

---

# The self-indication assumption almost stops the Doomsday argument. Almost. V2.0

By Paul Almond

---

10 January 2011

Website:

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

E-mail:

[info@paul-almond.com](mailto:info@paul-almond.com)

The Doomsday argument proposes that we can have a reasonable statistical expectation about the remaining lifespan of our civilization from our own position in history. According to the Doomsday argument, we should regard ourselves as typical observers, and it should be unlikely that we are extraordinarily early observers in history and more likely that our civilization will not survive for a long period of time. Many people consider the most promising objection to the Doomsday argument to be the self-indication assumption. A successful rebuttal of the Doomsday argument by the self-indication assumption would show that it is more likely that we live in a world which contains more observers due to our civilization lasting for longer. In a previous article, we gave an argument showing that the SIA does not work as a rebuttal of the Doomsday argument. This argument made reference to impersonal possible worlds. An improved argument is given here, which only makes reference to observer-centred worlds: impersonal possible worlds are not needed. The relationship between observers in observer-centred possible worlds and the statistics of your situation is analyzed. Your statistical situation is completely described by the set of observer-centred possible worlds centred on you, and the self-indication assumption would need to imply that the population or population density in a type of observer-centred possible world has an influence on the distribution of this set. This is shown to be the case for a general version of the self-indication assumption in which the observers' situations are not specified in detail, but it is not the case for the specific form of the SIA needed to refute the Doomsday argument.

# Table of Contents

1 Introduction .....	4
2 The Ultimate Reference Class .....	5
2.1 The Need for Something More Basic .....	5
2.2 The Ultimate Reference Class .....	5
2.3 The Superiority of the Ultimate Reference Class.....	6
3 A Practical Consideration: You cannot see all of the ultimate reference class. ....	8
4 Relating the Self-Indication Assumption (SIA) to the Ultimate Reference Class.....	10
4.1 The Self-Indication Assumption (SIA) must tell us about the distribution of you-centred possible worlds. ....	10
4.2 Secondary Reference Classes.....	10
4.3 An Analogy: Chessboard Worlds.....	11
4.4 What determines the degree of representation of an impersonal possible world in the secondary reference class? .....	14
4.5 Relating the Secondary Reference Class to the Ultimate Reference Class .....	17
5 Where does this leave the self-indication assumption (SIA)? .....	19
5.1 Does the number of observers in an observer-centred possible world matter? ...	19
5.2 Does the number of observers in the <i>future</i> in some impersonal possible world matter?.....	19
5.2.1 The Requirement for Observer-Centred Worlds .....	19
5.2.2 The Problem with Future Observers - Why the Self-indication Assumption (SIA) Fails to Refute the Doomsday Argument .....	21
6 Further Considerations .....	23
6.1 A Short Justification of the Doomsday Argument .....	23
6.2 Infinite Worlds and Level IV Multiverses .....	25
7 Conclusion .....	27
8 Acknowledgements.....	31
9 References .....	32

## List of Abbreviations

bit	binary digit
SIA	self-indication assumption

# 1 Introduction

The Doomsday argument is 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).

One objection to the Doomsday argument was developed by Dennis Dieks (1992) and is known as the self-indication assumption (SIA) – a name given to it by Nick Bostrom. The self-indication assumption argues that you should consider a civilization in which your own existence was likely more plausible than one in which it was not (Grace, 2010). This is supposed to offset the Doomsday argument and cancel it out.

Many people consider the SIA to be the leading candidate for a rebuttal of the Doomsday argument. This article will be arguing that hoping for the SIA to refute the Doomsday argument is somewhat optimistic. It will be shown that there are some grounds for accepting a limited version of the SIA, but this will not help the SIA to refute the Doomsday argument: in fact, the reverse applies. The version of the SIA which will be shown to be valid will be too weak to refute the Doomsday argument, but by showing that it has some limited validity we will have a situation in which we will have dealt with any intuitive expectation that something like the SIA should apply – and shown that it is inadequate for dealing with the Doomsday argument. This actually makes it *less* likely that there is a stronger version of the SIA out there, waiting to refute the Doomsday argument, as the intuitive support we might have for such an idea has already been examined and found to support something inadequate. Further, the analysis of the relationship between observers in possible worlds and the statistics of your possible situation that we will perform in all this will give enough understanding of the role that observers in possible worlds play in the statistics of your situation to show that it is clearly implausible to fit a version of the SIA which refutes the Doomsday argument into this anywhere. The SIA, as a rebuttal of the Doomsday argument, will be shown to fail.

This is a revised version of a previous article (Almond, 2011). This version is an improvement on the previous version because it dispenses with the idea of impersonal possible worlds – an idea that was relied on in the previous article.

## 2 The Ultimate Reference Class

### 2.1 The Need for Something More Basic

The SIA attempts to refute the Doomsday argument by suggesting that you should think in terms of “all the people who could have been you”. The problem with any such consideration is that it needs some justification. When we start to consider issues of “what could have been”, and the role that “what could have been” should play in any assessment of our current situation, if there is such a role, we need something more basic to which we can relate all this. We need some more basic method of assessing probabilities, at least in principle, which is free of built-in assumptions about “what could have been”. We can then see if considerations about “what could have been” are justified by such a method – which will ultimately come down to the question of whether they help us to apply it. The method we need is based on observer-centred possible worlds.

### 2.2 The Ultimate Reference Class

You know that you exist in reality as an observer. You know things about your local reality – for example, you will know whether, right now, you are looking at a red car or not – but there are things you cannot be sure of – for example, what is happening on distant planets that we have never observed directly or, more specific to this article, whether you are living in a civilization that will last for a long time. There may even be some big questions about the nature of local reality – about which you are not sure. The reason that you cannot be sure is that, at any time, you can only ever have a finite amount of information about your situation, and the type of information that you can have will be restricted according to your situation, while there is an infinity of possible situations: you can never have enough information to select one situation while ruling out an infinity of alternatives.

Suppose you make a list of every possible situation in which you could exist. Each entry in the list is a formal description of some possible situation in which you could be. The situations may vary a lot, but they all have one feature in common: each is consistent, in every detail, with what you know. Each of these descriptions is a possible world, but it specifically describes the world in relation to you, so it is a very specific kind of observer-centred possible world: it is a *you*-centred possible world.

An important feature of this list is that every you-centred possible world in it is a candidate for your situation. The list consists purely of observer-centred possible worlds that *might* be yours, and it contains the descriptions of all such possible worlds. It does not contain any observer-centred possible worlds that “could have been” yours but are not: if you could look at one of these observer-centred possible worlds and realize that it cannot possibly be a candidate for your situation, it is not a you-centred possible world and has no place on the list.

This list of you-centred possible worlds, then, gives you a set, or reference class, of all the possible situations in which you could be. We will call this the *ultimate reference class* and define it as follows.

### **The Ultimate Reference Class**

The ultimate reference class is the set of all formally expressed descriptions of you-centred possible worlds. A you-centred possible world is an observer-centred possible world that is a candidate for your situation, and it is consistent with everything that you know about your situation.

## **2.3 The Superiority of the Ultimate Reference Class**

You will never be able to see all of the ultimate reference class: it contains an infinity of you-centred possible world descriptions and if you programmed some computer to construct all these descriptions, it could never finish in any finite time. Nevertheless, there is a way in which we can meaningfully discuss “access” to the ultimate reference class.

Your situation will correspond to one of the you-centred possible worlds in the ultimate reference class, and there is no justification for preferring one such possible world over another. We might say that each you-centred possible world is equally likely to be yours, but it does not make sense to talk about the probability of selecting a single item from an infinite set. To make the set manageable, in the sense of being something on which you can perform statistical operations, suppose that we limit the set to containing only those you-centred possible worlds with descriptions containing  $n$  bits (binary digits) or less of information, and suppose that you can make  $n$  as large as you want. As  $n$  tends to infinity, this set will start to resemble the infinite set of the ultimate reference class. You can now meaningfully discuss the probabilities of being in various situations. For example, to obtain the probability that it will rain tomorrow, for the set of all you-centred possible worlds with some maximum description length,  $n$ , you can count the number of you-centred possible worlds in which it rains tomorrow, and the total number of you-centred possible worlds. The probability that it will rain tomorrow is the proportion of you-centred possible worlds in which it rains tomorrow – and you can make this result as accurate as we want just by increasing  $n$ . This can be applied to issues of civilization longevity too: to determine the probability that your civilization will still exist one thousand years in the future, you need to determine the proportion of you-centred possible worlds in which your civilization still exists one thousand years in the future.

The important point here is that, given some previous observations of reality, the ultimate reference class, fully defines the statistics of your possible situations. You can have access to it in the sense that, in principle, you can program a computer to generate

**The self-indication assumption almost stops the Doomsday argument. Almost. V2.0**

the ultimate reference class with some maximum description length,  $n$ , with  $n$  as large as available computer power and time permit, so in principle you can make the accuracy with which the output of your computer run represents the ultimate reference class as great as you want: we will call this having *full access* to the ultimate reference class. Having full access to the ultimate reference class, given some previous observations of reality, gives you as much information as you can possibly have about your situation, because the entire distribution of possible situations is being laid out before you, ready for counting, with as much precision as you want. If you have full access to the ultimate reference class, any philosophical argument about issues such as Occam's razor<sup>1</sup>, the Doomsday argument or the SIA are irrelevant, because you already have the full description of the statistics of your possible situation and they cannot change this. The ultimate reference class is *superior* to any other reference class or argument that might be made.

This could have clear implications for the SIA. The SIA asks you to think about situations that *could* have existed, but if you have full access to the ultimate reference class this is irrelevant: you already have access to all the situations in which you could exist.

---

<sup>1</sup> Here, we are considering the term "Occam's razor" as it is sometimes used to refer to various methods of theory selection based on some concept of information content or complexity, rather than to refer to an informal heuristic, as is typically the case in philosophy.

### 3 A Practical Consideration: You cannot see all of the ultimate reference class.

What has just been said in 2.3 may seem to be implying that ideas such as Occam's razor, the Doomsday argument and the SIA are irrelevant, as the ultimate reference class is going to determine things anyway, but this is not the case. Although the ultimate reference class *is* superior to any philosophical method or idea, we can never have full access to it. We can only ever have some information about the ultimate reference class. This is where various philosophical methods might become useful: they might tell us what the ultimate reference class might be like. If ideas such as Occam's razor, the Doomsday argument and the SIA are valid, you might use them to give you some information about the distribution of you-centred possible worlds in the ultimate reference class, in addition to anything you already know about it.

As an analogy, the ultimate reference class might be considered to be a line of closed boxes, each with various items in. You have a bag with some items in it, and you know that these match the contents of one of the boxes, but you do not know which one. You may have somehow been able to get some idea of what is in the boxes, but you know very little. You may use some philosophical idea to tell you what kinds of things are likely to be found in a lot of the boxes, and this would tell you what you might expect to find in your bag, but the boxes themselves always take precedence. A philosophical method can never change what is already in the boxes: it is only of any use in so much as it can give you an idea of what is in the boxes. Likewise, a philosophical idea which is supposed to give you an idea of expectation, such as Occam's razor, the Doomsday argument or the SIA cannot change the ultimate reference class. It can only ever be of any use in so much that it can tell you something about the ultimate reference class: about the distribution of you-centred possible worlds in it.

This has important implications for anyone proposing some idea that tells you about your expectations of reality. If an idea such as Occam's razor, the Doomsday argument or the SIA really tells us anything – then it should tell us about the ultimate reference class – about the distribution of you-centred worlds that are actually in it.

A correct formulation of any such argument would be saying something about mathematics as much it would be about the physical world. Although an argument about our expectation would involve some initial information about the physical world, the argument would not be using this to make claims about the physical world directly. Rather, it would be using that information to make a claim about a distribution which could actually be computed from that information, if you had sufficient computing power. In principle, you could program a computer with everything that you know about your situation and instruct it to generate every possible you-centred possible world description consistent with that information, subject to some maximum description length,  $n$ . A correct version of the Doomsday argument, or some argument about theory

selection, would actually be predicting the results of such a computation with a large value of  $n$ , and telling you that, if you did perform such a computation, you should expect a high proportion of the you-centred possible world descriptions generated by the computer to be ones which have some characteristic. In the case of the Doomsday argument, for example, it would be saying that they will tend to be ones in which your civilization does not last very long. Thinking of reference class in terms of predicting the results of computations, based on partial knowledge of the results, like this, may suggest some path to a formal method for dealing with the issue of what the members of a reference class should be.

Anyone proposing an idea that tells you about your expectations – such as the Doomsday argument, the self-indication assumption or Occam's razor, should be able to produce a justification of it that shows how it tells us something about the ultimate reference class, and if he cannot do this then there is no reason for thinking that the idea can tell you anything. We might imagine how such justifications could be made for various ideas. Occam's razor, for example, might be justified in terms of "simpler" theories having higher measure across you-centred possible worlds. The Doomsday argument might also be justified by an argument about the distribution of you-centred possible worlds, and this will be briefly discussed later in this article. This article, however, is about the SIA specifically, and if we are going to demand that ideas such as this are justified like this, it has clear implications for the SIA.

## 4 Relating the Self-Indication Assumption (SIA) to the Ultimate Reference Class

### 4.1 The Self-Indication Assumption (SIA) must tell us about the distribution of you-centred possible worlds.

The SIA tries to get you to think that it is more likely that you inhabit a possible world if there are more observers throughout history in that possible world, and therefore there are “more chances” for you to exist. All of these other people in a possible world, however are not you. The SIA nevertheless suggests that you should think that a lot of observers in a possible world, even though they are not identical to you, “could have been you”, and that this makes such a possible world more likely to be the one you inhabit. This contrasts with the idea of using the ultimate reference class: if you have full access to the distribution of you-centred possible worlds, then this fully describes the statistics of your situation and no considerations of the “could have been” kind are going to change it. The fact that we do not have full access to the ultimate reference class is the only reason that the SIA can have any chance to work. For the SIA to be a plausible objection to the Doomsday argument, its “could have been” is going to have to map onto the distribution of you-centred possible worlds in the ultimate reference class and tell us something about it. Can it do this?

### 4.2 Secondary Reference Classes

The ultimate reference class has been defined as the set of all formally describable you-centred possible worlds and your situation must be a member of this set. However, you may need to consider a larger set of observer-centred possible worlds which includes the ultimate reference class, but which includes other observer-centred possible worlds too. We will call such a set of observer-centred possible worlds a *secondary reference class*. A secondary reference class may need to be considered because you may have a set of observer-centred possible worlds about which you lack complete knowledge, and while you may know that some of these possible worlds are the you-centred ones, corresponding to your situation, you may not know which ones these are, so that you have to consider a larger set of possible worlds.

A secondary reference class, when it does need to be considered, should be defined as specifically as possible while ensuring that it contains all of the ultimate reference class. Ideally, we would only need to consider the you-centred possible worlds of the ultimate reference class, but if we cannot meet this ideal we should approach it as closely as possible, using a secondary reference class which involves adding as few observer-centred possible worlds as possible to the ultimate reference class.

If you live in the 21<sup>st</sup> Century, then an example of a secondary reference class would be the set of all observer-centred possible worlds corresponding to observers living in the 21<sup>st</sup> Century. This set is more general than the ultimate reference class for you, but any possible world centred on you must be one for a 21<sup>st</sup> Century observer, so this reference class must include every you-centred possible world, and must therefore include the ultimate reference class: it is just defined more generally than the ultimate reference class. Another example of secondary reference class, and a more general one, would be the set of all observer-centred possible worlds. This set includes every observer-centred world that can be conceived, and again must include every you-centred possible world, and therefore the ultimate reference class.

### 4.3 An Analogy: Chessboard Worlds

We will use an analogy to give an idea of how we should approach observer-centred possible worlds.

Suppose we use chessboards to represent observer-centred possible worlds. Each square may or may not have a piece on it. Each major piece (king, queen, bishop, knight, rook) corresponds to an observer, while pawns do not correspond to observers or anything of real interest to us: a pawn, if it is on a square, is just there.<sup>2</sup> We will assume that we are dealing with observers living in infinite worlds, so from the point of view of any observer, the chessboard is infinite – with the pattern of black and white squares extending endlessly in all directions.

If we look at such a chessboard we will see some squares with major pieces on them. Each such piece corresponds to an observer, and that observer will be in a particular situation that is defined by the type of piece that he is, the colour of square this piece is on and the arrangement of other pieces (major pieces corresponding to other observers and pawns) around him. If an observer-centred world were to be described for such an observer, therefore, the description would need to include all this information. The chessboard goes on without limit in all directions, however, so from the point of view of an observer, the chessboard extends infinitely in all directions from his own position, and the observer-centred world description would need to describe all this: this could never be done with a description of finite length. This may *seem* to disagree with Max Tegmark's view that a minimal amount of information should be needed to describe an infinite, Level IV multiverse (Tegmark, 1998), but there is no such conflict: Tegmark is viewing things in terms of the description of an entire, impersonal world, whereas here we are considering observer-centred possible worlds. An observer might live in a Tegmark ultimate ensemble, with minimal information needed to describe the whole of reality, but a description of an observer's actual situation needs to locate the observer in this ensemble, and that cannot be done with a finite amount of information. The best

---

<sup>2</sup> Pieces that are not observers (pawns) have just been introduced to add some realism: not everything in reality is an observer.

that can be done is what we are doing here: describing an observer-centred world with everything that can be said about an observer's situation given some maximum number of bits,  $n$ , in the description. (It should be noted that we are not considering time in this analogy, so the idea of anything like an initial state does not concern us.) With the description length is limited to  $n$  bits, the information provided will describe the arrangement of pieces on the chessboard in some finite area around an observer's square: an observer's situation could be described local to that observer. The amount of the chessboard covered by such a local observer-centred world description would depend on  $n$ , the maximum description length in bits.

Requiring descriptions to be shortened like this means that, for any value of  $n$ , there is only a finite number of possible observer-centred world descriptions. For example, suppose  $n$  is made large enough to allow an observer-centred world description to cover a grid of  $9 \times 9$  squares, with the observer being at the centre. The observer must be on a black square or a white square, so this gives two possibilities. The observer must be a major piece (king, queen, bishop, knight, rook), so this gives 5 possibilities. The grid of  $9 \times 9$  squares contains 81 squares, so without the observer's own square (as it was just dealt with) 80 squares need to be described. On each of these squares there can be another observer (a king, queen, bishop, knight, rook), a pawn or no piece at all, so this gives 7 possibilities for each square. There are therefore  $7^{80}$  possibilities for all 80 squares. Putting all this together, the number of different observer-centred possible worlds that can be described is  $2 \times 5 \times 7^{80} = 4.1 \times 10^{68}$ . This is, of course, an extremely large number of observer-centred possible worlds; however an infinity of such "chessboard-type" worlds is actually available: limiting the description length means that an infinity of observer-centred possible worlds has to be represented by just  $4.1 \times 10^{68}$  worlds. This does not mean that the set of  $4.1 \times 10^{68}$  worlds is useless: by examining the statistics of this set we can get an idea of the statistical properties of the set as the maximum description length,  $n$ , is increased, and by increasing  $n$  we can find out what happens as  $n$  tends to infinity. An important point, here, is that the infinite nature of each observer-centred possible world is not causing any contradictions or incoherence: it can be dealt with.

Suppose now that we have two specific types of observer-centred "chessboard" world.

- A **Type 1 Chessboard World** is an infinite chessboard, as described above, but with a restriction. Any piece (king, queen, bishop, knight, rook, pawn) or none can be placed on a black square, but no pieces of any kind are allowed on white squares: they must be unoccupied. Apart from this restriction, the pieces appear to be arranged randomly, with no obvious pattern, so any local arrangement of pieces that does not break these rules will correspond to an observer-centred possible world.
- A **Type 2 Chessboard World** is an infinite chessboard, as described above, but also with restrictions. Any piece (king, queen, bishop, knight, rook, pawn) or none can be placed on a white square, but major pieces (king, queen, bishop,

knight, rook, pawn) are not allowed on black squares: a black square must be occupied by a pawn or must be unoccupied.

The number of observer-centred possible worlds of each kind can now be computed. We will still assume that the maximum length of any observer-centred world description is being limited so that it can cover just a grid of 9x9 squares with the observer at the centre.

- For a **Type 1 Chessboard World**, there are 5 possible pieces for the observer (king, queen, bishop, knight, rook) on the central black square, and 7 possibilities (king, queen, bishop, knight, rook, pawn or none) for each of the 40 surrounding black squares in the 9x9 grid. The number of observer-centred possible worlds that can be described is therefore  $5 \times 7^{40} = 3.2 \times 10^{34}$ .
- For a **Type 2 Chessboard World**, there are 5 possible pieces for the observer (king, queen, bishop, knight, rook) on the central white square, and 7 possibilities (king, queen, bishop, knight, rook, pawn or none) for each of the 40 surrounding white squares in the 9x9 grid. In addition, there are 2 possibilities (pawn or none) for each of the 40 surrounding black squares in the 9x9 grid. The number of observer-centred possible worlds that can be described is therefore  $5 \times 7^{40} \times 2^{40} = 3.5 \times 10^{46}$ .

The Type 1 chessboard worlds make up  $3.2 \times 10^{34}$  observer-centred possible worlds in the secondary reference class and for the Type 2 chessboard worlds there are  $3.5 \times 10^{46}$  of them. When the description length of an observer-centred possible world is limited in this way, Type 2 chessboard worlds are therefore contributing more observer-centred possible worlds to the secondary reference class than those of Type 1.

This situation will continue to apply as the maximum description length of observer-centred possible worlds,  $n$ , is increased, and there will be many more observer-centred chessboard possible worlds of Type 2 than Type 1 as  $n$  tends to infinity. Type 2 chessboard worlds will make up more of the secondary reference class of observer-centred worlds.

A Type 2 chessboard world is made according to rules which mean that there are more ways of describing observer-centred worlds of that kind, subject to some maximum description length,  $n$ . We should now consider how this will appear to an observer in a Type 2 chessboard world. Because the rules which define the appearance of the observers are more conducive to the existence of observers, a Type 2 observer-centred chessboard world, with some maximum description length,  $n$ , will contain more observers than a Type 1 chessboard world. Type 2 chessboard observer-centred possible worlds make up a greater proportion of the secondary reference class than Type 1 ones. All else being equal, the number of observers is obviously relevant in determining the degree of representation of a kind of world within the observer-centred worlds of the secondary reference class.

If we are talking about numbers of observers, we need to consider the issue raised by the prospect of chessboard worlds containing *infinities* of observers – and we earlier said that we were imagining infinite chessboards in this scenario. This problem never arises with this kind of situation, because if a chessboard world is infinite we never see all of it in our consideration: we only deal with observer-centred possible worlds with some finite description length,  $n$ , as  $n$  *tends to* infinity: all of the observer-centred possible worlds that could ever be considered have finite numbers of observers. If we are not “seeing” all of the observers with any finite description length, however, this means that what we are thinking of as the number of observers in a possible world is really the density of observers. The number of observers is a reasonable simplification of what we should be thinking about, but density of observers is better: as will be shown later, this in turn is a simplification of a more general, information-theoretical consideration that we might make. The distinction is not too important: an observer, looking out from the centre of his chessboard observer-centred world, will see some number of local observers, and whether or not this is considered as an absolute number or a density should not matter too much.

There is another way of looking at this. If an observer-centred chessboard world contains a high density of observers, it will tend to overlap with many other observer-centred worlds. Suppose you are an observer in your observer-centred chessboard world, and you see another observer in your world. For a chessboard world, this means that you are a major piece (king, queen, bishop, knight, rook) on the central square of your observer-centred world, and you see another major piece (king, queen, bishop, knight, rook) on one of the surrounding squares. This other observer is actually an indication of another observer-centred world that can be made, centred on this other observer. Therefore, each observer you see in your own observer-centred chessboard world is an indication of another observer-centred world in the reference class that overlaps with it. In this article, we are trying to avoid discussion of *impersonal* possible worlds, but if you really need to think in terms of such things, you could regard this as meaning that when an observer looks out in his own observer-centred chessboard world and sees another observer, it indicates that both observers share the same impersonal world, and each of their observer-centred possible worlds with a limited description length is a different piece of it. This should mean that a type of observer-centred chessboard world in which observers can see lots of observers should make up a big proportion of the secondary reference class, because each such observer-centred world implies the existence of many others.

#### **4.4 What determines the degree of representation of an impersonal possible world in the secondary reference class?**

The analogy of chessboard worlds that was just discussed gives an idea of how we can approach observer-centred worlds in general.

Let us consider the secondary reference class of all observer-centred worlds that could possibly be described, subject to some maximum description length of  $n$  bits.

Our consideration of chessboard observer-centred possible worlds in 4.3 has shown that the extent to which a type of observer-centred possible world features in the secondary reference class of observer-centred possible worlds is related to the density of observers in such a possible world. For some maximum description length,  $n$ , if a given type of observer-centred possible world, centred on an observer  $O$ , has a larger number of observers in it as well as  $O$ , this implies that the way such possible worlds are constructed is conducive to the appearance of observers in them, and that many such observer-centred possible worlds can be constructed with the maximum description length,  $n$ , which in turn implies that observer-centred possible worlds of such a type will make up a lot of the secondary reference class of observer-centred possible worlds.

There is another way of looking at this. If an observer,  $O$ , looks out into his own observer-centred world and sees another observer, this is an indication that another observer-centred possible world of the same type as  $O$ 's can be described which is centred on this other observer, and if  $O$  looks out into his own observer-centred world and sees many other observers, this is an indication that many other observer-centred possible worlds of the same type as  $O$ 's can be made, each centred on one of these other observers, and it is reasonable to think that the amount of information needed to describe any one of these observer-centred possible worlds in enough detail to resolve it apart from others – to give a unique observer-centred possible world – has a good chance of being not too different from the information needed to describe  $O$ 's observer-centred world, and so will therefore comply with whatever maximum description length,  $n$ , is in effect, therefore featuring in the secondary reference class of observer-centred worlds.

This can also be viewed in terms of observer-centred worlds overlapping when there are many different ways of describing them – and therefore a lot of them in the secondary reference class. If  $O$ 's observer-centred possible world is of a type that is very common in the secondary reference class then some of these other observer-centred possible worlds are likely to overlap with  $O$ 's, and the owners of these other observer-centred possible worlds will appear in  $O$ 's observer-centred possible world as other observers.

Observer-centred worlds in which lots of observers are found should therefore account for a high proportion of the secondary reference class of observer-centred possible worlds.

In this article, we are trying to avoid use of impersonal observer-centred possible worlds, but readers who want to think in such terms can think of it like this: if  $O$  looks out into his own observer-centred possible world and sees another observer, then as well as indicating that two observer-centred possible worlds overlap, it indicates that both observers belong to the same impersonal personal world, which is contributing

these observer-centred possible worlds to the reference class – implying that densely populated impersonal worlds will tend to contribute more observer-centred possible worlds to the secondary reference class.

The issue of *infinite* observer-centred possible worlds needs to be considered, and it raises the issue of whether we should think in terms of *number* of observers or *density* of observers. Thinking in terms of number of observers may seem to cause problems with infinite possible worlds, but possible worlds are only considered with finite description lengths as the description length tends to infinity, so the problem of infinite observers is never actually encountered. Nevertheless, as the number of observers in an observer-centred possible world with some finite description length is only ever going to be the local number, it makes more sense to think in terms of density of observers. The distinction is not too important, though: we can just think in terms of number of observers for some maximum description length. *All else being equal*, the more densely populated a type of observer-centred possible world is, the more common it should be in the secondary reference class of observer-centred possible worlds.

It should be recognized, however, that even this is a simplification. What really determines the proportion of observer-centred worlds of some type in the secondary reference class for some maximum description length,  $n$ , is the numbers of observer-centred world descriptions that can be made for that type of possible world.

The density of observers in a typical observer-centred possible world of some type will be important, but the amount of information needed to describe a single observer's situation will also matter. For example, if there is a maximum description length,  $n$ , and a lot of information is needed to describe the basic physical laws of an observer-centred possible world of some type, then less information will be available to describe the situation that is specific to a particular observer, and there will be fewer ways to resolve observer-centred worlds of that type apart, and therefore fewer observer-centred worlds of that type that can exist in the secondary reference class of observer-centred worlds with some maximum description length,  $n$ . Some readers will recognize this as *Occam's razor*. What is really important about a high density of observers in some observer-centred possible world, is that it implies that unique observer-centred worlds of that type exist that can be resolved apart with relatively little information, allowing a lot of observer-centred world descriptions to be made for some maximum description length. We might therefore consider the idea of spatial density of observers as being relevant because of its connection with a more general kind of "information-theoretical" density of observers – how many observers can be found per bit of description in some observer-centred possible world, when the description length tends to infinity, and ignoring any issues relating to any spatial size of the world. This more rigorous way of viewing things is more flexible, because it allows us to consider observer-centred possible worlds in which there may not be any space: we might imagine some strange world in which observers are not separated spatially, but in some more abstract way.

The main point of what has been discussed here is that it seems to indicate that a greater tendency for observers to be found in an observer-centred possible world of a given type means that more observer-centred possible worlds of that type will exist in the secondary class. It may seem, so far, that this discussion is actually going to show the SIA to be true!

## 4.5 Relating the Secondary Reference Class to the Ultimate Reference Class

It has been shown how the degree to which a given type of observer-centred possible world is represented in the secondary reference class of observer-centred possible worlds is dependent on the density of observers that such possible worlds tend to have, subject to the qualifications just given in 4.4 and *all else being equal*. The secondary reference class now needs to be related to the ultimate reference class: the set of you-centred possible worlds. When that is done, a full view of the statistics of your situation will be available.

The ultimate reference class contains every observer-centred possible world centred on *you* as the observer, subject to a maximum description length of  $n$  bits: this means that every observer-centred possible world in the ultimate reference class has to be consistent in every detail with everything that you know about your situation. The secondary reference class that we have been discussing contains every observer-centred possible world centred on *anyone* as the observer, subject to a maximum description length of  $n$  bits. A you-centred possible is just a special case of a possible world centred on *anyone*, and this means that every you-centred possible world in the ultimate reference class also appears in the secondary reference class: the secondary reference class contains the ultimate reference class. The ultimate reference class consists of some of the possible worlds from the secondary reference class.

Suppose that you do not know anything about whether any particular observer-centred possible world in the secondary reference class of observer-centred possible worlds corresponds to you or some other observer. You know that every you-centred possible world in the ultimate reference class corresponds to a possible world in the secondary reference class and, for any maximum description length,  $n$ , you should view any given possible world in the secondary reference class as being as likely to be a you-centred possible world as any other.

This means that, *all else being equal*, for any maximum description length of observer-centred possible worlds,  $n$ , the number of observer-centred possible worlds of some type that correspond to observer-centred worlds in the ultimate reference class should be expected to be proportional to the number of observer-centred worlds of that type in the secondary reference class. The reasoning behind this is quite simple. If we regard the you-centred possible worlds as being randomly selected from the secondary reference class of observer-centred possible worlds, then each observer-centred

possible world of a given type should be regarded as “having a chance” of being a you-centred possible world, and more observer-centred worlds of a given type means more chances. To put this another way, if X% of observer-centred possible worlds are you-centred possible worlds, then the greater the number of observer-centred possible worlds of some type in the secondary reference class of observer-centred possible worlds, then, all else being equal, the greater the number of such possible worlds we should expect to be in the X% that are you-centred possible worlds.<sup>3</sup>

The number of observer-centred possible worlds of some type in the secondary reference class gives us, then, an idea of the extent to which it is represented in the ultimate reference class of you-centred possible worlds, but this is dependent on the density of observers that such possible worlds tend to have.

*All else being equal*, we should expect the degree to which a type of observer-centred possible world, is represented in the ultimate reference class of you-centred possible worlds to be proportional to the density of observers in that possible world or type of possible world. It should be noted that, as stated earlier, this is a simplification, and a more rigorous consideration to make is an information-theoretical one.

---

<sup>3</sup> Some readers may object to this by pointing out that all the observer-centred worlds of a given type may have physics that is completely different to yours, or may be possible worlds in which all inhabitants are compelled by some strange rule to have Mickey Mouse ears – clearly ruling out any such possible worlds as candidates for yours (we assume), but this misses the point: this is a statistical argument only, and the “all else being equal” qualification is important.

## 5 Where does this leave the self-indication assumption (SIA)?

### 5.1 Does the number of observers in an observer-centred possible world matter?

As has been discussed, it is more accurate to think in terms of density, rather than number, of observers, but this should not concern us too much: from a local perspective, density and numbers will be fairly indistinguishable.

The discussion given so far has actually shown that, *in a general sense*, the SIA is valid! The statistics of your possible situation are fully defined by the ultimate reference class of you-centred possible worlds, and observer-centred possible worlds of various types should be expected to be represented in the ultimate reference class according to density of observers.

Some qualifications should be given, here, however. What has been said previously about number of observers, density of observers and information-theoretical considerations apply: considering things in terms of density of observers is a simplification, and there is an “all else being equal” aspect to this. Nevertheless, there seems to be good reason to think that you are more likely to be in a densely populated observer-centred possible world than a less-densely populated one. It may seem that the discussion so far, by validating the SIA, amounts to a refutation of the Doomsday argument. However, we should now consider a further, important qualification of what has been said: the justification that has been given for the SIA is based on the assumption that we do not know which observer-centred possible worlds in the secondary reference class are you-centred possible worlds, but the proposal that the SIA refutes the Doomsday argument is suggesting that you “count” observers in the future of your civilization and the implications of this should be considered.

### 5.2 Does the number of observers in the *future* in some impersonal possible world matter?

#### 5.2.1 The Requirement for Observer-Centred Worlds

It has been shown that the SIA is valid in a general sense and that, *all else being equal*, you should consider yourself more likely to be in a densely populated you-centred possible world than a sparsely populated one. (As stated previously, this is a simplification, and an information-theoretical consideration should really be made.) The basic idea, here, is that, unless you know otherwise, any observer-centred world in the secondary reference class should be regarded as a possible candidate for one of the you-centred possible worlds in the ultimate reference class – and one of these

corresponds to your situation. An observer in a possible world is important because that observer might be you.

When we are considering types of observer-centred possible world for which we do not know much about the specific situations of observers, this makes sense. For example, one cosmological theory may lead to a different density of observers than another theory, based on things such as the expected density of planets suitable for life. A general version of the SIA would favour the theory with the greater observer density, though as has been stated this only applies in an “all else being equal” sense, and a proper, information-theoretical consideration would be more accurate. This would be because any of the observers in an observer-centred world corresponding to such a cosmological model *might actually be you*.

The requirement for an observer-centred possible world to be a *you*-centred possible world is a strict one. It must be consistent with every detail of what you know about your situation: otherwise, on finding out about an inconsistency, you could rule it out. For example, if you remember wearing a necktie with a particular pattern last week, any observer-centred world, to be a candidate for a *you*-centred possible world, must be owned by an observer who also remembers wearing a necktie with that pattern last week. Clearly, the probability of any single observer-centred world matching your experiences enough to be a *you*-centred possible world is very remote. However, for a large enough maximum description length,  $n$ , there will be a large number of observer-centred worlds – and  $n$  is considered as tending to infinity.

This is only statistics: the fact that a large enough collection of observer-centred possible worlds will be expected to contain a number of *you*-centred possible worlds does not mean that any subset of it has to do so. This only applies if we have no reason to think otherwise. If we know that all the observer-centred possible worlds in some collection of observer-possible worlds are inconsistent with your experience in any way then, no matter how large this collection is, you will know that none of them can be *you*-centred possible worlds – even though a naïve statistical consideration might say otherwise.

This is what determines how the SIA interacts with the Doomsday argument. The SIA suggests that, because there will more observers in the future of your civilization if it is going to last longer, the types of possible world in which your civilization lasts longer contain more observers and, therefore, that this weights things statistically in favour of thinking that you live in a possible world corresponding to a civilization that is going to last a long time. This is supposed to cancel out the Doomsday argument.

In the language used so far in this article, the SIA is suggesting that the probability of any type of observer-centred possible world being yours should be adjusted according to the typical local number of observers, or density of observers, in such a possible world, but the argument in this article has shown that, while there is some justification for such a

view, it can only be justified on the basis that the greater the typical density of observers in some type of observer-centred possible world, the more of the secondary reference class it takes up, each observer-centred possible world in the secondary reference class has a chance of being a you-centred possible world, the greater the number of observer-centred possible worlds, the greater the number of you-centred possible worlds that are expected – and it is the set of you-centred possible worlds that determines the statistics of your situation. An observer contributing to the number or density of observers in some type of observer-centred possible world that is owned by someone else affects the probability that you are in that type of world, but only because it indicates that that observer is in turn expected to correspond to his observer-centred possible world of the same type, which might be a you-centred possible world – because the observer might be you. What “might have been” is irrelevant. Your situation is fully described by the ultimate reference class of you-centred worlds, and an observer can only affect your statistics if that observer might correspond to a member of this reference class – if the observer could be you. If you actually know that an observer cannot be you – for example, wrong food preferences, wrong career history, *wrong period in history* – then the observer-centred world corresponding to that observer should be removed from the secondary reference class.

It should be apparent, now, that the idea that the SIA can work as a rebuttal of the Doomsday argument is in trouble.

### **5.2.2 The Problem with Future Observers - Why the Self-indication Assumption (SIA) Fails to Refute the Doomsday Argument**

Just taking the secondary reference class as it is may be fine when cosmological theories are being considered – when you do not know much about the observers in it – but the problem with trying to apply it to observers in the future of your civilization is that *you actually know that these observers are not you*, and the fact that you know this means that these observers should be removed from the secondary reference class and should not be considered as having any affect on the ultimate reference class and the statistics of your possible situation. The only reason for “counting” an observer that appears in someone else’s observer-possible world is that it suggests that this observer will have his/her own observer-centred possible world and that the observer might actually be *you* – which means there must at least be the possibility that the observer’s experiences are consistent with yours in every detail – but if the observer is living at a different time in history, this cannot be the case.

If you live in the 21<sup>st</sup> Century in London, wear a black business suit and drive a BMW, it should be obvious that someone who lives in the 27<sup>th</sup> Century in New London on Mars, wears a spacesuit, drives a flying car and sees a different date when he looks at the calendar is not a candidate for you. The observer-centred world corresponding to such a person can never be a you-centred possible world in the ultimate reference class, and the observer-centred worlds corresponding to such observers should be removed from

the secondary reference class. Such observers have no relevance to the general version of the SIA which has been justified in the previous discussion. Further, the previous discussion has shown how observers in types of observer-centred possible worlds relate to the ultimate reference class of you-centred possible worlds, and therefore to the statistics of your possible situation. Not only do observers living in your future play no part in this, but also the relationship between observers in observer-centred possible worlds and you-centred possible worlds has been discussed in sufficient detail to exclude any plausible way in which an observer living in your future could affect your statistics: an observer in your future is specifically invalid as a member of the reference class that does determine your statistics. At the same time, intuitive ideas about more observers making a possible world more likely have been dealt with, but this will not help the SIA to oppose the Doomsday argument at all, because the general version of the SIA that results specifically excludes the future observers needed to refute the Doomsday argument. The version of the SIA which was previously found to be valid only applies in an all else being equal way, when you have no reason to think that the observers being counted could not be you, but here all else is specifically *not* equal.

On the basis of this, the idea that the SIA can refute the Doomsday argument should be regarded as implausible. The SIA has been shown to be valid in a general sense, but in the specific sense of answering the Doomsday argument, the SIA should be regarded as dead.

## 6 Further Considerations

### 6.1 A Short Justification of the Doomsday Argument

The purpose of this article is not to prove the Doomsday argument correct, but merely to show that using the SIA against it is invalid. Nevertheless, a brief justification of the Doomsday argument will now be given, based on the ideas of the secondary reference class and the ultimate reference class that have been used in the discussion of the SIA.

We will consider a secondary reference class of observer-centred worlds, but we are not yet relating this to the ultimate reference class, so we will therefore allow it to contain all possible observer-centred worlds – even ones that could not possibly be you-centred worlds.

You have a particular position in history as an observer, and according to the Copernican principle, you should regard your situation as an observer as typical of that of other observers. There is statistical justification for thinking that your situation is typical of that of other observers in the secondary reference class of observer-centred worlds. This does not mean that all the other observers have to be in situations exactly like yours, or even in human civilizations, but it does mean that your own situation should not be exceptional. The common experience of an observer should be that of being in a civilization at about the same age and level of development as yours. The experience of being in, for example, a galactic civilization which developed interstellar travel twenty million years ago should not be expected to be typical; otherwise the question would arise of why *you* do not find yourself in such a situation if it is so common. This suggests that your observer-centred world should be typical of other observer-centred possible worlds in the secondary reference class: they should tend to be for observers in civilizations not too far ahead of yours in terms of age or level of development.

As discussed previously, if an observer, *O*, looks out into his observer-centred possible world and sees another observer, this indicates that an observer-centred possible world will exist for this other observer. Now, suppose that the observer-centred world owned by *O* contains a large number of observers in the distant future of *O*'s civilization. The existence of each of these observers in *O*'s possible world indicates that each of these observers corresponds to his own observer-centred possible world, so the result should be that a large number of observers living in the distant future of *O*'s civilization, from *O*'s point of view, end up in the secondary reference class of observer-centred possible worlds. Suppose now that this situation is typical of an observer-centred possible world in the secondary reference class: that a typical observer, *O*, inhabits an observer-centred possible world that contains many other observers living in the future of his civilization. The result should be that these other observers, living in very old civilizations, have their own observer-centred worlds and the secondary reference class of observer-centred possible worlds should contain a lot of such observers. However, it has been previously shown that the secondary reference class should not be expected to have a high

proportion of observers who are living in very old civilizations. The implication is that this is probably not the case: that the typical owner of an observer-centred possible world in the secondary reference class does not inhabit an observer-centred world containing many observers living in the future of his civilization, and the implication of this is that the typical observer lives in a civilization that is not going to last into the distant future.

This establishes the main conclusion of the Doomsday argument. The above discussion was not complicated with mention of limits on description length, and maximum description length tending to infinity: things have been discussed in those terms throughout enough of the article that the general idea should be clear, nor have the detailed, probabilistic aspects of the Doomsday argument been discussed.

This kind of justification for the Doomsday argument is not about your civilization specifically, but about the distribution of possible civilizations that could exist. It suggests that your civilization will end soon, but it also suggests that the typical experience of an observer is one of living in a civilization not significantly older or more advanced than yours, and that if other civilizations exist, these also tend to be short-lived. This version of the Doomsday argument may be somewhat nastier than the version assumed to apply by many people: it is as general in scope as the argument could be, and seems to follow the “*Nobody gets out alive*” principle of Japanese horror films.

As discussed previously, this kind of justification of the Doomsday argument is as much about mathematics as it is about the physical world. It is using some information about your situation to make an inference about the expected output of a hypothetical computer run to generate the set of possible descriptions of your situation. It is telling you that, if you did perform such a computation, you should expect a high proportion of the you-centred possible world descriptions generated by the computer to be ones in which your civilization does not last very long. This idea can be used to justify the application of the Copernican principle – the idea that your situation should be typical of that of other observers – that was just made. You might ask why your situation should be typical of others in the secondary reference class of observer-centred worlds, but the justification is merely that it is the only member of the reference class of observer-centred worlds about which you have any direct knowledge. If a computer is programmed to produce a list of observer-centred world descriptions, then your knowledge of your own situation corresponds to one item in this list. You have partial knowledge of the computer program’s output that can tell you something about the rest of the computer program’s output. What you know about the program that generates the list will be of relevance here. The kind of characteristics found in your own situation that you expect to find in other situations should also be relevant. This should give more of an idea of what we meant when we said, earlier, that the computational view of reference classes being applied here might suggest a path towards a formal method for dealing with reference class membership.

The Doomsday argument can tell you nothing about the expected outcome of any single random event. For example, suppose that you were about to toss a coin and if the result were “heads” a bomb would be armed, set to destroy your civilization within some time period, while there would be no such consequences if the result were “tails”. A naïve understanding of the Doomsday argument might suggest that you should think it more likely that the result would be “tails”, to comply with some statistical expectation of the end of your civilization, but the understanding of the Doomsday argument given here could justify no such belief as the version of the Doomsday argument justified in this way can only deal in distributions. (There is one qualification to this: if that if the coin is biased in some way, it might be the case that such a bias is common in possible worlds in general.) This should call into question the idea that William Eckhardt’s “shooting room” analogy of the Doomsday argument, used in his criticism of it, is actually a reasonable analogy of the Doomsday argument (Eckhardt, 1997).

## 6.2 Infinite Worlds and Level IV Multiverses

The issue of whether to consider things in terms of number of observers or density of observers was discussed previously. It was mentioned that the correct kind of consideration is an information-theoretical one, in which the number of observer-centred possible worlds of some type in the secondary reference class for some maximum description length,  $n$ , depends on the number of different observer situations which can actually be constructed subject to that limit on description length. Concepts such as “number of observers” or “density of observers” are valid, however, as simplifications when an “all else being equal” kind of consideration is being made.

The idea of “density of observers” is a better one than “number of observers” because of the issue of infinite possible worlds, although the hypothetical calculation being considered never needs to encounter an infinity of observers: it deals only with local numbers of observers anyway. An issue raised by all this, however, is the one of whether your ultimate reference class should be dominated by finite observer-centred possible worlds with finite numbers of observers or infinite observer-centred possible worlds with infinities of observers: in other words, should you expect to live in a finite world or an infinite world? It can be shown that you should expect to be living in an infinite world with an infinity of observers.

Let us consider the secondary reference class of observer-centred possible worlds with some very large maximum description length,  $n$ . This reference class will contain some observer-centred possible worlds with description lengths that are much shorter than  $n$  bits. These will correspond to possible worlds that are finite, so that as the maximum description length is increased, beyond some point, the description length does not increase: there is no more information to give about the possible world. Because the description lengths of such possible worlds are much less than  $n$ , however, only a very limited amount of such possible worlds can be described and they must occupy only a small proportion of the secondary reference class. The secondary reference class must

therefore be made up of possible worlds with much longer descriptions. This can be said no matter how large the maximum description length,  $n$ , is. The implication is that for any finite description length you can state, by allowing a larger maximum description length for the secondary reference class, you can always show yourself to be almost certainly living in an observer-centred possible world with a longer description – and as the maximum description length tends to infinity, this near-certainty tends towards certainty.

Suppose we now consider two possible worlds with very long description lengths, one with a very small number of observers and one with a large number of observers. The existence of the other observers in the second world implies that worlds of that type will tend to be more common in the secondary reference class. As the maximum description length is increased, it will always be possible to construct longer observer-centred world descriptions with more observers, so the number of observers that should exist in a typical observer-centred possible world in the secondary reference class should tend to infinity, suggesting that the number of observers in your world should tend to infinity.

Further, it can be argued that as the maximum description length,  $n$ , for observer-centred possible worlds tends to infinity, there will be opportunities for increasingly complex observer-centred world descriptions to exist in the secondary reference class. A type of observer-centred possible world with many different kinds of situations for observers will correspond to many observer-centred possible worlds. This suggests that the secondary reference class will be dominated by types of observer-centred worlds which provide increasingly large numbers of different ways for observers to exist as the maximum description length of an observer-centred possible world,  $n$ , increases: in other words, that the secondary reference class should be dominated by types of possible worlds that provide more scope for making descriptions that distinguish between observers to allow unique observer-centred possible worlds.<sup>4</sup> This may suggest that the kind of world in which you expect to be should become increasingly like the very general kind of multiverse which has been justified in other ways (Lewis, 1986; Tegmark, 1998, 2003, 2007; Mitra, 2010; Almond, 2010). Max Tegmark describes such a general multiverse as a Level IV multiverse. A full argument would be needed to demonstrate this, however, and various issues regarding description length would need to be considered.

---

<sup>4</sup> An objection to this could be that extending a type of possible world in ontologically diverse ways merely creates more world that has to be described locally, and so is no better than just adding more space, but this can be dealt with: building the secondary reference class is not about the ability to fully describe possible worlds, but is about the ability to make descriptions that resolve observers apart.

## 7 Conclusion

The self-indication assumption (SIA) is often regarded as the leading candidate for a refutation of the Doomsday argument, but an analysis of the intersection of the SIA and the Doomsday argument has shown that this idea is flawed.

The basis of any consideration of the statistics of your situation should be the ultimate reference class: the set of all you-centred possible worlds. Each you-centred possible world corresponds to an observer-centred possible world in a secondary reference class of observer-centred possible worlds. It has been shown that the density of observers in an observer-centred possible world will be of relevance to the probability that you are in it. A type of observer-centred possible world which tends to have a high density of observers will tend to correspond to a high proportion of observer-centred worlds in the secondary reference class, because it implies that the way in which such a world is constructed tends to give rise to observers, meaning that there should be more ways of describing possible worlds centred on such observers with some maximum description length,  $n$ . Another way of viewing this is in terms of each other observer seen by the owner of an observer-centred possible world, when he looks out into his own world, as implying the existence of another observer-centred possible world of the same type, centred on this other observer, so the existence of many such other observers implies the existence of many other observer-centred possible worlds of the same type. A further way of looking at this is in terms of the existence of other observers in someone's observer-centred possible world indicating that observer-centred possible worlds of that type are overlapping, indicating that many of them exist in the secondary reference class. Finally, though we are not recommending this as the best way of considering this, if an observer-centred possible world is part of an impersonal possible world with many observers in it, this implies that many other observer-centred possible worlds can be made that are part of that impersonal possible world, and these should occupy a lot of the secondary reference class of observer-centred possible worlds.

As each you-centred possible world corresponds to one of the observer-centred possible worlds in the secondary reference class, all this suggests that, *all else being equal*, the existence of a large number of observers in a possible world indicates that many of your you-centred possible worlds will be of this type, making it likely that you inhabit such a possible world. Consideration of infinite worlds suggests that we should probably think in terms of density, rather than number, of observers – although, as the secondary reference class of observer-centred possible worlds is only ever considered with a finite description length, even though this tends to infinity, we can only ever encounter finite numbers of observers. The most accurate type of consideration, however, is an information-theoretical one, in which the number of observer-centred worlds with some maximum description length,  $n$ , that correspond to some type of observer-centred possible world is considered, and the proportion of such worlds in the secondary reference class is considered as the maximum description length,  $n$ , tends to

infinity. The density of observers will determine the proportion of the secondary reference class corresponding to an impersonal world *all else being equal* – and issues such as the amount of information needed to describe the situations of observers will be relevant. For example, if a lot of information is needed to describe the laws of physics of some type of possible world, less information will be available to resolve observer-centred possible worlds apart for some maximum description length,  $n$ , so this will tend to limit the number of unique observer-centred possible worlds that can be made – and therefore the probability that you are in such a world. This should be recognizable as Occam’s razor.

This actually means that a general version of the SIA is valid. The “number” of observers in a possible world *does* matter. It also resolves one objection to the SIA: that it becomes incoherent when infinite worlds with infinities of observers are involved. As the number or density of observers is merely a simplifying idea, and the correct methodology – applying an information-theoretical consideration – only ever needs to deal with finite numbers of observers and finite possible world description lengths. Even though these numbers are allowed to tend to infinity, the situation remains mathematically coherent, even for infinite possible worlds. This general version of the SIA and Occam’s razor can actually be considered to be the same thing, looked at from different ends. A proper understanding of Occam’s razor tells us that, all else being equal, types of observer-centred possible world in which less information is needed to describe the basic physics should take up more of the secondary reference class of observer-centred possible worlds, while a proper understanding of the SIA tells us that, all else being equal, observer-centred possible worlds of a type that is common in the secondary reference class should be ones with a high density of observers, due to such possible worlds overlapping.

All this may actually seem to be *supporting* the SIA as a rebuttal of the Doomsday argument, but there is a serious problem. What has been shown to apply is a very general version of the SIA, and the existence of other observers in a type of observer-centred possible world only makes it more likely that you are in that type of possible world because each of those other observers indicates the existence of a further observer-centred world that could be yours: *each of those observers might actually be you*. This is relevant if you have only minimal knowledge about these observers, so that you know nothing that could tell you that an observer is not you; for example, a cosmological theory may give you some idea of the density of planets which could produce life, giving you some idea of the density of observers in a very general way. The SIA, however, when used as a rebuttal of the Doomsday argument is suggesting that you should count observers in the future of your civilization in some observer-centred possible world, and these are observers who specifically *are not you*. If you know that an observer could not possibly be you, there are no grounds for “counting” that observer in the general version of the SIA that has been shown to apply.

The general version of the SIA that has been shown to apply, then, does not refute the Doomsday argument. Worse still, obtaining this general version of the SIA has involved enough analysis of the relationship between observers in observer-centred possible worlds and the reference class of you-centred possible worlds that we should know enough about the role played by observers in observer-centred possible worlds with regard to the statistics of your situation to see that it is implausible for observers in the future of your civilization to fit into all this and have some effect, somehow. Further, while we may have some intuitive idea that more observers in an observer-centred possible world corresponds to a greater probability of being in that kind of world, this intuition is easily reconciled with the general version of the SIA, so the more specific version of the SIA needed to refute the Doomsday argument can be dispensed with without having to question this basic intuition.

The title of this article has implied that the SIA *almost* refutes the Doomsday argument, and in a sense, this is true: a general version of the SIA *does* work, that does tell us to count observers, but any application of it against the Doomsday argument fails because it involves the wrong observers.

The purpose of this article has not been to justify the Doomsday argument, but rather to rule out the SIA as a serious candidate for a rebuttal of it. Nevertheless, the Doomsday argument can be expressed in terms of the kind of reasoning used here. The Copernican principle would suggest that your situation should be typical of the situations of other observers, suggesting that the set of observer-centred worlds for all observers tends to contain observers in situations like yours, rather than observers in situations resembling the distant future of your civilization. However, if the typical observer-centred possible world was owned by an observer with many observers in the distant future of his civilization, this would indicate the existence of a large number of possible worlds owned by these observers, and therefore the existence of *many* observers in situations resembling the distant future of your civilization – something which has just been shown to be unlikely. This suggests that the typical observer-centred possible world is owned by an observer who does not have many observers in the distant future of his civilization, suggesting that the typical observer-centred possible world is owned by an observer in a civilization that does not have a distant future. This seems to suggest that your civilization is going to end soon.<sup>5</sup>

The ideas in this article can also be used to argue for an infinite reality with an infinity of observers, and possibly a reality of the kind termed a “Level IV multiverse” in Tegmark’s classification. For any finite world, there will be a finite number of observers with a finite maximum description length for the corresponding observer-centred worlds, and the maximum description length could always be made much larger, to include many

---

<sup>5</sup> There are some loopholes that we may discuss in a later article; however, none of them involve a simple, straightforward continuation of our civilization.

**The self-indication assumption almost stops the Doomsday argument. Almost. V2.0**

more observer-centred worlds, and types of worlds with more observers should always dominate the reference class. Further, as the description length tends to infinity, types of possible worlds which provide more ways for observer-centred possible worlds to be resolved apart would tend to occupy more of the reference class of observer-centred worlds, suggesting that it is likely that the observer-centred world corresponding to your situation is infinite, with the most general ontology possible – providing an infinity of different ways for observers to be produced.

## 8 Acknowledgements

Yvonne Deborah Finch, Katja Grace and Max Tegmark have been helpful in discussions relating to the subject matter of this article.

## 9 References

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/Substrate4.pdf> or <http://www.paulalmond.com/Substrate4.doc> [Accessed 1 January 2011].

Almond, P., 2011. *The self-indication assumption almost stops the Doomsday argument. Almost*. [Online] paul-almond.com. Available at: <http://www.paul-almond.com/SIA.pdf> or <http://www.paulalmond.com/SIA.doc> [Accessed 6 January 2011].  
(This is an earlier version of this article which relies on the idea of impersonal possible worlds.)

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.

Eckhardt, W., 1997. A Shooting Room View of Doomsday. *Journal of Philosophy*, May 1997, 94. (Also available online at: [http://williameckhardt.com.p2.hostingprod.com/yahoo\\_site\\_admin/assets/docs/William\\_Eckhardt\\_-\\_A\\_Shooting\\_Room\\_View\\_of\\_Doomsday.26095737.pdf](http://williameckhardt.com.p2.hostingprod.com/yahoo_site_admin/assets/docs/William_Eckhardt_-_A_Shooting_Room_View_of_Doomsday.26095737.pdf) [Accessed 3 January 2010]).

Gott, J. G. III, 1993. Implications of the Copernican principle for our future prospects. *Nature*, 363, pp.315-319.

Grace, K., 2010. *SIA and the Two Dimensional Doomsday Argument*. [Online] Meteuphoric. Available at: <http://meteuphoric.wordpress.com/2010/11/12/sia-and-the-two-dimensional-doomsday-argument/> [Accessed 2 January 2011].

Lewis, D. K., 1986. *On the Plurality of Worlds*. Oxford: Basil Blackwell.

Mitra, S., 2010. *Entangled states considered as physical representations of classical algorithms*. [Online] arXiv:1009.4472v1. Available at: <http://arxiv.org/abs/1009.4472> [Accessed 30 October 2010].

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](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\\_siam.pdf](http://space.mit.edu/home/tegmark/PDF/multiverse_siam.pdf) [Accessed 5 December 2010].)

The self-indication assumption almost stops the Doomsday argument. Almost. V2.0

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]).