
Minds, Substrate, Measure and Value

Part 5: Cryptic Ontology

By Paul Almond

6 December 2010

Website:

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

E-mail:

info@paul-almond.com

This is the fifth in a series of articles about the nature of reality, the nature of consciousness and the relationship between minds and the rest of the world. The previous articles were *Minds, Substrate, Measure and Value – Part 1: Substrate Dependence*, available at <http://www.paul-almond.com/Substrate1.pdf>, *Minds, Substrate, Measure and Value – Part 2: Extra Information About Substrate Dependence*, available at <http://www.paul-almond.com/Substrate2.pdf>, *Minds, Substrate, Measure and Value – Part 3: The Problem of Arbitrariness of Interpretation*, available at <http://www.paul-almond.com/Substrate3.pdf>, and *Minds, Substrate, Measure and Value – Part 4: The Cosmological Many-Interpretations View*, available at <http://www.paul-almond.com/Substrate4.pdf>. These argued that, to have a coherent view of probability, a measure view of consciousness is needed, in which an observer exists with some measure that is substrate dependent. An implication of this is that we must accept that all interpretations of reality which “find” an observer find a *real* observer, and that an observer exists in any situation with a degree of measure that depends on the substrate and the amount of information needed in an interpretation for it to find that observer. This was extended to be a cosmological position, implying that we are in a Type IV multiverse: a proposal with some similarities with Max Tegmark’s *mathematical universe* proposal. In this article, a further argument to support this view is given, based on the idea of determining probabilities when an observer is in one of two possible situations, one of which involves progressive encryption of the observer. It is argued that, in such a situation, probability becomes incoherent unless we adopt a measure and many-interpretations view of consciousness, and that this supports the case made so far in this series.

Table of Contents

1 Introduction	3
2 Scenarios Involving Progressive Encryption.....	5
2.1 The Cryptic Ontology Debate at <i>Machines Like Us</i>	5
2.2 Scenario 1: An Encrypted Observer	5
2.2.1 The Scenario.....	5
2.2.2 Discussion	5
2.3 Scenario 2: Encrypted Pornography	6
2.3.1 The Scenario.....	6
2.3.2 Discussion	6
2.4 Is the question about Fred meaningful?	6
2.5 Scenario 3: Probabilistic Version of the Encrypted Observer.....	7
2.5.1 The Scenario.....	7
2.5.2 Discussion	8
3 Conclusion.....	10
4 Acknowledgements.....	11
5 References	12

1 Introduction

This is the fifth in a series of articles exploring the relationship between minds and physical systems (substrates) on which they are based. The previous articles in this series are as follows.

Minds, Substrate, Measure and Value, Part 1: Substrate Dependence. 2007.

(Available at: <http://www.paul-almond.com/Substrate1.pdf> or, <http://www.paul-almond.com/Substrate1.doc>.)

Minds, Substrate, Measure and Value, Part 2: Extra Information About Substrate Dependence. 2007.

(Available at <http://www.paul-almond.com/Substrate2.pdf> or <http://www.paul-almond.com/Substrate2.doc>.)

Minds, Substrate, Measure and Value, Part 3: The Problem of Arbitrariness of Interpretation. 2008.

(Available at <http://www.paul-almond.com/Substrate3.pdf> or <http://www.paul-almond.com/Substrate3.doc>.)

Minds, Substrate, Measure and Value, Part 4: The Cosmological Many-Interpretations View. 2010.

(Available at <http://www.paul-almond.com/Substrate4.pdf> or <http://www.paul-almond.com/Substrate4.doc>.)

In the first three articles, an argument was made for the *many-interpretations view*: the idea that all interpretations of a physical system should be treated equally, and that any interpretation, no matter how extreme, that describes a mind should be regarded as corresponding to a *real* mind. The fourth article widened the scope of the argument and proposed the *cosmological many-interpretations view*. The many-interpretations view, as described previously, had made a strange kind of special case of minds. The cosmological many-worlds view treats all of reality – all objects – in the same way, so that any object, and not just a mind, that can be found by some interpretation of a physical system actually exists. The cosmological many-interpretations view, therefore, receives some of its support from the many-interpretations view, of which it is a generalization. Extra support for it also comes from generalizing and formalizing our existing concept of *emergence*.

The argument for the cosmological many-interpretations view, with its multiverse implications, was discussed in a debate between me and Steve Grand (who was sceptical of the argument) on the *Machines Like Us* website in 2009, and this is available at <http://www.machineslikeus.com/forum/cryptic-ontology/> (Grand & Almond, 2009). In that debate I presented an argument for the many-interpretations view, and by implication the cosmological many-interpretations view, that involved progressive

encryption of minds. In previous articles, I have used thought experiments involving reference class, in which you could be in a number of different situations, and are about to make an observation, with the results of the observation being different in each situation. It was argued that, in such thought experiments, if you assume that simplistic functionalism applies, there is no coherent way for you to assign probabilities. The only viable way of resolving this issue, and making probability coherent, is to take a measure view of consciousness, which leads to the many-interpretations view. The argument from progressive encryption of minds uses reference class in the same way: it sets up another situation, this time involving an observer who is encrypted to varying degrees, in such a way that probability becomes incoherent. This article is about this argument.

2 Scenarios Involving Progressive Encryption

2.1 The Cryptic Ontology Debate at *Machines Like Us*

Part of the argument I am making is based on the arbitrariness of the computational interpretations required in the classic functionalist position. In a discussion with Steve Grand on the *Machines Like Us* website, which Steve Grand gave the appropriate title *Cryptic Ontology* (a title which is so appropriate I stole it for this article), I gave two scenarios involving progressive encryption (Almond & Grand, 2009). The first of these involved an encrypted observer, while the second involved encrypted pornography.

2.2 Scenario 1: An Encrypted Observer

2.2.1 The Scenario

The first progressive encryption scenario which I gave in the *Cryptic Ontology* debate was as follows.

"Fred has had himself uploaded into a computer. He has himself encrypted to hide. He has himself more strongly encrypted the next day, and the next day still more strongly encrypted. As this goes on it becomes harder to find Fred. The only way you could find Fred is by doing a really complex interpretation of the computer system. Eventually, this interpretation would be so extreme it is even more complex than Fred himself! Has Fred ceased to exist? I would say no - as far as Fred is concerned, nothing changes. However, to use Steve's kind of language he has more or less moved "into another universe". He is separated from us not by light years or wormholes but by something more extreme - a kind of gulf of complexity of interpretation. We could say "Fred is hiding in the computer - if you analyse it deeply enough you will find him!" but this may seem an act of faith to some people - and the same reasoning could be used to say that the Mona Lisa, or a simulation of another universe is hiding in the matter of my desk. As far as Fred is concerned, reality is going on as normal and Fred could apply complex interpretations to the structure of his own world that describe other objects, and so on. If anyone objects to this my question is simple: at what point does Fred's encryption become so extreme that he ceases to exist - and why?"

We should assume, here, that the encryption is of a type such that, as it becomes more extreme, a longer computer program is required to "find" the observer.

2.2.2 Discussion

In this scenario, an observer is encrypted to a progressively greater degree. This raises the issue of when the observer stops being an observer. Does he/she remain conscious

during such a process? If you think that the observer does remain, the problem is that when the observer is very extremely encrypted, evidence of the observer could only be found by using a very extreme interpretation – but this would start to get very similar to the way in which a very extreme interpretation could be used to “find” an observer when there is no observer even there. On the other hand, if you think that the observer ceases to exist during this process – *when?* When we have done a lot of encryption? How much? When we have done any encryption at all? The problem with that is that encryption of something, merely means that some algorithm will need to be run to find it, and some kind of interpretative algorithm will always be needed – even for an object which is standing there in front of your eyes – so anything could be said to be “encrypted” to some minimal degree. The arbitrariness of all this suggests that a *measure* view is needed.

“Uploaded” characters actually hide in this way in Greg Egan’s novel, *Permutation City* (Egan, 1994).

2.3 Scenario 2: Encrypted Pornography

2.3.1 The Scenario

The second progressive encryption scenario which I gave in the *Cryptic Ontology* debate was as follows.

“The government declares all pornography a work of Satan, and people can be imprisoned for possessing it - even encrypted pornography - no matter how strongly encrypted - and regardless of whether they possess the means to decrypt it. Anyone who really thinks about this issue would see that it would make it illegal to own a copy of Microsoft Word, the Bible, a desk, or a meat pie - because by applying a complex enough logical manipulation to these items you could "decrypt" anything you wanted. Any sensible administration would have to impose some kind of practical limit on what they meant by saying one thing was stored encrypted in another system - and it should be obvious that such a limit would be arbitrary and not about anything profound.”

2.3.2 Discussion

As well as encryption of minds, we can consider encryption of other things, and such consideration suggests that the issues raised by encryption of minds extend to them too. This is similar to the case made in the fourth article of this series that a generalization of emergence naturally leads to the cosmological many-interpretations view, and is further support for it.

2.4 Is the question about Fred meaningful?

Of the two scenarios, I think that the first one, in 2.2, about Fred’s progressively more deeply encrypted mind, is the more interesting. Some people may object to this

scenario by saying that it is asking a pointless question, and that we can never really know whether Fred, or some computational instantiation of him, is conscious or not. I disagree with this. We should have a view of consciousness which allows us to have some idea whether or not a given system is conscious. Nevertheless, it is possible to make a variation of the progressive encryption of minds scenario which makes it an issue of probability and reference class, rather than one of wondering if someone else is conscious. This approach should be familiar to readers of the previous articles in this series, and it will be obvious that my intention is to show that, when we do this, and do not acknowledge the issues of *measure of observers* and *many interpretations* being possible, our view of probability becomes incoherent – yet when we do acknowledge these issues, yet more support is given for the argument built in this series so far – that we live in a “mathematical”, Type IV multiverse, using the classification system proposed by Tegmark (Tegmark, 2003). I will now provide this altered, probabilistic version of the first scenario.

2.5 Scenario 3: Probabilistic Version of the Encrypted Observer

2.5.1 The Scenario

The first scenario, that of the encrypted observer, in 2.2 can be extended to make a probabilistic version of it, an example of which is as follows.

You know that your thoughts now are based on one of two computer programs, each running in a different computer, Computer A or Computer B. Each program is running your mind in a virtual reality. The two computers are physically the same – apart from with respect to the differences in software mentioned here (which will, of course, have to manifest themselves as physical differences between the machines). You may have got into this situation because your brain was scanned into software on the two computers, or you may have always existed on them as a result of being a manufactured AI system. How you got into this situation is irrelevant. The main point is that the only way you know of that your thoughts could be caused is as a result of programs running in A or B.

You are standing in front of a table (in the virtual reality). There is a closed box on the table. In the program running on Computer A, when you open the box, you will find a red ball inside the box. In the program running on Computer B you will find a blue ball inside the box.

This is obviously a question of what you think the probability is that you will see a red ball. Clearly, with no reason to distinguish between Computer A or Computer B, you should say that the probability is $1/2$. We now go further.

We start to progressively encrypt the software running on Computer B, just as in the earlier example from my debate with Steve Grand, so that it gets progressively harder to see that there is a program running a mind in a virtual reality. At any stage in the

progression of this encryption you will know that your experiences are due to one of two computers, A or B, but the software on computer B is encrypted to some degree, making it harder to see that it is there.

At the start, the software running on computers A and B is the same, apart from the difference in ball colour. After a small amount of encryption, it is just a bit more complex to determine what the software on computer B is doing, and then the software on computer B is encrypted a bit more, and so on.

What happens to your probabilities now?

2.5.2 Discussion

One objection to the first scenario, in 2.2, of the encrypted observer, could be that it *does not matter* whether an observer continues to exist or not while being progressively encrypted— that the question may be meaningless or unanswerable and that we need not concern ourselves with it. This scenario deals with this, using the method that has been used previously in this series of turning the issue into one of *reference class*.

What happens to the probabilities as the encryption of Computer B's software becomes more extreme? Does the version in Computer B continue to be a candidate for your situation? Does it continue to have half or your assignment of probability? Does the probability of seeing a red ball stay the same, no matter how deeply encrypted the "blue ball" program is? Does it decrease as the encryption increases? Does the encryption on Computer B reach some point at which it is not a possible situation for you, and the probability of seeing a red ball abruptly goes from $\frac{1}{2}$ to 1 with some small increase in the encryption on Computer B?

If you think that the chance of seeing a red ball is initially $\frac{1}{2}$, because there are two programs – two possible situations – from which to choose, and it stays at $\frac{1}{2}$ as the encryption in Computer B gets more extreme, until the encryption in Computer B abruptly becomes too extreme, then you have the problem of declaring an arbitrary cut-off point – some arbitrary point at which an interpretation is so extreme that it is not valid. There is no "hard" philosophical reason for choosing this cut-off point and where to put it would be a subjective decision. The problem with this is that this subjective decision is now affecting your chance of being in a particular situation. If, when the encryption has reached a certain level, you think there is effectively no program running on Computer B that could be a candidate for your situation then you will assume your experiences are entirely dependent on what is happening on Computer A and that you have a 100% chance of seeing a red ball when you open the box. However, someone who thinks the encryption in Computer B has not become so extreme that you can eliminate it as a possibility, and that it is equally relevant, will say that you have a $\frac{1}{2}$ probability of seeing each ball. In the absence of any objective way of choosing a cut-off point, neither of you has any authority and probability is incoherent. *With even one*

possible situation in which probability becomes incoherent, probability as a whole is incoherent. This approach is therefore untenable.

If you think the chance of seeing a red ball is initially $\frac{1}{2}$, and stays at $\frac{1}{2}$, no matter how extreme the encryption of the program on Computer B is, you are taking the position that you would have to regard Computer B as running a certain program even when only a very contrived interpretation would reveal this. This is equivalent to accepting that extreme interpretations which produce programs are valid, which is not too far from what I am saying, although I do not agree that the probabilities stay the same.

If you think that the chance of seeing a red ball starts at $\frac{1}{2}$, with the programs on Computer A and Computer B making equal contributions to the reference class of possible situations, but that as the degree of encryption on Computer B increases, it becomes progressively less likely that you are in that situation, and progressively more likely that you will see a red ball when you open the box, you are effectively adopting something like the many-interpretations position. If you do not think there is some cut-off point, when the probability of seeing a blue ball suddenly goes to zero, you are accepting that any degree of encryption on Computer B can still correspond to an observer that has some statistical significance, and you now have the problem that observers could be found by equally extreme interpretations when nobody has encrypted anything, because you can find anything in anything with a sufficiently extreme encryption.

The point of all this is that the only way for probability to be coherent in this situation, and give anything like usable results, is to apply some kind of measure approach to the consciousness of observers when encryption like this is involved, with the measure of an observer being related to the amount of information needed in an algorithm to find the observer. Once we have done this, naïve functionalism has been found inadequate again, and we have gone, once more, down the road of accepting many interpretations, in as much as they relate to observers. This is further support for the *many-interpretations view*, discussed in the first three articles of this series, and therefore, by implication the *cosmological many-interpretations view* which follows from it.

3 Conclusion

An argument has been given that, when encryption of a mind is occurring with varying degrees, probability becomes incoherent. This is shown by means of a scenario in which you could be in one of two different situations, where a different outcome of some experiment is expected in each, and each situation involves you existing on a different substrate, with progressive encryption occurring with the version of you on one substrate, meaning that your presence becomes increasingly obscured. You are asked for an estimate of the probability that the experiment has a particular outcome. If you think that the encrypted version of you does not feature in your reference class when computing this probability, you have the problem that encryption could be almost not there – and when occurring to a small degree would be indistinguishable from the way in which some processing is routinely needed to determine the existence of objects anyway. If you think that the encrypted version of you features in your reference class no matter how extreme the encryption is, you have the problem that extreme interpretations could “find” some version of you that could be said to be encrypted in practically anything. The only way out of this mess is to adopt some kind of measure view of consciousness and to accept that this measure decreases when the length of the algorithm needed to describe an interpretation that finds some version of you increases. This is yet more support for the many-interpretations view discussed in the first three articles of this series.

The fourth article of this series argued for the *cosmological* many-interpretations view. The many-interpretations view was about many interpretations and minds, but the cosmological many-interpretations view extended it to be about things in general. Part of the argument for the cosmological many-interpretations view was based on formalization and generalization of our existing concept of emergence, but part of it was also based on the need to generalize the many-interpretations view, so as to avoid making a special case of minds, with the implication that a Type IV multiverse results, in which all interpretations that describe objects describe real, physical objects. Such a view has some similarities with the mathematical universe proposed by Max Tegmark (Tegmark, 1998, 2007). Because part of the support for the cosmological many-interpretations view comes from the many-interpretations view, the support that the argument about progressive encryption of minds gives to the many-interpretations view also extends to support for the cosmological many-interpretations view, strengthening the case that we are in the kind of multiverse discussed in the fourth article of this series.

Objections can be made to the progressive encryption of minds argument and answers will be given later to these, as well as to other objections that can be made against the argument in this series in general.

4 Acknowledgements

Michael Fridman has been helpful in discussions relating to this subject. Steve Grand, OBE was helpful, by taking part in e-mail discussions, in 2009, which ultimately became the *Cryptic Ontology* debate with me (<http://www.machineslikeus.com/forum/cryptic-ontology>). The title, “Cryptic Ontology”, which Steve Grand gave to that debate has been reused by me, here. Numerous people who have taken the time to make objections to the ideas in this series on <http://www.lesswrong.com> have also been helpful in showing the objections that are likely to be made and where defence of the argument is needed.

5 References

(The first four references, below, are for the previous articles in this series.)

Almond, P., 2007. *Minds, Substrate, Measure and Value, Part 1: Substrate Dependence*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/Substrate1.pdf> or <http://www.paul-almond.com/Substrate1.doc> [Accessed 30 October 2010].

Almond, P., 2007. *Minds, Substrate, Measure and Value, Part 2: Extra Information About Substrate Dependence*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/Substrate2.pdf> or <http://www.paul-almond.com/Substrate2.doc> [Accessed 30 October 2010].

Almond, P., 2008. *Minds, Substrate, Measure and Value, Part 3: The Problem of Arbitrariness of Interpretation*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/Substrate3.pdf> or <http://www.paul-almond.com/Substrate3.doc> [Accessed 30 October 2010].

Almond, P., 2010. *Minds, Substrate, Measure and Value, Part 4: The Cosmological Many-Interpretations View*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/Substrate3.pdf> or <http://www.paul-almond.com/Substrate3.doc> [Accessed 29 November 2010].

Egan, G., 1994. *Permutation City*. London: Millennium. (Fiction).

Grand, S. & Almond, P., 2009. *Cryptic Ontology*. [Online] Machines Like Us. Available at: <http://www.machineslikeus.com/forum/cryptic-ontology> [Accessed 31 October 2010]. (A debate between me and Steve Grand, in which the idea of a multiverse based on all possible interpretations, and the justification for it, was discussed. Steve is sceptical of all this.)

Tegmark, M., 1998. Is the theory of everything merely the ultimate ensemble theory? *Annals of Physics*, 270, pp.1-51. (Also available online at: http://arxiv.org/PS_cache/gr-qc/pdf/9704/9704009v2.pdf [Accessed 4 September 2010]).

Tegmark, M., 2003. Parallel Universes. *Scientific American*, May 2003, pp.40-51. (Also available online at: http://space.mit.edu/home/tegmark/PDF/multiverse_sciam.pdf [Accessed 5 December 2010].)

Tegmark, M., 2007. The Mathematical Universe. *Found.Phys*, 38, pp.101-150. (Also available online at: <http://arxiv.org/abs/0704.0646> [Accessed 5 December 2010]).