MWI is true or not. You may know in other ways, such as from physics experiments beyond the scope of this article or from philosophical arguments. It merely means that none of your knowledge can come from repeated quantum suicide experiments.

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

A survivor of repeated quantum suicide experiments should consider the issue of *low measure marginalization*. If you have survived many quantum suicide experiments, such that the measure for the kind of situation in which you are, relative to that of other situations you might have been in earlier in your life, is very low, the plausibility of any “conventional” explanation of your situation – whether MWI or not – may be much lower than the plausibility for some strange explanation – such as the idea that you are delusional or in a computer simulation.

One issue with all this may be the idea of subjective survival. Should someone who is about to undergo a sequence of quantum suicide experiments fear death, or should he eliminate any futures in which he ceases to exist from his consideration, and in which he cannot observe the outcome, only considering those futures in which he is alive to make observations? Some people may think that the argument in this article suggests that it is wrong to think you will survive, as the reduction in measure means you should find it less likely to be in an observer moment on the other side of a quantum suicide experiment. This would be taking the argument in this article too far, however. Although finding yourself existing after a long series of quantum suicide experiments should be considered unlikely, this is assuming that all times in your life are treated equally. In reality, we require that events must satisfy some criteria to be considered “part of our future”, and this will reduce the set of possible events that we would consider part of our future. As an analogy, we might imagine a very old person in an observer moment near the end of his life. He may find it very unusual and improbable to be in such an unusual kind of observer moment, and he should consider it just as unusual to find himself in a slightly later observer moment in five minutes, slightly closer to his death, but this does not mean that he will not expect to find himself in this slightly later observer moment, and that he will expect to be “magically teleported” back to some earlier time in his life. We can also make an analogy with the Doomsday argument. Advocates of the Doomsday argument suggest that it is very unusual to find yourself at such an early period of history, and that this suggests that the world will end soon; however, I am unaware of anyone seriously suggesting that, in the absence of such an end of the world, you should not expect to be experiencing the “next” observer moment, but should expect to find yourself in the “more probable” distant future: we do not view our future as being “selected” like that. We should be careful, therefore, of reading too much into considerations of measure like the one in this article when it comes to future expectations. It remains the case, however, that when you actually find yourself experiencing an observer moment after some sequence of quantum suicide experiments, you cannot reasonably view your memories of events as priori knowledge, and you should consider the likelihood that you are experiencing this moment, now, with these memories. As has been shown, with MWI whatever degree of likelihood you find here will be no different from the likelihood of surviving the sequence of quantum suicide experiments in the absence of MWI. Quantum suicide experiments will therefore indicate nothing about the likelihood that the many-worlds interpretation of quantum mechanics is true – not even to the person undergoing them.

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

5 References

- Almond, P., 2008. *Civilization-Level Quantum Suicide*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/CivilizationLevelQuantumSuicide.pdf>, <http://www.paul-almond.com/CivilizationLevelQuantumSuicide.doc> or <http://www.paul-almond.com/CivilizationLevelQuantumSuicide.htm> [Accessed 9 December 2010].
- Bostrom, N., 2005. *Self-Location and Observation Selection Theory*. [Online] The Anthropic Principle. Available at: <http://anthropic-principle.com/preprints/self-location.html> [Accessed 8 December 2010].
- Carter, B., 1983. The anthropic principle and its implications for biological evolution. *Philosophical Transactions of the Royal Society of London*, A310, pp.347-363.
- Elga, A., 2000. Self-locating belief and the Sleeping Beauty problem. *Analysis*, 60(2), pp.143-147. (Also available at: <http://www.princeton.edu/~adame/papers/sleeping/sleeping.html> [Accessed 10 December 2010]). (The problem was discussed earlier in unpublished work by Arnold Zuloff.)
- Everett, H., 1957. Relative State Formulation of Quantum Mechanics. *Reviews of Modern Physics*, 29, pp.454-462.
- Gott, J. G. III, 1993. Implications of the Copernican principle for our future prospects. *Nature*, 363, pp.315-319.
- Marchal, B., 1988. Mechanism and Personal Identity. *Proceedings of WOCFAI 91* (Paris. Angkor), pp.335-345. (Also available at: http://iridia.ulb.ac.be/~marchal/publications/M&PI_15-MAI-91.pdf [Accessed 10 December 2010].)
- Moravec, H., 1988. The Doomsday Device. *Mind Children: The Future of Robot and Human Intelligence*. Harvard: Harvard University Press. p.188. (Also available at: <http://books.google.com/books?id=56mb7XuSx3QC&lpg=PP188&pg=PA188> [Accessed 10 December 2010].)
- Price, M. C., 1995. *The Many-Worlds FAQ*. [Online] The Anthropic Principle. Available at: <http://www.anthropic-principle.com/preprints/manyworlds.html> [Accessed 12 December 2010]. (Also available at: <http://www.hedweb.com/everett/everett.htm> [Accessed 12 December 2010] and <http://kuoi.com/~kamikaze/doc/many-worlds-faq.html> [Accessed 12 December 2010].)

Repeated quantum suicide would *not* suggest that MWI is correct – even to the person doing it.

Tegmark, M., 1997. *The Interpretation of Quantum Mechanics: Many Worlds or Many Words?*. [Online] arXiv:quant-ph/9709032. Available at: <http://arxiv.org/abs/quant-ph/9709032> [Accessed 8 December 2010]. (Also available at: http://xxx.lanl.gov/PS_cache/quant-ph/pdf/9709/9709032v1.pdf [Accessed 12 December 2010].)

Tegmark, M., 1998. *Quantum Suicide*. [Online] The Universes of Max Tegmark. Available at: <http://space.mit.edu/home/tegmark/crazy.html> [Accessed 9 December 2010]. (The relevant text here was originally an e-mail sent on Sat Nov 28 13:20 EST 1998, and has been reproduced by Max Tegmark on the page in this link, which is on his personal website. The webpage mentioned here is undated, so I have used the year that Tegmark gives for the e-mail.)