

ON WHAT IT MEANS TO HAVE THE SAME CONCEPT

Wylie Breckenridge

This is an essay about concept identity. In particular, it is about this question: What do I mean when I judge, as I sometimes do, that one of your concepts is the same as one of mine? I have a thought, for instance, that I express by saying "Peter likes chocolate". You have a thought that you express by saying "Mary hates chocolate". I judge that the concept you expressed by "chocolate" is the same as the concept that I expressed by "chocolate". What do I mean?

I will consider attempts to provide an answer in terms of the extension of our concepts in various contexts and possible worlds. I will argue that there is one sense of "the same" for which quite a good answer can be given, but there is another sense for which it can't - a sense whose meaning cannot be explained extensionally.

What is a concept?

I will be saying a lot about concepts, so I want to begin by setting out as clearly as possible what, for the purposes of this essay, I will take them to be.

I will take a concept to be something that only one person can have. If I have a concept, you cannot have the numerically same concept. It's mine. If you have a concept, I cannot have the numerically same concept. It's yours. That doesn't mean, however, that our concepts can't be tokens of the same type. I think that they can. In fact, I must think that they can in order to judge, as I do, that our concepts are the same. I will return to this issue in the next section. For now, I want to say what *my* concepts are.

I will define concepts in terms of thoughts and components of thoughts. If I express a thought by saying "Peter likes chocolate", let me refer to it either as the thought that Peter likes chocolate, or just as "Peter likes chocolate". I believe that my thoughts have *components*. I believe, for example, that "Peter likes chocolate" has a component that I express using the word "chocolate". Let me refer to this component as "chocolate". In general, let me refer to a thought component by the word that I use to express it. I will take my thoughts to be (mental) events, so that the time at which they occur is essential to them: if I think sometime today that Peter likes chocolate and think so again sometime tomorrow, I will have had numerically distinct thoughts, and each component of the first will be numerically distinct from each component of the second.

I believe that my thought components (all of them - the ones that belong to thoughts that I have already had, and the ones that belong to thoughts that I will have in the future) fall naturally into disjoint classes. I believe that the component "chocolate" of my thought today that Peter likes chocolate is in the same class as the component "chocolate" of my thought tomorrow that Peter likes chocolate, and also in the same class as the component "chocolate" of my thought today that Mary hates chocolate. That class is what I call my concept CHOCOLATE (after the word that I use to express its members) (I shall write concept names in upper case). In general, by one of my concepts I mean any one of the classes into which I believe my thought components naturally fall. That is what I will take my concepts to be. What I take *your* concepts to be is more-or-less the same, just a few obvious changes.

I have defined my concepts in terms of thoughts, thought components, and an ability that I have to sort thought components into classes. Don't I owe an explanation of these things as well? In particular, don't I owe an explanation of the last one? Why do I think that the component "chocolate" in "Peter likes chocolate" belongs to the same class as the component "chocolate" in "Mary hates chocolate", but to a different class than the component "Mary" in "Mary hates chocolate"? It seems natural to try and explain this in terms of what concepts they express - to say, for example, that they belong to the same class because they express the same concept. But it would be circular for me to do that - I have said that they express the same concept because they belong to the same class. So how do I explain it?

I don't think that I need to. I think that when I talk about thoughts, thought components, and this ability to sort thought components into classes, it should be fairly clear what I am talking about. Clearer, at least, than when I talk about concepts. Moreover, I think that the claims I have made about them are plausible, and even common-sensibly true. So I'm happy to define concepts in terms of them. I'm not after a rock-solid definition, just one that makes it clear what I'm talking about. None the less, I'll give an analogy that I use to think about concepts. It might help you to understand what I'm getting at in this essay. I think of my concepts as stamps (rubber stamps, not postage stamps). I use them to stamp-out thoughts by stamping-out thought-components. I have one stamp for each concept: a stamp for PETER, a stamp for CHOCOLATE, a stamp for LIKES, a stamp for PETER LIKES (which I make by taping together PETER and LIKES), and so on. To have the thought that Peter likes chocolate, I grab the stamps PETER, LIKES, and CHOCOLATE (or, maybe, PETER LIKES and CHOCOLATE) and stamp-out the thought. When I claim that two thought-components belong to the same class I am claiming that they were stamped out with the same stamp. How do I know that they were? Easy. I know which stamps I used to stamp them out.

To be sure, let me say some more about concepts. Each of my thought components belongs to exactly one concept (in terms of my analogy, is made by exactly one stamp). For example, the component "chocolate" of "Peter likes chocolate" belongs to my concept CHOCOLATE and to no other. Sometimes (when it will not cause confusion) I will say that a thought component simply *is* the concept to which it belongs, and that when I express the component I am expressing the concept. Distinct thought components may belong to the same concept (may be made by the same stamp). For example, "chocolate" in "Peter likes chocolate" belongs to the same concept as "chocolate" in "Mary hates chocolate" (they both belong to CHOCOLATE). Sometimes I will say that they simply *are* the same concept (CHOCOLATE), and that when I express them I am expressing the same concept. I will also say that the thoughts to which they belong *employ* or *use* or *contain* the same concept. Distinct thought components may also belong to different concepts (may be made by different stamps). For example, "Peter" in "Peter likes chocolate" belongs to a different concept than does "Mary" in "Mary hates chocolate". Sometimes I will just say that they are different concepts.

I believe that "Peter likes chocolate" employs six concepts: PETER, LIKES, CHOCOLATE, PETER LIKES, LIKES CHOCOLATE, and PETER LIKES CHOCOLATE. I name them after the words that I use to express them: I express the first by saying "Peter", so I call it "PETER"; I express the second by saying "likes", so I call it "LIKES"; and so on. I believe that "Mary likes chocolate" also employs six

concepts: MARY, LIKES, CHOCOLATE, MARY LIKES, LIKES CHOCOLATE, and MARY LIKES CHOCOLATE. I believe that the concept LIKES in the first case is identical with the concept LIKES in the second. By this I mean that the component in the first that I express using the word "likes" belongs to the same concept as the component in the second that I express using the word "likes" (they are made with the same stamp). I also believe that CHOCOLATE in the first case is identical with CHOCOLATE in the second, and that LIKES CHOCOLATE in the first is identical with LIKES CHOCOLATE in the second. I believe that no other identifications can be made between the concepts employed by the first and the concepts employed by the second.

As these examples suggest, I take it that concepts can themselves have other concepts as components. For example, PETER LIKES has PETER and LIKES as proper components, and PETER LIKES as a non-proper component. Moreover, I take it that a concept can be a complete thought - that is, that concepts are not restricted to being proper components of thoughts. For example, I take it that PETER LIKES CHOCOLATE is a concept. Perhaps there is some class of concepts which have no proper components and from which all other concepts can be composed. Call these, if there are any, *simple* concepts. It would no doubt be interesting to know that such concepts exist, and perhaps they would assist and/or simplify some theories of concepts. But I am not concerned about that here. I will treat all concepts as equal in this respect.

What is it to have the same concept?

Now to return to the question that is my main concern: What do I mean when I say that your concept is the same as mine?

I claim that there are at least two senses in which I make such judgments. I'll call them the *identity sense* and the *semantic sense*. If I judge them to be the same in the first sense then I'll say that they are *identically the same*, or that they are the same concept. If I judge them to be the same in the second sense then I'll say that they are *semantically the same*, or that they have the same meaning. I think that the two senses cut across each other. That is, I think that it's possible for me to judge that our concepts compare in each of the following four ways:

		Identically the same?	
		Yes	No
Semantically the same?	Yes	Same concept Same meaning	Different concept Same meaning
	No	Same concept Different meaning	Different concept Different meaning

So what are these two senses, and why do I think that they cross cut? Some examples of judgments that I make should be enough to answer both questions.

From here on I will sometimes consider the way in which I compare my own concepts to investigate the way in which I compare mine and yours. I believe that they are closely related. Perhaps they are at root the same - perhaps comparing my own concepts amounts to me putting myself in your shoes, or perhaps comparing mine with yours amounts to me putting you in mine. I don't know, but I will assume that the way in which I compare my own concepts can illuminate the way in which I compare mine and yours.

I will start, then, with some examples of how I judge my own concepts to be the same or different in each sense. I have a concept NATURAL NUMBER. I believe it is identically the same as itself and semantically the same as itself - same concept, same meaning. I also have a concept PRIME NUMBER. I believe that NATURAL NUMBER and PRIME NUMBER are identically different and semantically different - different concepts, different meanings. They are the easy examples. Now a harder one. My concept NATURAL NUMBER used to have a different meaning. I used to think that being a natural number meant being a positive integer. Now I think that being a natural number means being a non-negative integer (a subtle but non-trivial difference). I believe that the concept I used to have and the concept that I now have are identically the same but semantically different - same concepts, different meanings. That I think this way is illustrated by what I say. I say that my concept NATURAL NUMBER has changed, not that I have replaced an old concept NATURAL NUMBER₁ by a new concept NATURAL NUMBER₂. And I say that my concept NATURAL NUMBER used to have a different meaning, not that I used to have a concept with a different meaning. Now I've given an example for all but one of the four cells in my table. The fourth example will come later.

Next, some examples of how I judge my concepts to be the same as or different to yours in each sense. You and I are arguing about whether or not 0 is a natural number. I say that it is, you say that it's not. Scenario 1: We agree that a natural number is defined to be any non-negative integer; you think that 0 is negative; I point out that it's not; you agree, and change your mind about it being a natural number. Scenario 2: I discover after some questioning that you think a natural number is defined to be any *positive* integer (rather than non-negative), and that's why you think that 0 is not a natural number; I convince you that it's better to define it to be any non-negative integer instead; you come to agree, adjust your definition, and change your mind about 0 being a natural number. Scenario 3: I decide after some questioning that you have been talking about prime numbers (well, what I call "prime numbers" - you, for some reason, have been calling them "natural numbers"); I explain what I have been talking about, you start talking about the same thing, and agree that 0 is a natural number.

In each scenario I am comparing your concept that you express as "natural number" with my concept that I express as "natural number". In the first I decide that they are identically and semantically the same. You disagree with me about 0 because you mistakenly think that it is negative. In the second I decide that our concepts are identically the same, but semantically different. You disagree with me about 0 because you disagree with me about what it is to be a natural number - your concept has a different meaning for you than mine does for me. In the third I decide that our concepts are identically and semantically different. You disagree with me about 0 because you have been arguing about whether or not it is a prime number - about whether or not it has divisors other than 1 and itself - and I have been arguing about whether or not it is a

natural number - about whether or not it is a non-negative integer. You have been employing a different concept with a different meaning. Again, it remains for me to give an example of the fourth possible judgment (different concepts, same meaning). That will come later.

Because there are two senses in which I judge our concepts to be the same, my original question splits into two:

- (Q1) What do I mean when I say that your concept is *identically* the same as mine?
- (Q2) What do I mean when I say that your concept is *semantically* the same as mine?

By saying that they are identically the same, I can't, according to the way that I have defined a concept, mean that your concept is numerically identical with mine in the way that I say two of my own concepts are numerically identical: my concepts are classes of my thought components, your concepts are classes of yours, my thought components are mine and mine alone, yours are yours and yours alone. So maybe it's a bad name. But I wanted it to sound stricter than being semantically the same, and I couldn't think of anything better. I don't think it will cause confusion.

I'm interested in answers to Q1 and Q2 that give necessary and sufficient conditions under which I judge that our concepts are the same in each sense. That is, I'm interested in finding conditions that are necessary and sufficient for me to make these judgments:

- (Id) Our concepts are identically the same.
- (Sem) Our concepts are semantically the same.

Of course, if we do find such conditions then it doesn't follow that we have found out what I *mean* by Id and Sem. After all, just because a native says "Gavagai" if and only if he sees a rabbit, it doesn't follow that "Gavagai" means he's seen a rabbit - it might mean, for example, that he's seen a rabbit's ear. But finding such conditions is a necessary step towards finding their meanings, and if we can't do that then we can't find their meanings either. This first step is the one that I am interested in.

One possibility comes quickly to mind:

Condition 0

You use the same word to express your concept as I use to express mine (that is, you use the same word to express the thought components that belong to your concept as I use to express the thought components that belong to mine).

I need only reflect on the way that I express my own concepts to realise that words are not a reliable guide to which concepts a person is expressing or to what meanings those concepts have, and that condition 0 is therefore no good. It is not *sufficient* for either Id or Sem. Why not? I have concepts that I believe are identically and semantically different, and yet express using the same word "goal". One I express when I say "I

kicked a goal today", but not when I say "My life has no goal". The other I express when I say "My life has no goal", but not when I say "I kicked a goal today". So just because you express your concept using the same word as I do, I will not conclude on that basis that our concepts are identically or semantically the same. Nor is condition 0 *necessary* for either Id or Sem. Why not? I have a concept that I express sometimes using the word "beer" and sometimes using the word "coldie". I express it using "beer" when I say "I'd love a beer", and I express it using "coldie" when I say "I'd love a coldie". In each case, I believe I'm expressing the one concept with the one meaning. So just because you express your concept using a different word than I do, I will not conclude on that basis that our concepts are identically or semantically different. So condition 0 is neither sufficient nor necessary for either Id or Sem.

No matter, this is not the sort of condition I'm interested in anyway. I'm interested in ones that are given in terms of the *extensions* of our concepts. What is the extension of a concept? I need to say what I mean. I believe that every one of my concepts has an extension. What kind of thing the extension is depends on the concept. The extension of PETER is a person: Peter. The extension of LIKES is a class of ordered pairs: the class of all ordered pairs (x, y) such that x likes y (examples are (me, television), and (dad, beer)). The extension of CHOCOLATE is a substance: chocolate. The extension of PETER LIKES is a class: the class of all things y such that Peter likes y. The extension of LIKES CHOCOLATE is also a class: the class of all things x such that x likes chocolate. The extension of PETER LIKES CHOCOLATE is a truth value: true if Peter likes chocolate, false otherwise. That should be enough to make it clear what I mean by the extension of a concept.

Here's the sort of condition that I'm interested in:

Condition 1

The extension of your concept is the same as the extension of mine.

Again, I need only consider my own concepts to see that this condition is no good. It is not sufficient for Id. Why not? I have concepts that I think are identically different and yet have the same extension. One is LIVING ANIMAL WITH A HEART, and the other is LIVING ANIMAL WITH A KIDNEY. The extension of the first is the class of all living animals that have a heart. The extension of the second is the class of all living animals that have a kidney. I believe that as it turns out, every living animal that has a heart has a kidney, and every living animal that has a kidney has a heart, so these classes are the same and the concepts have the same extension. Identically different concepts, same extension. So just because your concept has the same extension as mine I will not conclude on that basis that they are identically the same. Condition 1 is not necessary for Id. Why not? Compare my old concept NATURAL NUMBER with my new one. I think they are identically the same and yet have different extensions - 0 is not in the extension of the old one but it is in the extension of the new one. Identically the same concepts, different extensions. So just because your concept has a different extension to mine I will not conclude on that basis that they are identically different. Thus, condition 1 is neither necessary nor sufficient for Id.

Condition 1 is not sufficient for Sem. Why not? I believe that my concepts LIVING ANIMAL WITH A HEART and LIVING ANIMAL WITH A KIDNEY have the same extension, and yet I judge them to be semantically different. Semantically different

concepts, same extension. So just because our concepts have the same extension I will not conclude that they are semantically the same. One might think that it is, however, *necessary* for Sem - that I will judge our concepts to be semantically the same *only if* I think they have the same extension. This because it seems plausible that a concept's meaning determines its extension. But we can see that this is wrong by considering a concept like **HERE**, whose extension (what it refers to) is *context* dependent. It is quite possible for me to think that **HERE** has the same meaning for me and you - the location of the context in which it is thought/uttered - and yet think that the extension of your **HERE** is different to the extension of mine. This will happen whenever your location is different to mine. Then we have a case of semantically the same concepts, different extensions. Another example. It is possible for me to think that **WATER** has the same meaning for both of us - the substance, for example, that has the same underlying structure as watery stuff (the clear flowing liquid in rivers and oceans, etc.) - and yet think that your **WATER** has a different extension to my **WATER**. Suppose we disagree about what that underlying structure is - I think it's H₂O, you think it's XYZ. Then if there is a black tarry substance in the world whose underlying structure is XYZ, I will agree with you that it is in the extension of your concept, but I will think that it's not in the extension of mine. I will think, that is, that we have semantically the same concepts with different extensions. So condition 1 is not, after all necessary for Sem.

Concepts like **HERE**, whose extension is context-dependent, suggest that we need our condition to accommodate that dependency:

Condition 2

In every context, the extension of your concept is the same as the extension of mine.

What do I mean by the extension of a concept in a context? It's a straightforward generalization of the idea of extension that I described earlier. Consider again my six concepts **PETER**, **LIKES**, **CHOCOLATE**, **PETER LIKES**, **LIKES CHOCOLATE**, and **PETER LIKES CHOCOLATE**. The extension of **PETER** in context C is a person: whoever is Peter in the world of C when **PETER** is employed in C (why do I need to say "in the world of C"? Because we will soon be allowing contexts to vary across possible worlds). The extension of **LIKES** in context C is a class of ordered pairs of things in the world of C: the class of all ordered pairs (x, y) such that x likes y when **LIKES** is employed in C. The extension of **CHOCOLATE** in context C is a substance: whatever is chocolate in the world of C when **CHOCOLATE** is employed in C. The extension of **PETER LIKES** is a class: the class of all things y in the world of C such that Peter in the world of C likes y when **PETER LIKES** is employed in C. And so on.

HERE is not the problem for condition 2 that it was for condition 1. Although I believe that in different contexts your **HERE** and my **HERE** might have different extensions, I believe that in the same context they have the same extension.

If by *every* context we mean to include contexts in other possible worlds as well as contexts in the actual world, then **WATER** is not the problem for condition 2 that it was for condition 1. The disagreement between you and me about the extension of **WATER** is the result of our disagreement about the underlying structure of watery stuff, which can be seen as a disagreement about which context (in particular, which world) we are in. I think we are in a context in which watery stuff is H₂O, you think we are in a context

in which watery stuff is XYZ. But we both agree what the extension of WATER is in every context. In particular, I believe that when your WATER and my WATER are employed in the same context they have the same extension.

Also, by including contexts in other possible worlds, the concepts LIVING ANIMAL WITH A HEART and LIVING ANIMAL WITH A KIDNEY are not the problem for condition 2 that they were for condition 1. Even though they have the same extension in the actual world (and, therefore, in every context in the actual world), I believe there are possible worlds (and, therefore, contexts in those worlds) in which they don't. I believe that in some worlds some living animals have a kidney but not a heart, in some worlds some living animals have a heart but not a kidney, and in some worlds there are both sorts of animals.

NATURAL NUMBER is still a problem for condition 2. Recall, I believe that my old concept and my new concept are identically the same but have different extensions. This shows that condition 1 is not necessary for Id. But I also believe that their extensions are context-independent (i.e. they are independent of the context in which they are employed in all possible worlds), so I also believe that in every context they have different extensions. This shows that condition 2 is not necessary for Id either. But as far as we can tell at the moment, condition 2 is both necessary and sufficient for Sem (we have no counterexamples at the moment).

Not so fast. I have concepts that I believe are semantically different and yet have the same extension in every context. They are WATER and WATERY STUFF. By WATERY STUFF I mean any substance that is clear, colourless, free-flowing, etc. By WATER I mean any substance that has the same underlying structure as the watery stuff around here (i.e. in my context). These, I believe, are different meanings. That's why I judge them to be semantically different concepts. The extension of WATERY STUFF in a context C includes exactly those substances in the world of C that are watery - i.e. are clear, colourless flowing liquids, etc. The extension of WATER in a context C includes exactly those substances in the world of C that have the same underlying structure as the watery stuff in that world. These extensions are the same in every context. So I have two concepts that I believe are semantically different and yet have the same extension in every context. Therefore, just because you have a concept that has the same extension as mine in every context, I will not conclude on that basis that it is semantically the same as mine. So condition 2 is not sufficient for Sem.

We can fix this problem by comparing the extensions of the two concepts not just in each context, but in every context-world pair:

Condition 3

In every context-world pair, the extension of your concept is the same as the extension of mine.

What do I mean by the extension of a concept in a context-world pair? Again, it's a straightforward generalization of the extension of a concept in a context. Consider a context C and a world w. The extension of PETER in (C,w) is a person: whoever is Peter in w when PETER is employed in C. The extension of LIKES in (C,w) is a class of ordered pairs of things in w: the class of all ordered pairs (x, y) such that x likes y when LIKES is employed in C. The extension of CHOCOLATE in (C,w) is a

substance: whatever is chocolate in w when CHOCOLATE is employed C . And so on. It might seem unnecessary to specify the context in which a concept is being employed when evaluating its extension in some world, because it will only depend on the world in which it is being evaluated, and not on the context in which the concept is being employed. But we will see in the next paragraph that that is not so.

How does condition 3 avoid the problem that condition 2 has with WATER and WATERY STUFF? Although I believe they have the same extension in every context, there are some context-world pairs in which I believe they have different extensions. I believe, for example, that I am in a context in which water is H_2O (i.e. the underlying structure of the watery stuff in my context is H_2O). I also believe there is a possible world in which the underlying structure of watery stuff is XYZ, and in which the substance whose underlying structure is H_2O is black and tarry. I believe that in that world the extensions of WATERY STUFF and WATER are different - XYZ is in the first but not the second, and the black tarry stuff is in the second but not the first. So my semantically different concepts WATER and WATERY STUFF fail condition 3, as they should if it hopes to be sufficient for Sem.

We are now making some progress with finding necessary and sufficient conditions for Sem - for me to judge, that is, that our concepts are semantically the same. Furthermore, comparing the extensions of concepts on a context-world pair basis turns out to be quite informative about concept meanings. It promises to explain, for example, the notions of *a priori truth* and *a posteriori necessary truth*. I will not go into that here. Note that we still have some work to do, though. I have, for example, concepts that I believe to be semantically different and yet have the same extension in all context-world pairs. They are TRIANGLE and TRILATERAL. Being a triangle and being a trilateral mean different things to me - the first means having three angles, and the second means having three sides. That's why I judge them to be semantically different. But I believe they have the same extension in every context-world pair. No matter in which context I employ them, the extension of TRIANGLE in any possible world is the class of all three-angled plane figures in that world, and the extension of TRILATERAL in any possible world is the class of all three-sided plane figures in that world. But I believe that, as a matter of necessity, a plane figure has three angles if and only if it has three sides. So the extension of these concepts is the same in all possible worlds and for all contexts. So this problem needs to be addressed before we can claim to have a necessary and sufficient condition for Sem. It is plausible to think, however, that we might make some progress by continuing in the same direction - by adding more dimensions to the way in which we compare the extensions of the concepts.

The prospects are not looking so bright, however, for finding a necessary and sufficient condition for Id. We saw that NATURAL NUMBER shows that condition 1 is not necessary for Id - I judge my old concept to be identically the same as my new one and yet believe that they have different extensions. We saw that condition 2 suffers the same problem, and it is easy to see that condition 3 does as well. So none of them is necessary for Id. The problem is that although I judge my old concept and my new concept to be identically the same, I believe they are extensionally different in the actual world, in every context, and in every context-world pair. The reason is simple: 0 is everywhere. As long as we have 0, there seems to be no way of getting their extensions to agree, as there must be if we want an extensional condition that is *necessary* for Id.

There is a problem trying to find a *sufficient* condition as well. In fact, there can't be one. Why not? My argument is quite simple. You say that you're about to teach me two new concepts that you call "BLIM" and "BLAM". Or maybe you say that you're about to tell me a story involving "blim" and "blam", terms that are new to me. As soon as you do so, I create two concepts of my own, concepts that I judge to be identically different. Apart from following you in calling them "BLIM" and "BLAM", I know nothing at all about them. In particular, I know nothing at all about their extensions. But I still judge them to be identically different. Andrew, having overheard you, says "I know all about blim and blam." I judge my BLIM to be identically the same as his BLIM, and identically different to his BLAM. But I know nothing at all about the extension of my BLIM or his BLIM or his BLAM. So the conditions under which I judge his BLIM to be identically the same as my BLIM can't have anything to do with their extensions.

There are objections that I must address.

First, it might be objected that I *don't* have two new concepts yet - I don't have a concept until I know what it means, or at least until I have one belief that employs it. I think that's wrong. On the one hand, I have many concepts whose meaning I don't know - CAUSATION, for one. It's not that I know what it means for *me* but not what it means for *anyone else*. I don't even know what it means for *me*. On the other hand, if I can't have a concept until I have a belief that employs it, then I couldn't have any concepts at all: if I don't have the concept then I can't have any thoughts that employ the concept (by my definition, at least), and so I can't have any beliefs that employ the concept, because beliefs are just thoughts. Or maybe I should just say that I *do* have at least one belief that employs each concept: that my mum either is or is not in the extension of the concept, for instance. So I reject this objection.

Second, it might be objected that I only judge my BLIM and my BLAM, or my BLIM and your BLAM, to be identically different because I assume (perhaps without realising) that their extensions are in some yet-to-be-discovered way different. I think that's wrong too. I think that you could *tell* me that their extensions coincided exactly, in all contexts and all worlds and whatever you like, and yet they would still be identically different to me. You could even tell me that they have the same *meaning* and I would still take them to be identically different. That possibility provides the fourth example that I have been promising: BLIM and BLAM would then be concepts that are identically different but semantically the same - different concepts, same meaning.

Of course, if I did come to believe that they had the same extension and/or the same meaning then I might well decide that I can do without one and just keep the other. Perhaps we do that all the time. Perhaps that's what we did with THE MORNING STAR and THE EVENING STAR. And perhaps, for that reason, the only sense in which I'm interested in whether or not your concept is the same as mine is the semantic sense - perhaps all that I need to know to effectively communicate with you is that our concepts are semantically the same, although not necessarily identically the same. Even so, there is still one sense in which I judge our concepts to be the same - the identity sense - for which there are no conditions on the extensions of our concepts that are necessary and sufficient for me to make such judgements. It follows that what I mean when I judge our concepts to be identically the same cannot be explained in terms of the extension of our concepts.

In fact, it cannot be explained in terms of *anything* that you can tell me about your concept. Nothing you tell me about it - its extension, its meaning, or anything else - forces me to judge them to be identically the same. Nor does anything you tell me force me to judge them to be identically different. There is nothing you can say that forces me either way. The reason is simple. I can, if I want, *massively* change what I believe about one of my concepts without changing its identity. I might change its meaning, I might change its extension, I might change the truth-value of every statement that employs it. But I can, if I want, keep its identity. Or, if I want, I can scrap the concept altogether and create a new one, with the meaning and extension that I want, and use it in place of the first in all of my thoughts. There may be pragmatic reasons for choosing one course of action over the other, but it's a choice that I always have. Likewise, I can always choose to judge that one of your concepts is identically the same as mine, even though we *massively* disagree about it. When I judge our concepts to be identically the same I mean nothing more, I suggest, than that I have *decided* to take them to be identically the same. It is a claim about what I have decided to do, and *nothing* about your concept or mine can make it true or false.