

2018

The Similarity of Division

Gedalia Zemel
Bard College

Recommended Citation

Zemel, Gedalia, "The Similarity of Division" (2018). *Senior Projects Fall 2018*. 37.
https://digitalcommons.bard.edu/senproj_f2018/37

This Open Access is brought to you for free and open access by the Bard Undergraduate Senior Projects at Bard Digital Commons. It has been accepted for inclusion in Senior Projects Fall 2018 by an authorized administrator of Bard Digital Commons. For more information, please contact digitalcommons@bard.edu.

The similarity of division

Senior Project submitted to
The Division of Social Studies
of Bard College

by

Gedalia Zemel

Annandale-on-Hudson, New York

Dec 2018

Outline	4
Chapter 1: an introduction to our distinction.	8
Defining the key terms	8
Some examples to illustrate our distinction.	9
Some more kinds of physical claims	12
Claims about universal laws such as claims about gravity	12
Chapter 2: the mind independent external world	16
Mind independent external world	18
Chalmers	18
Solids are made of empty space	18
Does “being in a simulation” even mean anything?	21
Confusion about words	22
Being in a simulation doesn’t mean anything? It sure meant something in The Matrix	23
Which world is simulated?	24
Back to the structure of Chalmers's argument	27
Chapter 3: Ayer	34
Outline of this section	34
Outline of similarities with Ayer	34
Differences with Ayer	35
Problems Ayer introduces	35
Intro to Ayer	36
Comparison with Ayer	40
Claims like (3): claims about universal laws such as claims asserting the existence of gravity	48
Claims like (4): claims about loose laws such as general rules governing how mammals learn.	50
Using Ayer’s loose verificationist principle to solve for (3) and (4)	51
The verificationist principle	52
Chapter 4: A closer look at our comparison with Ayer.	61
A summary of how our claim currently compares to Ayer’s claim.	61
A summary of how we will further distinguish our claim from Ayer’s claim.	63
A closer look at the differences between Ayer’s claim and our own claim.	65
Chapter 5: review and further discussion.	78
Further discussion	80
Bibliography	84

Outline

I introduce a fairly simple distinction about one way to divide up truth claims. That is to say, I propose that claims can be divided up into one of two disjoint sets. The goal of this project is not to argue in favor of that distinction *per se*, but to consider our distinction's implication in comparison to other similar distinctions.

The distinction I propose is that one way we can divide up claims, is into the set of "claims about physical things" and into the set of "claims about non-physical things" and that those two sets of possible claims will be disjoint.

Obviously, our distinction requires further definition but the distinction itself is fairly simple. Most of our discussion will be spent considering the implications of our distinction and comparing the implications to those of similar distinctions, namely with the distinction between meaningful claims and meaningless claims put forward by Ayer in *Language Truth and Logic*.

In what way is our distinction similar to Ayer's distinction and in what way is our distinction different than Ayer's distinction? Are there conditions¹ under which Ayer's

¹ By conditions under which Ayer's distinction will be more similar to our own, I refer to Ayer's definition of rationality. Given Ayer's definition of rationality it might follow that Ayer's distinction (between meaningful and meaningless beliefs) will be applied differently depending on which standards of belief formation one considers to be reliable.

distinction will be more or less similar to our own? What are the advantages and disadvantages of Ayer's distinction in comparison to our own?

These are the kinds of questions we will try to answer. There are two main ways we will compare and contrast our distinction with Ayer's distinction.

By "compare and contrast" I mean firstly to compare how other distinctions divide up a set of claims into two disjoint sets of claims².

The second way that we will compare our distinction to that of Ayer is by examining the other philosophical commitments presumed in our distinction. When you stake out an area around one philosophical position you are often, both implicitly and explicitly, making claims about other philosophical positions simply as a necessary consequence of your initial claim. So does our proposed distinction force us to take positions we don't want to take? How do the presumed commitments of our distinction compare to the commitments of Ayer's distinction?

The actual chapters of our paper will proceed as follows.

² As we've said, our distinction divides up claims into two groups. We can similarly use Ayer's distinction to divide up potential claims into two groups. So then, we examine how the membership of the two groups of claims produced by our distinction compares to the two groups of claims produced by Ayer's distinction.

In the first chapter we will introduce our claim.

In the second chapter we introduce a major problem for our claim, the apparent presumption a mind independent external world. We introduce Chalmers's claims about "realness" in the case that we are inside of a simulation. We point out some disagreements we have with Chalmers but eventually we use Chalmers's argument to resolve our problem. We explain why our claim does not presume the existence of a mind independent external world.

In the third chapter we introduce Ayer's distinction between meaningful and meaningless claims. We point out that, with the exception of tautological claims, Ayer will seemingly divide up claims the same way that we will divide them up. That is we will both divide up a set of claims into two identical sets. We call this similarity the 'similarity of division'.

In the fourth chapter we reexamine the 'similarity of division'. We take a closer look at Ayer's use of 'probability' and 'rationality'. We then present a case where, given Ayer's specific usage of 'probability', the similarity of division will not hold.

In the fifth chapter we review the paper and suggest some points for further discussion.

Chapter 1: an introduction to our distinction.

Some claims are about the physical world; we call these claims about the physical world “physical claims”. Other claims are not about the physical world; these are nonphysical claims.

We say that a physical claim is a claim about the world itself whereas a non-physical claim is *only* about the way we categorize things in our minds and not about the world itself.

Defining the key terms

What exactly do we mean by “claims about the physical world”; what are these “physical claims”? Let’s begin by introducing the criteria for “physical claims”.

The criteria for a “physical claim” is any claim where the world will itself be different corresponding to the truth value of the claim. That is, supposing two possible worlds, one for each of a claim’s two possible truth values, the world where the claim is true will itself be different from the world where the claim is false.

So then, a physical claim is one where each of the truth values of the claim corresponds to different possible worlds.

To determine if a claim is a physical claim we need only ask ourselves “will there be some difference in the world itself corresponding to the truth value of the claim”. If the answer to that question is “yes”, that a world where the claim is true will itself be different than a world where the claim is false, then we say that the claim is a physical claim. If the answer to that question is no, that a world then the claim is true will not itself be different based on the truth value of the claim, then the claim is a nonphysical claim.

Our explanation of our distinction and the criterion behind it aren't even close to clear enough. We need some examples.

Some examples to illustrate our distinction.

(1) : “there is a chair over there”

To determine whether this claim is a physical claim we will simply ask, for this claim about the location of a chair, “will a world where this claim is true look different than a world where this claim is false”.

Yes! It would appear that a world where the claim is true will in fact look different than a world where the claim is false. In a world where the claim is true *there will be a chair* at the specified location. In a world where the claim is false *there won't be a chair* at the specified location³. There will in short be a difference in the placement of at least one chair corresponding to the truth value of the claim: there will in this case be a difference in the world corresponding to the truth value of the claim.

So then, because there is a difference in the world itself corresponding to each truth value of the claim it would follow that the claim is a physical claim. So then we can say that (1) was a physical claim. Let's move on to another example.

(2): "killing puppies is evil"

Again, to determine whether this claim is a physical claim we need only to ask, "will a world where this claim is true look different than a world where this claim is false". Let's consider what the world will look like corresponding to the two possible truth values of (2).

³ There is some stuff going unsaid. For example, we assume competent language speakers. Likewise we assume that when two competent language speakers disagree about whether there is a chair over there they are disagreeing about an observable matter of fact and not about the definition of "chair" or "there". That is to say that claims like "there is a chair over there" are true in the case of a there being a chair but false in the case of there not being a chair.

In comparison a disagreement over whether "evil" exists in the world doesn't hinge on something observable in the world. That two competent language speakers disagree about such a claim isn't about something in the world but about how the category "evil" is applied. That is to say that the truth value of claim about the existence of evil (again between language speakers) doesn't hinge on something in the world itself.

In a world where the second exemplar claim is false I will categorize the act of killing a puppy as non-evil. That is, I will mentally categorize the action of killing a puppy as non-evil.

What will a world where the second claim is true look like? In a world where the claim is true I will categorize the act of killing a puppy as evil.

Nothing in the world itself changes based on the truth value of this second exemplar claim. When two people argue about the truth value of this claim they aren't (in most cases) arguing about something in the world itself but about what things fall into the mental category of evil. Because there is no difference in the world itself corresponding to the truth value of the second claim we say that the second claim is not a physical claim. The second claim is a claim about non-physical things, or a non-physical claim.

The application of our criterion is simple. To determine if a given claim is a claim about physical things, or a claim about nonphysical things we need only to answer two questions and then make a single comparison. We must ask "what will a world where this claim is true look like" and "what will a world where this claim is false look like". From there we need only compare the responses of the two questions. If the responses to both questions is the same, that is to say that there will be no difference in the world

corresponding to the truth value of the claim, then the claim is a claim about nonphysical things. If however the response to both questions is different, that is that there is some difference in the world corresponding to each of the two truth value, then the claim is a claim about physical things.

Some more kinds of physical claims

There are two more kinds of physical claims that should at least be mentioned here. For the time being we will not apply our criteria to these two kinds of physical claims. That will have to wait for the end of the project. What we will do though is introduce each of the two kinds of case with an exemplar case and briefly explain how these two cases are different than our first exemplar case with the chair. Again, the application of our criteria to these two cases will come at the end of the project. Our goal here is merely to acknowledge the existence of these two major categories of physical claim without immediately jumping into them.

Claims about universal laws such as claims about gravity

(3) Any claim about gravity for example. That is to say any claim about a universal law. That a world where a certain claim about gravity is true will itself be different than a world where that claim about gravity is false seems obvious enough at first. If the claim is false, I don't know, people start floating around. If the claim is true then things go on as usual. Of course though, things aren't that simple. For example suppose some understanding of gravity makes lots of accurate predictions but makes a

single incorrect prediction. Do we dispose of the theory or do we assume auxiliary hypothesis which are consistent with the theory?

Whereas the application of our criterion in the chair case was seeming simple, the application of our criterion is more difficult in the gravity case. That is to say that it's difficult to approximate what the world would look like for each of a universal claim's truth values.

I should point out that we are not in any way claiming that there are in fact any such universal laws. We are merely acknowledging that universal laws such as gravity might exist, and that if they did in fact exist, that claims about universal laws would be claims about physical things.

The application of our criterion to such cases will be a challenge. However, as I said, we will deal with that towards the end of our project. Let's quickly look at the final case. The final case will be claims like:

(4) "mammals are particularly sensitive to aversive stimulus related to things that they eat".

In case (4) we propose a kind of rule, but with (4) unlike with (3) the rule is less stringent: the rule is only a general rule but that there are exceptions to the rule. That is

for example, the truth value of the hypothesis clearly isn't contingent on any specific event in the world. For example suppose the following case:

- I drink some Coca Cola
- I become violently ill
- This happens multiple times
- I then later drink Coca Cola, again I sense no aversion, I have no Gag reflex etc.

Surely in such a case we would all agree that we have not falsified (4). In other words, (4) expresses a rather uncontroversial claim. However, there is no single event, or even reasonably small sets of events that could falsify that claim. So then the question becomes, how would a world where (4) is true look different than a world where (4) is false?

We will wait until we have dealt with Ayer to examine these last two kinds of claims. Put simply according to Ayer the problems of (3) and (4) are equally applicable to (1). That is according to Ayer, no claim can ever be conclusively verified or falsified. We agree with Ayer on this point. In fact we take it a step further and claim that even tautologies can never have their truth values conclusively settled.

Given then that claim by Ayer, a claim which we agree with, the problems of (3) and (4) will be equally applicable to (1), and to all physical claims. That is, we will see that there is no way to conclusively verify or falsify a physical claim.

Right now we have a more pressing contradiction to deal with: our apparent presumption of mind independent external worlds.

Chapter 2: the mind independent external world

Before we can go on there is one problem that must be dealt with; it is the philosophical elephant in the room. The problem is, it would seem, that our distinction presumes the existence of a mind independent external world.

In our description of the two kinds of claims we seem to be presuming the existence of a mind independent world. After all we do say that claims about physical things are claims about the world itself; however, in the case of claims about nonphysical things we say that the claim is not about the world itself, rather that it is merely a claim about how we mentally categorize things in our mind.

So then it would seem to follow from the explanation of our distinction that physical things exist in the world itself. Conversely it would likewise seem to follow that non-physical things don't exist in the world itself, but only exist in our minds. In other words, it seems to follow from our distinction that things that exist in the world itself are different than things that exist in our mind. In other words, we seem to be indirectly saying that things which exist in the world itself exist outside of our minds. We seem to be implying that the physical world exists outside of our mind. *We seem to be presuming the existence of a mind independent external world.*

So what? Say that our distinction does in fact presume the existence of a mind independent external world, is that really a problem? Yes.

The problem with supposing the existence of a mind independent external world is that doing so would, firstly, force us to take a position on heavily contested philosophical ground. Now, that would be acceptable if we had some knock down argument in favor of the existence of a mind independent external world. We don't have any such argument. And so taking a strong position in such an uncertain area without any strong auxiliary arguments would be a bad argumentative move. Put simply we don't want our argument to hinge on a single contentious claim for which we have no supporting argument.

Additionally, I wouldn't even be asking you to take my word for it without a supporting argument. I have no clue if there is a mind independent external world; it would take me awhile to work out if "mind independent world" is even a meaningful concept. I don't want to stake our project's claim on a single auxiliary claim which may or may not even mean something. If we want our distinction to be even slightly plausible we need to separate our distinction from claims about mind independent external worlds.

Mind independent external world

Of course, as you might have guessed, our claim doesn't actually imply the existence of a mind independent external world. Our claim is in fact totally neutral on all questions of mind independent external worlds; to explain why we need to borrow an argument from Chalmers.

Chalmers

Prima facie it seems to me that the salient point in Chalmers's argument (the one we are about to look at) is about philosophy of language and doesn't really have much to do with "mind independent external worlds".

That is Chalmers is making a general principled claim about the nature of meaning. So Chalmers is pointing out a kind of confusion around certain kinds of abstract ideas. At the same time though Chalmers, on the object level appears to be dealing with claims very much like our own. Chalmers applies a general principle to a specific kind of case (that we are living in a simulation) and we borrow that principle to explain why our distinction doesn't presume the existence of a mind independent external world.

Solids are made of empty space.

Chalmers is interested in a hypothetical case where we are in fact living in a simulation. Were it the case that we are in fact living in a simulation, Chalmers asks, would it then follow that our universe isn't "real"?

In such a case there would seemingly be reason to believe that our world is not real. After all the world simulating our world, that is the world our simulation is happening inside of, is the real world (or at least one step closer to the real world). What we naively thought of as chairs and tables would then really be bytes, or magic, or whatever makes our simulated world exist. Really though the table is simulated; the table is not real. The table isn't made from real stuff, the table is made from simulated stuff, the table is made of bytes. We are simulated and unreal, or at least that's how it might seem upon first analysis.

Chalmers though disagrees with such an assessment of this hypothetical simulated world. According to Chalmers, even in the case that we were living in a simulation it would not then follow that our world is not real; rather it would follow that what we really meant when we spoke about realness is different than what we had once thought.

In other words, when we said "real 1" we really meant "real 2". This concept of "by x we really meant y" (where in this case x is our understanding of "real" prior to knowledge of the simulation and y is our understanding of "real" after knowledge of the

simulation) is the subject of our interest in Chalmers. Chalmers introduces this principle though by going back to an historical case.

It was once the case, Chalmers says, that we thought that things like tables and chairs were totally solid.

If we looked close enough we were able to observe that things like tables and chairs, solid things, are in fact almost entirely made up of empty space. Yet we still believe that things like chairs and tables are solid. The discovery that empty space made up most of tables and chairs did not make tables and chairs any less solid.

We did not decide that chairs were not solid. Instead we continued to believe that chairs were solid but that our understanding of solid was, at a deeper level, different than we expected it to be. Whereas we once thought that solid things were things without any empty space, we later realized that solid things are simply things without easily visible empty space.

So Chalmers says, the same is true for the computational simulation hypothesis:
"The Computational Hypothesis says that physics as we know it not the fundamental level of reality. Just as chemical processes underlie biological processes, and microphysical processes underlie chemical processes, something underlies

microphysical processes. Underneath the level of quarks and electrons and photons is a further level: the level of bits."(40).

In other words, in the same way that we say that chairs are still solid (despite the fact that they are mostly made up of empty space) so too we would say that that which we take to be real (what we call our "reality") is still real. In the case of the chair we realized that what it means to be solid is different than what we thought it meant to be solid; in the case that we were in a simulated reality we would realize that reality is different than we thought it was. What it means to be real, at the deeper level, is different than we *expected* real to mean.

Does "being in a simulation" even mean anything?

Our interest though isn't in such a case. Our interest is in the conceptual mechanism that allowed Chalmers to say "by x we really meant y". We don't really care about the object level case, in fact, I'm not sure it's even sensical. If we're going to take Chalmers's conceptual framework though we should at least be clear about the object level case that we disagree with. That is we agree with Chalmers on the nature of meaning but not on the case he applies it to. We totally disagree with Chalmers on the case of the simulation. As it stands I don't, to my best understanding at least, think that this idea of finding out that we are in a simulation is at all sensical. What would that knowledge of being inside of a simulation look like?

Confusion about words

Some people say that the question of consciousness is pretty much nonsensical, that not only do we not know the answer, but we can't even imagine what an answer would look like. That is the question of what physical processes give rise to experience doesn't make sense in that we don't even know what a *possible answer* could look like.

In other words we suppose some entity x exists. We believe that x is the answer to the question of how physical things give rise to experience. However, there is no x , or if there is an x we have no idea what it could possibly look like. We don't know the *kind* of thing that it is. We could not invent any x that could in theory explain how physical things give rise to experience. The idea of being in a simulation is a lot like this x . That is like x we can seemingly have meaningful conversations about being inside a simulation, but also like x , when look closely we see that there is no possible answer to the question "what does it mean to be inside a simulation".

When speaking about simulations we consider the idea that we are in a simulation without any obvious problems. That is, we assume that the idea of being in a simulation means something (in other words we assume that it's a sensible idea). We assume that there is a case that we would classify as being in a simulation, but what would such a case look like? I claim that in the same way that we don't know the kind

of answer which would explain the physical-experience problem so to we don't even know the kind of criteria which would qualify as "being in a simulation".

Being in a simulation doesn't mean anything? It sure meant something in The Matrix!

What? Of course we know what it's like to be in a simulation! For example the matrix, which Chalmers's paper on realness inside of a simulation is about, captures it perfectly. Neo is born in a simulation and until he meets Morpheus, that simulation is where he lives. But eventually Neo meets Morpheus, he takes the red pill, goes to the pay phone and exits the simulation.

This appears pretty sensical and totally counter to my claim that "being in a simulation is nonsensical". After all Neo is working in an office, he is in a simulation. When he takes the red pill and disappears through the phone to Zion he is isn't in a simulation, right?

No. Let's call the supposed "simulated" reality where Neo is a 9-5 office worker "office world". We will call the so called "real world", the one where the robots and Zion are, "Zion world". Now let me ask you, what makes you think that Zion world is non simulated and that office world is simulated? Yes, yes, Morpheus told Neo that office world was simulated, and that Zion world was not simulated, but we're doing philosophy here so we need to do better than that; we can't take Morpheus's word for it.

Which world was simulated?

At first glance we want to point out that Neo began in one simulated universe and stepped through the payphone into a non-simulated universe, right? OK, so he was in one world and he stepped into another and so the first world must be a simulation.

Right?

Well, suppose the following counterfactual case: suppose that office world was the non-simulated world and that Zion world was the simulated world. Would the movie look any different?

I don't think so. We assume that Neo stepped into the phone booth in *simulated office world* and woke up in *non-simulated Zion world*. However, we have no reason to suppose that Neo really stepped into the phone booth in *non-simulated office world* and stepped out in *simulated Zion world*. That is, there is nothing to differentiate the simulation from the non-simulation. Nothing of course except for Morpheus's word, but that's not nearly enough for us.

In other words, everything we saw in the movie was equally compatible with either world being the simulated world. Because we are not able to distinguish a case where Zion is simulated and office world is non-simulated from a case where the roles are reversed, we have no reason to believe that one world as opposed to the other is

simulated. At this point we have only shown that one cannot distinguish the simulated world from the non-simulated world but we still assume that one of the worlds, even if we don't know which one, is simulated, right?

No. There is nothing to distinguish going from a simulated world to a non-simulated world from the case of going from one world to another world (in the same universe) via some, I don't know, let's say transport beam. Likewise we couldn't distinguish moving from simulation to non-simulation from, let's say, moving from one universe to another. To the best of my knowledge though I have no idea what could be thought of as being in a simulation. I don't know what it would mean to live in a simulation.

We might want to say that we know that one world is simulated and that the other world isn't simulated because the movement of bytes in one world corresponds to every state and change of state in the other world. . This doesn't work because there are many counterexamples which can demonstrate that just because there are bytes corresponding to everything in a given reality it does not follow that those bytes are simulating the reality matching those bytes.

First, supposing that it's possible for some universe to simulate another universe, then the same should be true of our universe. Our universe should seemingly be able to

simulate other universes. More specifically our universe should be able to simulate our own universe itself.

That is suppose that someone announced tomorrow that they are going to create reality z, a simulated world that looks exactly like our own. In such a case we would all agree that we exist outside of that specific reality z. That is, there are bytes corresponding to everything in our universe and yet we know that that specific reality z isn't simulating our own reality. So we see that just because bytes correspond to things in one reality it doesn't follow that that specific reality is itself a simulation.

Now I know that some people will claim that while such a case wouldn't demonstrate that we are inside of that specific reality z that we are still very likely, inside of another reality z one level higher up. This is not however a problem, because we have still not explained what it means to be inside of a simulated reality. We know that the mere correspondence of bytes doesn't mean that one reality is simulated but perhaps there are other conditions?

We could go on a lot more with this; although this line of reasoning explains our disagreement with Chalmers it isn't complete. However, we have provided the basic outline of our disagreement with Chalmers although we didn't consider every possible counter example. The point is that we disagree with Chalmers's object level case and question whether "being inside of a simulation" is even a meaningful term.

Back to the structure of Chalmers's argument

Forgetting about what “being in a simulation” even means, our interest is in Chalmers’ more general principle. Chalmers said that when we discovered that solids were made up of empty space our understanding of what it really meant to be solid changed. I propose that Chalmers's argument can be used to resolve our apparent claim on the existence of an external world: Chalmers's argument explains why our distinction doesn’t actually commit us to the existence of the external world.

Let’s look back at the basic structure of Chalmers’s argument. Once we have the structure of his argument we can apply it to our case. As it stands now our distinction seems to imply the existence of a mind independent external world. We will use Chalmers's argument to explain why that is not the case, why our distinction is neutral on the existence of mind independent external worlds and why our distinction is totally compatible with the possibility that there is no mind independent external world. Let’s take a look at the basic structure of Chalmers’s case.

First Chalmers introduces a term which names both a set of things, and which also corresponds to a specific definition. In one of Chalmers's examples the term was “real”. The term named things like ourselves as well as things like the universe we live

in. In Chalmers example the term also had a definition which excluded things simulated in bytes. The set of things we apply the term to have the features of that term's definition: namely our universe was real in that it wasn't simulated because we assumed that our universe wasn't simulated.

But then, in Chalmers example we found out that the set of things named by the term don't have the properties which the term was supposed to define: we found out that the things included in "real" (our universe) are in fact simulated. All of the things we called "real" didn't fit into our definition of real.

Nevertheless we realized that the term does in fact name all the things we thought it named, but that it's definition was at a deeper level different than we thought it was. Our universe was still real. All that changed was our definition of "real": "real" we discovered included simulated universes.

In both Chalmers's example of realness as well as his example of solidity, the term has both a naming function and a definition. The naming function serves to distinguish one kind of thing from the other. For example "solid" serves to distinguish things like chairs from things like water. The definition serves to explain why things fall under the term, like how the definition of solid can be used to determine that a chair is solid but that fumes are not solid.

When we found out that the things which fall under the naming function don't at a deeper level fall under the definition, we changed the definition to better fit the things it named. I propose that we can do the same thing to our distinction if in fact there is no mind independent external world. **More specifically I propose that the purpose of our distinction is to distinguish claims like those about chairs from claims like those like about evil; the purpose of our distinction is not to make claims about the essential nature of physical things.**

In the example of "real" the term had a function and a definition. Later we realized that the definition didn't match the function so we changed the definition. In our case though we have a function and a definition which matches that function; however the definition is only meant to fulfill its function and it does so without making claims about mind independent external worlds. In the case that there is no mind independent external worlds, however, our definition can be changed to include mind independent external worlds.

We first introduced our distinction to perform a certain function. There are two kinds of things that we want to separate from each other. We want to separate claims which we think of as corresponding to things in the world from claims which don't correspond to things in the world.

In other words we want to separate claims about “evil” and perhaps certain kinds of aesthetic claims, as well as certain kinds of formal claims, from claims that we can simply look at the world to determine their true value. The purpose of our distinction is merely to serve that function: to separate these two kinds of claims from one another.

With that in mind let’s explain why our distinction doesn’t assume the existence of a mind independent external world. Let’s suppose then that there is no mind independent external world. Would it then follow that the things we call “physical things” are in fact no different than the things we call “non-physical things”?

I don’t think so.

You see, even if there in fact is no mind independent external world, our distinction would still serve its function. It’s function is to distinguish claims about things like chairs from claims about things like evil. We say that physical claims are about things that exist in the world itself and that non-physical things exist only as mental categories because that’s a useful way to distinguish claims about things like chairs from claims about things like evil. The definition doesn’t say anything about mind independent external worlds at all; it simply performs its function of distinguishing these two kinds of claims. Were it the case that there were no mind independent external world we would still define physical claims and non-physical claims the same way

because they would still serve the important function of distinguishing two kinds of claims from one another.

In the case that there was no mind independent physical world there would still be things which are physical and things which are nonphysical: there would still be things that exist in the world itself and things which exist only as mental categories. The function of those terms would not change. They would still differentiate two kinds of things and two kinds of claims from one another. All that would change is that, at a deeper level, what we think of as being physical and existing in the world would, in a small way, be different than we might have expected.

Whereas we might have at first thought that physical things were things that existed in the world and not in our thoughts, we would, given no mind independent external world, simply change our deeper definition in a very small way. We would simply say that “physical world” is the physical part of the mental world. That is, given the existence of a mind *dependent* external world we would simply explain the distinction between physical claims and nonphysical claims as follows.

Both physical things and non-physical things, everything, would exist only in the mind. What we call physical and non-physical would both just be subsets of things which are really mental. We would then divide up the world which exists, again, only in the mind, into the subset of that mental world which is physical and into the subset of

that mental world which is non-physical. Just because there is no physical world doesn't mean that evil and chairs are both thoughts in the exact same way. At the very least we can admit that you can cut a chair in half with a saw but that you can't cut evil in half with a saw. Physical things would be different than mental things in the same way that they always were.

In other words, even if we assume that there is no mind independent external world, there would still be claims about chairs and about evil, and only one of those two types of claims could still be evaluated by looking at the world. All that has changed now is that when we say 'physical world' we now mean the part of the mental world that has the properties related to physical things. Likewise when we say "non-physical things", or mental things, we mean the area of the mind dependent world which can't be evaluated by looking at the world. When we would speak about physical things we would simply mean the kind of claims that we could look at the (essentially mental) physical world to resolve.

So then, we make no commitments to the existence of a mind independent physical world. That we distinguish between physical things and mental things is merely a functional definition; there is an important difference between those two kinds of things. The important difference between those two kinds of things, these two kinds of claims, is that we can determine their truth values in different ways. That we determine

their truth values differently holds the same even in the case that there is no mind independent external world.

When we distinguish physical things from mental things we are not implying that physical things are not at a deeper level mental. We are merely saying that there is a difference between physical things and mental things, and that the difference would remain even in the case that there was no mind independent external world. Indeed all that would change in such a case is that physical would mean the part of the mental world that is physical, while non-physical would mean the part of the mental world that *only* has to do with things like mental categorization. So then our distinction is totally neutral on the existence of mind independent external world.

There is no problem of assuming the existence a mind independent external world. Our distinction doesn't assume the existence of a mind independent external world. Our distinction is totally compatible with the existence of the mind independent external world. This resolves our discussion of the mind independent external world. Now we move onto our comparison with Ayer.

Chapter 3: Ayer

Outline of this section

In his explanation of logical positivism, Ayer introduces a criterion to determine the meaningfulness of a claim. The criterion is Ayer's loose verificationist principle which he uses to distinguish meaningful claims from nonmeaningful claim. In this section we will compare our distinction to Ayer's loose verificationist principle. We will first consider similarities, then differences, finally we will consider some points brought up by Ayer which might present problems for our distinction. First though an outline of the three parts of this chapter.

Outline of similarities with Ayer

We introduce Ayer's verification principle, which, like our distinction, can divide up the set of all possible claims into two groups. I will argue that the membership of those two groups is, with the exception of tautological claims, mostly identical to the two groups produced by our distinction⁴.

⁴ That is in other words that, for a given set of non-tautological claims, our distinction and Ayer's verification principle will divide up the set of possible claims into two identical groups.

Differences with Ayer

After we have reviewed the similarities between our distinction and Ayer's verification principle we will proceed to the differences. An important difference I will argue is in the attribution of meaning and meaninglessness. That is, while our distinction makes no commitments on the meaningfulness of nonphysical claims, Ayer holds that all non-tautological but non-physical claims, are meaningless.

One other major difference rests in Ayer's understanding of rationality and the function of factual claims which Ayer uses to redefine probability. It will follow from those three things that Ayer might allow claims to be grouped differently depending on what method of forming beliefs one considers reliable. In this case it can in theory follow that Ayer's test will group claims very differently than we would group them. It might follow that our distinction only groups claims similarly to Ayer's when one, like Ayer, considers science a reliable method for assessing beliefs.

A more trivial difference will be come from the way Ayer categorizes tautologies. However, that isn't of much interest to us.

Problems Ayer introduces

Once we have finished our comparison with Ayer we will consider some problems for our distinction brought up indirectly by Ayer.

One problem for example will come from Ayer's dismissal of the strict verificationist principle. Ayer points out that no empirical claims can even in theory be completely falsified/verified. Ayer's criticism of the strict verification principle would prima facie be applicable to our distinction. Our distinction will have to answer to that criticism. We will need to consider our argument that claims about physical things can have their truth value determined simply by comparing them to the actual world. It would seem, correctly, that given Ayer's argument, comparison between claim and world isn't sufficient to settle a claim's truth value.

Intro to Ayer

Before we begin our comparison with Ayer we need to introduce his criterion.

Ayer's introduces a kind of test meant to be applied to a claim: "We need only formulate the criterion which enables us to test whether a sentence expresses a genuine proposition about a matter of fact, and then point out that the sentences under consideration fail to satisfy it."(35). That is on the basis of Ayer's test we are able to determine if a given claim about a matter of fact is or is not meaningful. I should quickly note, I mean that all non-tautological claims are either meaningless or are matters of fact. So then, put another way, we can simply apply the test to a given claim and determine if the claim is in fact meaningful. Let's take a look at the actual test.

Ayer broadly⁵ introduces his test: “We say that a sentence is factually significant to any given person, if, and only if, he knows how to verify the proposition which it purports to express—that is, if he knows what observations would lead him, under certain conditions, to accept the proposition as being true, or reject it as being false.”(35) . What Ayer means is that a given claim is meaningful only when a future sensory experience can in theory confirm or disconfirm the belief. Put another way, a claim is only meaningful if it is less consistent with some potential future experiences than it is with others. A meaningful claim should, in short, put some kind of a constraint on expected future experiences such that it is compatible with some sensory experiences but not with others.

However, “If, on the other hand, the putative proposition is of such a character that the assumption of its truth, or falsehood, is consistent with any assumption whatsoever concerning the nature of his future experience, then, as far as he is concerned, it is, if not a tautology, a mere pseudo-proposition.” (35). In other words, if a given claim is consistent with every possible future sensory experience then the claim is meaningless. Let’s take a step back.

⁵ I say that ,“Ayer broadly introduces his test”, because Ayer adds more and more nuance to the test as he goes on and because what we are now presenting is an unnuanced introduction to Ayer’s test. There is a lot of nuance for example in how Ayer defines “verify”. Likewise we couldn’t fully understand what Ayer means by “observations would lead him” unless we know about Ayer’s unusual usage of “probability” which in turn comes from his definition of rationality, which in turn...you get what I’m doing.

The point I’m trying to make in this footnote is that we won’t even attempt to provide a fully nuanced view of Ayer’s test, although we will eventually provide an understanding of Ayer that is nuanced in areas related to our distinction . Afterall, the entire “language truth and logic” is an explanation of the nuanced version of Ayer’s test and the test’s application and would be way too much to include: “the adoption of this procedure is an essential factor in the argument of this book, it needs to be examined in detail” (36).

To this point we have only given a general overview of Ayer's test, or criterion. While there are some important nuances that we will add as we go on, for the time being let's try to get a good understanding the current unnuanced version of the test. Let's consider two example cases, and illustrate how Ayer's test is applied to the exemplar case. Here I will use the same two exemplar claims that were previously used to demonstrate the application of our own criterion.

(1): "There is a chair over there".

To apply Ayer's test we need to ask ourselves, are there any future observations that would make us more or less likely to believe (1)? In this case there are some such obvious potential observations. For example, if I observe that there is no chair at the specified location I will be much less confident that (1) is true. Conversely if I look at the specified location and see a chair I will be more confident that (1) is true. (1) is thus meaningful because it is consistent with some observations but is not consistent with other observations.

Put another way, there are predictions following from (1), namely the prediction amounting to the sensory experience of seeing a chair; these predictions can be tested by observing the specified location of the chair.

I should just point out so that there is no confusion that according to Ayer, sensory experience can never *conclusively* verify/falsify any claim. When we say that a certain sensory experience isn't consistent with a certain claim we simply mean that it increases the probability that the claim is false; it doesn't mean that we have conclusively demonstrated that the claim is false. Likewise, there is never any case where we can settle a claim as conclusively verified. The exception to these rules are tautologies which are, according to Ayer, true or false with complete certainty⁶.

Let's look at our second exemplar claim.

(2): "killing puppies is evil"

⁶ Ayer says that tautologies have definitive truth values. I'm not sure if I agree with Ayer on this point. The disagreement stems from the following hypothetical counter example, although there are many other similar counterexamples.

We use the term "working memory" to broadly describe the amount of variables a person is able to manipulate at a given time. That definition isn't perfect but it's more than close enough. Now considering a tautology like "all unmarried men are bachelors" requires us to manipulate a couple of variables. We have to take 'bachelor' and ask ourselves, what is that terms definition. We then remember that bachelor is defined as unmarried man.

The problem is that the previous operations require us to competently manipulate two variables; something which is broadly speaking measurable. So then, it follows that our ability to assess tautologies is contingent on a matter of fact. Our ability to assess tautologies is contingent on our working memory. However, that our working memory can correctly manipulate two variables can never be conclusively verified. It is afterall an empirical claim. So then our ability to assess tautologies is contingent on something which can never be conclusively verified. It would then seem to follow that tautologies cannot be conclusively verified. *However then, since tautologies are not themselves empirically falsifiable, though contingent on empirically falsifiable claims, it will then follow that tautologies are not merely trivially meaningful claims, but that tautologies are meaningless in the same way that, according to Ayer, ethical claims are meaningless.*

Some quick notes about (2). We say “killing puppies is evil” where “evil” is *only* a value judgment, but not a predictive tool such as it might be in the case that a violent mob is killing “evil” people, or to use a more realistic example in that case that “evil” people receive social sanction. In other words, in (2) we talking about the attribution of a moral value.

Again, to apply Ayer’s test we need to ask ourselves, are there any future observations that would make us more or less likely to believe (2)? More appropriately are there any predictions made by (2) such that their apparent success/failure will alter the probability of (2) being true? Prima facie there are no such predictions. I can think of no observation that claims about evil would preclude. (2) is thus is consistent with all experiences, so then we say that (2) is meaningless.

This concludes our overview of Ayer’s criterion. We will now compare Ayer’s criterion to our own and in doing so I will argue that Ayer’s criterion is similar to our own.

Comparison with Ayer

In our thesis we proposed that that all claims can be divided into two sets. They are the set of claims about physical things, and the set of claims about non-physical

things. We proposed a test to determine which of the two types a given claim is. Our test for a claim's type was simply to ask of any given belief, "will there be some difference in the world itself corresponding to the truth value of the claim". If there is a difference in the world itself corresponding to the truth value of the claim, we said that the claim is a claim about physical things, that it is a physical claim. However, if there is no difference in the world corresponding to the truth value of the claim, we said that the claim is about non-physical things, that the claim is a non-physical claim. And so for any set of claims we can use this test to divide up the claims into these two sets. Ayer divides up the set of potential claims in a similar way.

Like us, Ayer introduced two types of claims along with a test to divide potential claims: "We need only formulate the criterion which enables us to test whether a sentence expresses a genuine proposition about a matter of fact, and then point out that the sentences under consideration fail to satisfy it"(35). Ayer, like us, can divide up a set of claims into two sets: the set of meaningful claims, and into the set of meaningless claims.

I propose for the time being⁷, that, with the exception of tautological claims, the membership of Ayer's two sets will be identical to the membership of our own two sets⁸.

⁷ We will see that this might not be the case. If we use Ayer's version of rationality to define "probability" it might follow that the similarity between our two sets and Ayer's two sets will be totally contingent. Contingent that is on what a given person considers to be reliable methods of producing predictions.

⁸ That is in other words, that for a given set of claims both Ayer's criterion and our own criterion will divide up that set of potential claims into two sets, and that for the membership of each one of those two sets, its membership will be identical to the membership of one of Ayer's two sets.

That this is true seemingly follows from each of our criterion. Ayer's criterion for the meaningfulness of a claim, put one way, is that there be some possible sense experience that is inconsistent with the given claim: " 'Would any observations be relevant to the determination of its truth or falsehood?' . And it is only if a negative answer is given to this second question [the second question being the question in the previous sentence] that we conclude that the statement under consideration is nonsensical"(38) That is in other words that for a given claim, there should be some sense experience it is consistent with and other sense experience that it is not consistent with. In other words that our sense experience of the world should be different based on the truth value of the claim.

That criterion for the meaningfulness of a claim is, you might notice, very similar to our criterion for physical claim: a claim is a physical claim if some area of the world will be different corresponding to the truth value of that claim. This isn't exactly like Ayer's criterion but it's close. While our criterion holds that a claim is physical if something in the world changes based on that claim's truth value, Ayer's criterion holds that a claim is meaningful if some potential sense experience is inconsistent with the claim.

More specifically, the membership of our set of claims about physical things will be identical to the membership of Ayer's set of meaningful claims; likewise the membership of our set of claims about non-physical things will be identical to the membership of Ayer's set of non-meaningful claims.

So then our criteria for physical claims hinges upon the relation between a claim's truth value and the possible worlds corresponding to it. Ayer's criteria hinges upon the relation between a claim's truth value and future sensory experience. So then our individual criterions might appear to be different, after all a sensory experience is a different thing than the actual state of the world, right? Maybe. That will depend on our answer to the question, are there any parts of the physical world that can be different corresponding to different truth values but which nevertheless produce no differing potential sensory experience?

That last paragraph was a mouthful, we will need to break it down into a simpler form. Suppose the set of all possible physical claims. It will be true that every member of that set, every physical claim, will have some difference in the world corresponding to each of its truth values. What we want to know then is, for all of those potential differences in the world corresponding to each truth value, is there any pair of such corresponding worlds which cannot, even in theory, be distinguished with a sensory experience? This is what we mean when we ask, will every difference in the world correspond to only some potential sensory experience but not to other potential sensory experiences? Put more simply, can every difference in the world potentially be distinguished by sense experience?

Looking for claims that will help us answer this question, one claim comes to mind.

The claim I'm thinking of is the question, "a tree falls in the forest and no one is around to hear it, does it make a sound"⁹? We can alter that question a bit so that it's making a claim: "when a tree falls in the forest and no one is around to hear it, the tree makes a sound". Also let's assume that by sound we don't mean the experience of sound, but the vibrations. Let's see if Ayer will handle this claim the same way we will handle it

According to Ayer the claim will be meaningful only if there are some potential sensory experiences which are inconsistent with the claim. There is, however, no obvious sense experiences which is inconsistent with the claim. After all, the claim is specifically about a case where there is no one around to have immediate sensory experiences of the falling tree. According to Ayer then the claim will prima facie be meaningless¹⁰. Remember, if our criterion classifies this claim as a physical claim, we

⁹ The idea to use this example to illustrate empiricism/positivism wasn't my idea. I saw it in a blog post "Making Beliefs Pay Rent (in Anticipated Experiences)" by Eliezer Yudkowsky. There he uses the example to argue that we should reduce beliefs to what they can predict and that we should make beliefs "pay rent" by making accurate predictions. The point is that I got the idea to use this example from Yudkowsky even though he uses it to make another empiricist claim.

¹⁰ We said that, according to Ayer, the claim, "when a tree falls in the forest and no one is around to hear it, the tree makes a sound", would be meaningless. We said that the claim would be meaningless because there is no sense experience that can be inconsistent with that claim given that the claim precludes any sense experience. After all the claim is about a sound which no one is around to hear. So then it seems like, according to Ayer, this claim would be meaningless.

However one might object that in theory there might still be sense experiences that are inconsistent with this claim. For example, one might imagine that there might have been a recorder at the scene of the falling tree. We only said that no one was around to hear the tree fall, but there might still have been an

will have demonstrated that our criterion is different than Ayer's criterion. How then would we apply our criterion to this claim?

To apply our criterion to this case we need to ask ourselves, will a world where this claim is true itself be different from a world where this claim is false? More specifically though we said that the world where the claim is true will only be different than the world where the claim is false *if* we are able to look at the world and determine which of the two possible worlds we are in. However, we could only distinguish which of the two worlds we live in on the basis of sensory experience. In other words, our criterion, like Ayer's criterion, really distinguishes between claims which are consistent with all sensory experience from claims which are inconsistent with some sensory experience.

audio recorder recording the sound of the falling tree. Given then that possibility, there is at least in theory some future sense experience that is inconsistent with the claim. That is the sense experience of listening to a potential audio recording of the fallen tree might be inconsistent with the claim, for example in the case that a recording of the fallen tree has no falling tree noises. So it would seem that our exemplar claim is in fact potentially inconsistent with some sense experience. It should then follow that the claim about the falling tree is meaningful according to Ayer.

To this can simply respond by amending our claim. The purpose of the exemplar case was to provide a possible physical claim that is meaningless. That is to say that we want to provide a claim that might be about the world itself but is definitely not inconsistent with any sense experience. That's not to say that the claim is a physical claim, but simply that the claim might be a physical claim but will require further analysis.

So then, if one is going to claim that our exemplar claim is still in theory inconsistent with some sense experiences we can amend the claim so as to be consistent with any sense experience. For example we might say, "when a tree falls in the forest and no one is around to hear it, neither person nor recorder, the tree makes a sound". Better yet we might say that, "when a tree falls in the forest and no one is around to hear it, nor are there any potential sensory experiences that connected to the sound of the falling tree, the tree makes a sound". In doing so we would preserve the purpose of our example. The claim still might be a physical claim but is at the same time clearly not inconsistent with any sensory experience. The claim is still meaningless. We can thus modify our exemplar claim in the case that one argues that the claim is inconsistent with some sensory experiences.

Let's review then where this leaves us. We wanted to know if our distinction was unlike Ayer's. Specifically we wanted to know if there are any claims about the world itself which are nevertheless consistent with all sensory experience. Such a case would have demonstrated that our criterion was in certain cases different than Ayer's criterion. However we remembered that we define a difference in the world itself as a difference such that we are able to determine which of the two worlds we are living in by observing the world. In other words, our criterion, like Ayer's, really distinguishes claims on the basis of sensory experience. Our distinction is then to this point, though worded differently, still similar to Ayer's criterion, that is it is similar in the sense that both criterions divide up potential claims into matching groups.

In short then it does indeed seem to follow that for every difference in the world corresponding to a claim's truth value there must also be a sensory experience. In short then, all meaningful claims would appear to also be claims about physical things.

Before we really get into the major substantive differences between our criterion and Ayer's criterion there is one more comparison we need to make. We need to consider Ayer's "loose verificationist principle". But first, let's recall one unresolved problem from the first chapter. The problems brought up by our previously unresolved problem will lead into our discussion of Ayer's "loose verificationist principle". That is in other words we will review an previously unresolved problem before introducing Ayer's "loose verificationist principle"; the principle will then help us resolve our problem.

You might remember from our introduction that we decided to hold off on the discussion of two kinds of physical claims until later in the paper. We called these claims “(3)” and “(4)”. Claim (3) included all claims about universal laws such as claims about gravity. Claim (4) included claims about non universal laws which are still mostly true, like claims that mammals are extremely sensitive to aversive stimulus that are related to food; what’s salient about claims like (4) is that they are mostly true but certainly incorrect at times.

There was however a challenge to applying our criterion to these two cases. The problem was that it was difficult to imagine how exactly a world where claims like (3) and (4) are true will be different than a world where claims like (3) and (4) are false.

Sure in cases like (1), the case of the chair, it’s seemingly easy to distinguish a world where (1) is true from a world where (1) is false. But in the case of both (3) and (4) it’s a bit more difficult to imagine how exactly a world where the claim is false will look different than a world where the claim is true.

Let’s review why it is difficult to distinguish a world where claims like (3), or (4) are true from a world where claims like (3), or (4) are false. We begin with (3).

Claims like (3): claims about universal laws such as claims asserting the existence of gravity

Consider claims like claim (3). That is claims about universal laws, such as for example, claims asserting the existence¹¹ of gravity. Remember, we said that claims like (3) are claims about physical things. However, to meet our standard for “claims about physical things” there must be some difference in the world corresponding to the truth value of the claim. That is, for (3) to be a claim about a physical thing it must be the case that we can determine the truth value of the claim by looking at the world. Yet, in the case of claims like claims about gravity, it’s hard to pinpoint how exactly a world where the claim is true will be different than a world where the claim is false. Let’s apply our criterion to (3.1) and see what problems come up.

(3.1) “gravity exists”, also assume that “gravity” represents a specific predictive formula that is meant to make accurate predictions for a certain set of cases.

First we can ask ourselves, what will a world where (3.1) is false look like? We might then say that in a world where (3.1) is false, that objects will move in such a way that we cannot use relations of mass to predict their motion. Sure, there are less extreme cases where we might say that (3.1) is false, however, we can all agree that in

¹¹ We use ‘existence’ but ‘existence’ might be a bit too strong of a word. I want to make clear I’m not claiming that gravity is a thing that exists in the same way that say a chair might exist. When I see that gravity exists I simply mean something like “this gravity abstraction is an accurate predictive tool”.

the previous case (3.1) is clearly false. In such a world we can clearly say that there is no universal law of gravity.

Then we ask ourselves, what will a world where (3.1) is true look like? To this we can say, that in the case that some formula for gravity can predict all movement based on relations of mass, that (3.1) is clearly true. That is, all predictions based on gravity will be correct without any need for auxiliary hypothesis.

So then we might claim, (3.1) is clearly a physical claim. After all there is a big difference between these two worlds. And so we might think that claims like (3.1) can be simply settled; such claims are claims about physical things because we can distinguish the truth value of the claim by looking at the world.

However, we notice that this (3.1) is different than (1), the claim of the chair.

‘ (1) “there is a chair over there” ‘

In that case we said that if the claim was true, the chair would be in the specified location. If the claim was false there would be no chair in the specified location. So then we can seemingly always look at the world to determine the claim’s truth value. In the case of the chair we were *seemingly* always able to distinguish the world where (1) was

true from the world where (1) was false. In the case of (1) we took it for granted that we could distinguish a world where (1) was true from a world where (1) was false.

However, in the case of claims like (3.1) it's much harder to distinguish the truth value of the claim by looking at the world. Save for extreme cases such as a possible world where some formula for gravity always makes perfect predictions without any auxiliary hypothesis, it's unclear how we can look at the world to determine the truth value of claims asserting gravity. In cases like (3.1) it's unclear if we can look at the world to determine the truth value of the claim.

In summary the problem of cases like (3.1), claims about universal laws, is that it's unclear whether one can distinguish the truth value of the claim by looking at the world. Whereas in our simplified case (1) it seems to be that case that we can always simply look at the world to determine the truth value of the claim, it's unclear whether looking at the world can always resolve claims like (3.1). I for one couldn't provide the exact conditions under which (3.1) has a specific truth value.

Claims like (4): claims about loose laws such as general rules governing how mammals learn

Let's move onto our consideration of cases which propose general rules which allow for some exceptions.

(4) “mammals are particularly sensitive to aversive stimuli related to things that they eat”

The problem with claims like (4), that is loose rules, is that it's unclear how exactly a world where the claim is true will be different from a world where the claim is false. For example (4) would be true even if it didn't hold for a single person; it would even be true in the case that it didn't hold for a single species. How then, given that we acknowledge possible small exceptions to the rules, can one distinguish a world where (4) is true from a world where (4) is false? How can we say that claims like (4) are physical claims, when we are unsure how to distinguish a world where (4) is true from a world where (4) is false? For claims like (4) to be claims about physical things we must be able to determine the truth value of the claim by looking at the world, however, it's uncertain which observations would be sufficient to determine its truth value.

Using Ayer's loose verificationist principle to solve for (3) and (4)

In the previous sections we discussed a problem with assigning the status of “physical claim” to claims like (3), and (4). The problem was that for a claim to be a “physical claim” we must be able to distinguish a world where the claim is true from a world where the claim is false. For claims like (3), and (4), however, we are unsure how to go about distinguishing worlds where they are true from worlds where they are false.

Up until this point we have taken for granted that this problem facing (3), and (4), isn't applicable to (1). That is, in the case of (1) we are always able to distinguish a world where (1) is true from a world where (1) is false. To do that we need only observe the placement of the chair; that information alone is enough to determine the truth value of (1). So then (1), unlike (3) and (4), meets the criterion for being a physical claim.

In what follows we will review Ayer's loose verificationist principle. Ayer's loose verificationist principle will force us to answer some questions we have to this point left unanswered. As we will see our criterion, like Ayer's, requires something like his loose verificationist principle. Once we have clarified the kind of verificationist principle we will use we will see that the questions facing (3) and (4) are also applicable to (1). Then with some clarification we will settle the question facing (1), (3), and (4).

The verificationist principle

According to Ayer a claim can only be meaningful if there are potential sensory experiences which are inconsistent with the claim. The "loose verificationist principle" we're about to look at describes what it means for a sensory experience to be inconsistent with a claim.

Ayer's introduces his criterion, "he knows what observations would lead him, under certain conditions, to accept the proposition as being true, or reject it as being false"(35): under certain conditions, certain sensory experiences will mean that the claim is false, while other sensory experiences will mean that the claim is true. Claims which meet this criterion will be meaningful while claims which fail to meet this criterion will be meaningless. And yet some questions remain. What does Ayer mean when he speaks about accepting or rejecting a claim?

Ayer says first that for a claim to be meaningful there must be potential sensory experiences which can verify the claim and make us "accept the proposition as being true". Does Ayer then mean to say that for a claim to be meaningful there must be potential sensory experiences which conclusively determine that the claim is true?

Likewise Ayer says that for a claim to be meaningful there must be potential sensory experiences which can falsify the claim and make us "reject it as being false". Does Ayer then mean to say that for a claim to be meaningful there must be potential sensory experiences which conclusively determine that the claim is false?

The answer to both of the previous questions is a resounding no. According to Ayer a claim can neither be conclusively verified nor conclusively falsified. The kind of verification Ayer is talking isn't verification the strong sense of the term: "A proposition is said to be verifiable, in the strong sense of the term, if, and only if, its truth could be

conclusively established in experience”(37). Rather, Ayer means a weaker standard of verification:“ it is verifiable, in the weak sense, if it is possible for experience to render it probable’ ” (37). In short Ayer holds that a claim is meaningful not only in the case that it can be conclusively verified and conclusively falsified by potential sensory experience. Instead Ayer holds that a claim is meaningful if some potential sensory experiences render the claim more probable while other potential sensory experiences render the claim less probable.

More relevant though to our concerns though is why Ayer thinks a strict verificationist principle can't be the criterion for meaningfulness. The reason for that is simple. According to Ayer, no non-tautological claim can ever have its truth value conclusively determined by sensory experience: “the principle that a sentence can be factually significant only if it expresses what is conclusively verifiable is self-stultifying as a criterion of significance. For it leads to the conclusion that it is impossible to make a significant statement of fact at all.” (38) .

That is to say that no claim will ever meet the criterion of the strict verificationist principle. The strongest reason that no claim can ever meet the strict verificationist principle is because any sensory experiences which are inconsistent with a claim will “presuppose the existence of certain conditions”(38). Allow me to give an example of what I mean by that.

Suppose that some predictive formula for gravity predicts that such and such object will be in such and such place at a given time. Furthermore, suppose that the object isn't in that place at the specified time. At face value we would say that our formula was falsified. But we can't say that with certainty. The object we were predicting might have been an optical illusion. There might have been a massive invisible object pulling our object in way that is totally consistent with gravity. Our depth perception might be off. Our understanding of time might be off. So on and so forth. The point is that for any one of these possible cases our formula would not have been predicted. A similar line of reasoning can be applied to verification.

I think Ayer's point is correct. We can never conclusively determine the truth value of the claim, and that means trouble for our criterion.

We said that the criterion for a physical claim is that one can distinguish a world where the claim is true from a world where the claim is false. However, given that we are in agreement with Ayer, we can no longer distinguish a world where a claim is true from a world where that claim is false. After all, if we cannot conclusively determine the truth value of the claim then we also cannot perfectly distinguish a world where a claim is true from a world where a claim is false.

This would hold true, for cases like (3) and (4), where we have already determined that we can't perfectly distinguish worlds where they are true from worlds

where they are false. This would also hold true for cases like our super simple introductory claim, (1). For one thing you can't perfectly determine if the chair is in the specified location; I mean, what happens if you move the chair over 10 atoms, 100 atoms, 101 atoms. It's not clear when the chair would be at the specified location and when it wouldn't be at that location. It's not clear which exact atoms constitute the chair: there are some number of atoms such that their removal would stop the chair from being the chair, what that exact number is though, I sure don't know. So on and so forth. Even the case of the claim about a chair isn't as black and white as we've been assuming.

Where does that leave us? Well we've determined that we agree with Ayer. Strict verification is impossible, we must adopt some form of loose verification. In our terms, for a claim to be physical we need not be able to conclusively determine the truth value of the claim by looking at the world. In what way then must we be able to determine the truth value of physical claims by looking at the world, given that by "able to distinguish between possible world" we really mean something along the lines "able to distinguish between possible worlds in a way which may or may not be correct"? Given that we are not using a strong verificationist criterion, what does it even mean to distinguish a world where a claim is true from a world where a claim is false?

What we mean by "for a claim to be a physical claim we must be able to distinguish a world where the claim is true from a world where the claim is false" is

broadly, that some possible worlds must be consistent with the claim while some possible worlds must be inconsistent with the claim. Note, that by “consistent”, and “inconsistent” we don’t mean with finality; that is we don’t mean that for a claim to be a physical claim that there must be both a possible world that is conclusively inconsistent with the claim and where the claim must be false with finality, as well as some possible world that is conclusively consistent with the claim and where the claim must true with finality. Rather, we mean that there must be a possible world where the claim is inconsistent in the sense that the claim is improbable as well as a possible world that is consistent with the claim, that is to say a possible world where the claim is more probable.

Put as simply as possible using everyday language, for a claim to be a physical claim there must be some possible state of the world which would increase the likelihood that the claim is true as well as some possible state of affairs which would increase the likelihood that the claim is false. When we say that for a claim to be a physical claim we must be able to distinguish a world where the claim is true from a world where the claim is false, we mean that we must be able to distinguish a world where the claim is *likely* to be true from a world where the claim is *likely* to be false. This requires further discussion on what we mean by probability but that will come later.

Let’s recap where this leaves us. Initially we had a problem with cases like (3) and (4). In the cases of (3) and (4) we weren’t sure how exactly to distinguish a world

where the claim is true from a world where the claim is false. Given however, that (3) and (4) were both claims about physical things, and given that for a claim to be a physical claim we must be able to distinguish a world where the claim is true from a world where the claim is false, the status of (3) and (4) as physical claims was in question. We noted that in the case of (1), however, we could seemingly always distinguish a world where (1) was true from a world where (1) is false.

Then we looked at Ayer who claimed that we can never conclusively determine the truth value of claims about matter of fact. We generally agreed with his point: in the case of claims about physical things we can never conclusively determine the truth value of the claim. So then, when we spoke about distinguishing a world where a claim is true from a world where a claim is false in the case of physical claims we didn't mean that we must be able to do so with finality. We merely meant that we must be able to distinguish a world where the claim is likely to be true from a world where the claim is likely to be false.

Now that leaves us to return to the cases of (1), (3) and (4). Contrary to our assumptions, in the case of (1) we cannot conclusively determine the truth value of the claim by looking at the world. As we said the chair might be an optical illusion, the exact criterion of "chairness" is unknown, so on and so forth; the point that we can never conclusively determine the truth value of any physical claim¹², including (1). So then, in

¹² Which is not to say that we can conclusively determine the truth value of non-physical claims, but that is a separate discussion.

the case of (1), just as in the case of (3) and (4) we cannot conclusively determine the truth value: we cannot conclusively distinguish a world where (1) is true from a world where (1) is false.

That we cannot conclusively distinguish a world where claim is true from a world where a claim is false though doesn't mean that there are no physical claims. We need not be able to determine the truth value of physical claims with finality. When we speak about distinguishing a world where a claim is true from a world where a claim is false we mean distinguishing a world where the claim is likely to be true from a world where a claim is likely to be false: by "distinguish between possible worlds" we don't mean that one must be able to do so with perfect precision.

Now that's all well and good but the question remains how is our claim different than Ayer's claim, at least for non tautological claims?

To this point we have determined that, with the exception of tautological claims, we will group the set of possible claims the same way that Ayer would group those claims; sure we might put different names on our two subsets of claims but their membership will be identical.

Now, there would be nothing wrong if it were the case that we agreed with Ayer. However, that is not the case at all. To this point we have been steelmanning the view

that our distinction is similar to that of Ayer. As we shall soon see, in this next chapter, our claim is pretty different than that of Ayer in some pretty major ways .

Chapter 4: A closer look at our comparison with Ayer.

A summary of how our claim currently compares to Ayer's claim.

To this point we have attempted to focus on the similarities between our distinction and between Ayer's distinction. If I had to summarize the major similarities and differences between our distinction and Ayer's distinction we would have to focus on three points; two of the points are differences between our distinction and Ayer's distinction while one of those points is a similarity between the two distinctions.

The major similarity between Ayer's claim and our own is in the way both claims will divide up the set of potential claims. As we have pointed out many times, for the set of potential claims (excluding tautologies), our distinction and Ayer's distinction will divide up the claims into two identical sets.

That is, that the membership of our set of physical claims will be identical to the membership of Ayer's "meaningful claims". Likewise the membership of our set of non-physical claims will be identical to the membership of Ayer's "meaningless claims". Of course, as we have noted many times, there is one way in which the membership of our two sets will be different than the membership of Ayer's two sets: tautologies.

A quick side note: we will refer to this similarity, the one where Ayer divides up potential claims into sets identical to our own, as ' the "similarity of division" '13. Now back to the differences between Ayer's distinction and our own distinction.

While we categorize tautologies as non-physical claims, Ayer categorizes tautologies as meaningful, if trivially so, claims. It follows then that for tautological claims, that we will not divide up the claims the same way as Ayer. While we otherwise have said that the membership of the set of physical claims is identical to the membership of the set of non-physical claims, that is untrue in the case of tautological claims; tautological claims will be members of the set of Ayer's meaningful claims yet they will not be members of our set of physical claims.

So then, that we categorize tautological claims differently than Ayer categorizes them is the first noteworthy difference between our distinction and Ayer's distinctions. It is one way in which the similarity of division doesn't hold.

The second major difference between our distinction and Ayer's distinction is, as we have previously noted, Ayer categorizes all non-empirical claims as meaningless

¹³ This is done for simplicity's sake. We will reference this "similarity of division" a bunch and we don't want to say the full sentence "the similarity between the membership of the two sets produced by our distinction's division of claims and between the membership of the two sets produced by Ayer's distinction's division of claims", so for simplicity's sake we shorten things.

while we don't make any claims about the meaningfulness of physical claims¹⁴. Our thesis isn't about the kind of claims which are meaningful. We are totally neutral on the meaningfulness of non-physical claims. We, unlike Ayer, allow for the possibility that non-physical claims are meaningful although we make no commitments either way.

Now this is all well and good. As things stand now, our distinction is pretty similar to Ayer's. At the end of the day we mostly categorize claims the same way as Ayer. Sure, we categorize tautological claims differently than Ayer categorizes them. Sure, we allow for the possibility that non-physical claims are meaningful while Ayer does not. At the end of the day though the fact remains that, with one exception (tautological claims), Ayer's distinction categorizes claims the same way as our distinction...or does it? To this point we have been doing our best job to steelman the claim "Ayer's distinction is similar to our distinction": as much as we can we have attempted to present the case that our distinction is similar to Ayer's distinction. Now that changes; we will begin to distinguish Ayer's claim from our own.

A summary of how we will further distinguish our claim from Ayer's claim.

¹⁴ In other words, Ayer holds that non-physical claims are meaningless; while we on the other hand make no claims about the meaningfulness of non-physical claims.

In what follows we're going to have to go on a bit of a winding path, collecting a couple of important premises, until eventually we have all the pieces we need. Once we've organized our argument we will attack the major similarity between Ayer's claim and our own.

Let's summarize what I am about to argue. My argument isn't super complicated at all but it requires looking at a handful of moving pieces (premises and arguments) before any pay off. For that reason I'm giving this simple summary upfront; my hope is that this initial summary will clarify why we are hopping around arguments which seemingly don't distinguish our claim from Ayer's claim in any obvious way.

First let's put things super simply. Our criterion divides up claims on the basis of our ability to compare claims to the world. Ayer's criterion divides up claims on the basis of whether the claim can be used to better anticipate experience. To this point our criterion has divided up claims identically to Ayer's division of claims. However, that is only because Ayer assumes that the scientific method is the best available method to evaluate claims which anticipate experience. However, Ayer will concede that the best available method is subjective and changes from person to person. In fact Ayer will concede that there could in theory be some other method of evaluating empirical claims which he would believe is superior to the scientific method¹⁵.

¹⁵ Ayer says "We trust the methods of contemporary science because they have been successful in practice. If in the future we were to adopt different methods, then beliefs which are now rational might become irrational from the standpoint of these new methods."(100). This implies that we trust science because it has been the most successful method of producing empirical beliefs, but it is implied, were

Putting things together, given that Ayer holds that empirical claims are ones which allow us to anticipate experience, and given that he also holds that the best way to anticipate experience is partially subjective, it will then follow that for certain methods of evaluating anticipatory claims the division of possible claims will change: that is the membership of this set of “empirical claims”, or “claims which are used to improve our ability to anticipate experience”, the membership of the set of meaningful claims in short, will change from person to person.

The consequence of the previous vague blob of words will be that the similarity of division is contingent on a person believing that the scientific method (or something approximating the scientific method) is the best way to evaluate empirical claims. However one might believe that any number of methods are best for evaluating claims; for some of those methods though the similarity of division will not hold. Ok, that was a bunch. Let’s take a breath before explaining what we just said.

A closer look at the differences between Ayer’s claim and our own claim.

there another empirical claim evaluating method, Ayer would concede that that method is preferable to the scientific method.

First, we need to clarify what Ayer means by (non-tautological) “meaningful claims”, and “empirical claims”, and by all the many names Ayer uses to refer to the category of meaningful non-tautological claims. Ayer mostly defines these empirical claims as a claim which is more consistent with some potential sensory experience and less consistent with some other potential sensory experiences. However, Ayer also provides another necessary condition for membership as an “empirical claim”, that a claim “enable us to anticipate the course of our sensations” (97).

First we need to explain how we know that Ayer holds that a necessary condition for membership in “physical claim” is that a claim must in some way allow us to better anticipate the course of our experience: that this, let’s call it, “potential anticipatory utility” of a claim is a necessary condition for classification as an “empirical claim” isn’t obvious and requires a closer reading. It’s not especially complicated but it requires us to more closely examine Ayer piece by piece. What do we really mean by “empirical claim”?

Now you might say that we already know what Ayer means with terms like “empirical claims”. An empirical claim, you will remember, is any claim which is consistent with some possible sensory experience but inconsistent with other possible sensory experience: “Empirical propositions are one and all hypotheses, which may be confirmed or discredited in actual sense-experience”(93-94). You would be right in saying so; however, we will see that, at a deeper level, there is a more exact meaning to

“empirical claim”. Ayer’s “empirical claim” is also a claim which allows us “to anticipate the course of our sensations”(97). Let’s take a closer look.

Ayer says that “all propositions which have factual content are empirical hypotheses; and that the function of an empirical hypothesis is to provide a rule for the anticipation of experience. And this means that every empirical hypothesis must be relevant to some actual, or possible, experience, so that a statement which is not relevant to any experience is not an empirical hypothesis, and accordingly has no factual content” (41). An *essential feature* of an empirical hypothesis is that it provide a rule for the anticipation of experience.

In short Ayer holds that every empirical claim must in some way allow us to anticipate the course of our experience. That is, it a necessary condition for membership as an empirical claim that a claim in some way allow us to anticipate the course of our experience. Furthermore, not only is it a necessary condition for membership as an empirical claim that a claim allow us to anticipate the course of our experience, it is in fact the entire reason that we make empirical claims, “What is the purpose of formulating hypotheses? Why do we construct these systems in the first place?” The answer is that they are designed to enable us to anticipate the course of our sensations. The function of a system of hypotheses is to warn us beforehand what will be our experience in a certain field—to enable us to make accurate predictions”(97). Ayer continues , “The hypotheses may therefore be described as rules which govern our

expectation of future experience.”: that is when thinking about empirical hypothesis we need only to think about which experiences they predict.

The point is, that it is a necessary condition of an empirical hypothesis that it allow us to anticipate the course of our experiences. Now that we have established this anticipatory criterion we need to take a closer look at the rules governing it. That is Ayer provides rules for how to evaluate these hypothesis which govern the course of our experience. Once we have those we can get to our point; that given the anticipatory criterion, and given Ayer’s further rules for evaluating claims which allow us to anticipate experiences, it will then follow that under certain conditions the classification of a claim’s meaningfulness will change. But we’re getting ahead of ourselves. We still need to look at Ayer’s rules for evaluating these anticipatory claims.

So we’ve established that this anticipatory criterion is a necessary condition of empirical claims. Put differently, it is a necessary condition of an empirical claim that it, in some way, enable us to determine what the probability of certain potential sensory experiences will be.

But we require further explanation of this “probability”: “It is necessary now to make clear what is meant in this context by the term ‘probability’”(99).

Ayer continues by explaining the rules under which a claim can be used to predict the probability of certain future sensory experiences. When Ayer speaks about the probability of an empirical claim he means “Roughly speaking, all that we mean by saying that an observation increases the probability of a proposition is that it increases our confidence in the proposition, as measured by our willingness to rely on it in practice as a forecast of our sensations and to retain it in preference to other hypotheses in face of an unfavorable experience. And, similarly, to say of an observation that it diminishes the probability of a proposition is to say that it decreases our willingness to include the proposition in the system of accepted hypotheses which serve us as guides to the future.”(100). So then when Ayer says that it is a necessary condition of an empirical claim that it allow us to anticipate the course of experience, to figure out which sensory experience is probable, he means that in the sense that there must be sensory experience which makes us more or less willing to rely on the hypothesis to predict our future experience.

But there is more, Ayer uses his concept of “rationality” to further explain what he means by ‘probability’.

When Ayer speaks about the probability of a claim he means the probability in the sense that we are likely to rely on the claim in the future. Furthermore when Ayer

speaks about the probability of a claim he doesn't only mean how likely we are to rely on it to make predictions in the future. He means, how likely we are to *rationaly* rely on a claim to make predictions in the future. That is, by "probability" he means our likelihood to rely on the claim to make predictions *given* the method which we consider to be the most reliable method for predicting future experience. That is there are times when we don't use the best available methods to predict our future experience:

"Although we acknowledge that certain standards of evidence ought always to be observed in the formation of our beliefs, we do not always observe them. In other words, we are not always rational"(100). When Ayer says that the purpose of a claim is to better assess the probability of future experience he means the probability of the claim given that we use whichever hypothesis evaluating technique we consider to be most reliable: the "rational" probability.

So then, grounding the last paragraph in the text, Ayer explains what he means by the probability of a claim. He means the likelihood of the claim given our use of the method for evaluating empirical claims which we consider to be most reliable, that is to say "rational": "To say that an observation increases the probability of a hypothesis is not always equivalent to saying that it increases the degree of confidence with which we actually entertain the hypothesis, as measured by our readiness to act upon it: for we may be behaving irrationally. It is equivalent to saying that the observation increases the degree of confidence with which it is rational to entertain the hypothesis .And here we may repeat that the rationality of a belief is defined, not by reference to any absolute

standard, but by reference to part of our own actual practice.” (100-101), and by rational Ayer more clearly explains what he means “we define a rational belief as one which is arrived at by the methods which we now consider reliable”(100).

Now, this is where things get interesting. I will now give an example where, given what we have now said about rational probability, a claim will be meaningless while at the same time being a physical claim. That is, we will show, that under certain conditions there will be claims which we classify as physical claims but which Ayer will classify as a non-physical claim. Remember that this is just a thought experiment. We are certainly not claiming that the person in our thought experiment exists. We are simply using a thought experiment to create a case where Ayer will, in certain conditions, classify physical claims differently than we will classify them. The thought experiment is itself far fetched and silly but that doesn't matter.

Suppose a hypothetical person who believes that the most reliable method for predicting certain kinds of future experience is by the purity/impurity of their past thoughts. Let's just suppose that we define purity and impurity using the Torah. So pure thoughts are thoughts about Torah study and charity and things of that nature. Impure thoughts are for example thoughts about certain kinds of bugs, and thoughts about touching dead bodies. Let's suppose he believes that the purity of his thoughts best predicts all especially bad, and good things that happen to him.

Now let's consider a hypothetical case; a case with a hypothetical claim that needs our categorization. We will present the case and then consider how the claim will be categorized. That is, first we look at how we will categorize the claim, then how Ayer would categorize the claim given his belief that the scientific method is "rational", then we will consider how Ayer would categorize the claim from the perspective of the person in our thought experiment.

Let's suppose that hypothetical person is using a rope to pull up a piano from the sidewalk into their second-floor window. The rope breaks. The piano falls and breaks into a hundred pieces. This upsets our hypothetical person and makes him feel very bad, which of course makes him realize that it fell because he had an impure thought. He doesn't know which thought it was, but the piano wouldn't have fallen if there had been no thought so there must have been one. Down below on the street a person says "the piano fell because the knot on the rope came apart unfurled and untied". The claim in question then is "**the piano fell because the knot on the rope came apart unfurled and untied**". Now, let's consider how we will categorize the claim.

First we consider how we would categorize this claim. We would ask ourselves, would we be able to distinguish a world where this claim is true from a world where the claim is false? We would, of course, be able to distinguish a world where the claim is true from a world where the claim is false. We could, for example, observe the knot at the moments right before the piano fell. If for example a knot becomes undone then we

would be in a world where the claim is true. If however, we observed that the knot did not become undone but instead the rope ripped far away from the knot then we would say that we are not in a world where the claim is true. So then, because we can distinguish a world where the claim is true from a world where the claim is false, we will say that this claim is a physical claim¹⁶.

Now we move onto Ayer's criterion given his belief that the scientific method is the most reliable method (rationale) for assessing claims like the one above. We need only to ask ourselves whether there are sensory experiences that would make us more or less likely to rely on this claim to predict future experience. The answer is of course that there are such experiences. According to Ayer, who believes that the scientific method is the most reliable way to predict experience, there are of course experiences which can confirm or disconfirm the claim: we can at the very least observe the knot at the point that it came apart and thus potentially make better predictions about knots in the future. So then according to Ayer, given his subjective rationality, this claim will be a meaningful claim.

So far so good. We say that the claim is physical, Ayer would say that the claim is meaningful. Up until this point our criterion categorizes the claim identically to Ayer. This is where things diverge. Let's look at how Ayer would evaluate the claim relative to our exemplar person's most trusted method for anticipating experience.

¹⁶ Ignore any background assumption about causation, which we are certainly making. What matters here is that the claim can be tested against experience. This will be true even if we are making some faulty assumptions about the existence of causation.

Our hypothetical person would say that the piano must have fallen because they had an impure thought at some time in the past. He would say that the knot coming apart had nothing to do with it. He would be puzzled in fact that we think the falling piano had anything to do with the rope. That is, even if the knot had not come apart the piano would have phased through the rope, made corporeal like the ghosts in Harry Potter, and broken into one hundred pieces on the ground. That is the knot might have come apart, it might not have come apart, but those things have no causal relation to the piano falling down¹⁷.

Now would Ayer say that, relative to this person, the claim is meaningless? I think so. According to Ayer a claim is meaningful only the case that some potential sensory experience would confirm or disconfirm the truth value of the claim. However for the claim “the piano fell because the knot on the rope came apart unfurled and untied” there is no potential sensory experience which will confirm or disconfirm the truth value of the claim. To our hypothetical person, the truth value of the claim cannot be changed by some sensory experience. Because our hypothetical person believes that reference to his past thoughts is the most reliable method for anticipating future experience it follows that there is no sensory experience which can falsify/confirm the claim in question. So then according to Ayer, for this person, the claim in question would then be a meaningless claim.

¹⁷ Note that this is different than an auxiliary hypothesis; this standard under which the best way to predict bad things is our hypothetical person’s most trusted way of predicting the future.

What we have shown here is that our similarity with Ayer is contingent on the method a person believes to be most reliable for anticipating experience. In the case that a person who, like Ayer, believes that the best method for anticipating experience is the scientific method, then Ayer's criterion will indeed divide up non-tautological claims the same way as we divide up those claims. However, for a person who uses other methods to anticipate the course of their experience, such as our hypothetical person, Ayer's criterion will divide up non-tautological claims into sets which don't perfectly match the two sets produced by our criterion.

There is though one major question that remains. You might ask, how would we assess our hypothetical claim relative to our hypothetical person. Would we also not say that, relative to our hypothetical person, the claim is a non-physical claim? That is, couldn't the person say, given our criterion, that the hypothetical claim is not, given their personal beliefs, a physical claim? I'm not sure. Let's see.

The criterion for membership as a physical claim is that there must be something in the world that can distinguish a world where the claim is true from a world where the claim is false. Now our hypothetical person would say that we cannot distinguish a world where the claim was true from a world where the claim was false. After all, our hypothetical person believes, that in both worlds the piano was drawn down by past impure thoughts and that the rope had nothing to do with it; that is the piano would have

been drawn down even if it was suspended by steel cable; it would past through the cable. So then it would seem to follow that the hypothetical person could perhaps say that we cannot distinguish a world where the claim is true from a world where the claim is false. How could we; the movement of the piano after all has nothing to do with any knots?

Let us consider how we would analyze the claim. We would take the claim, “the piano fell because the knot on the rope came apart unfurled and untied” to mean that in a world where the knot came apart that the piano would fall but in the world where the knot did not come apart the piano would not have fallen. Now, of course, the hypothetical person could object and claim that the piano fell because of impure thoughts. That has no bearing on the claim though. The claim is simply that in a world where the knot was tied the piano would not have fallen but in a world where the knot was not tied that the piano would have fallen. Now our claim might be totally misguided but our claim is one which could in theory, if perhaps incorrectly, be tested against the world. That our hypothetical person objects and proposes an alternative hypothesis has no bearing on the claim itself. That, however, is *uniquely not the case with Ayer*. Prima facie anyone would take the hypothetical claim to be about a testable claim; even if they thought that the claim makes incorrect assumptions. Ayer though, with his further criterion of rationality, prevents an everyday understanding of the claim in the case of our hypothetical person. According to Ayer, we must evaluate empirical claims by our own subjective standard, not simply by the content of the claim itself. So while the claim

was about a testable, if perhaps incorrect from the perspective our hypothetical person, empirical claim, Ayer will insist that our hypothetical person evaluate the claim, not on the basis of what the claim itself says but by additionally using his own standards, standards which then jeopardize the claims status as an empirical claim.

So then we have demonstrated is that in certain cases Ayer's criterion will divide up the set of non-tautological claims differently than our criterion would divide up the set of claims. Sure our example was pretty far-fetched. We used a wacky example to test the conceptual edges of Ayer's case. Are there any less wacky cases which we could have used instead of our very wacky case? I would assume so but I'm not sure.

Initially, the largest similarity between Ayer's criterion and our criterion was that the two criteria will divide up a set of non-tautological claims into two identical sets. What have demonstrated in this chapter is that that similarity, what we have called the "similarity of division", is contingent on what one believes to be the most reliable method of anticipating experience. Given certain possible subjectively trusted methods the similarity of division will not hold and with that the main similarity between Ayer's criterion and our own will not hold.

Chapter 5: review and further discussion.

We began the first chapter by introducing a distinction between two types of claims: what we called “physical claims”, and “nonphysical claims”. That is claims such that there will be a difference in the world corresponding to their truth value, and claims such that there will not be a difference in the world corresponding to their truth value. We used claim (1), the claim about the chair, as an example of a physical claim. We used claim (2), the claim about the morality of killing puppies, as an example of a nonphysical claim.

Also we noted that there were two other kinds of physical claims. There were (3), claims about universal laws, such as claims about gravity. There was also (4), claims which clearly had some exceptions, such as claims about conditioning in mammals. We noted that for (3) and (4) that the conditions under which they are true was unclear and as such it was unclear whether they met the conditions for membership as physical claims. We put (3) and (4) aside. We would explain why (3) and (4) were physical claims at a later point.

In the second chapter we introduced a major problem for our claim. We seemed to be presuming the existence of a mind independent external world. To explain why we were not presuming the existence of a mind independent external world we introduced

Chalmers's paper where he argued that, should it turn out that our reality is simulated, it would not then follow that our reality isn't real. We objected to some of Chalmers's claims about simulations, namely we argued that "being in a simulation" was not a meaningful term. Eventually though we returned to our problem where we used Chalmers's argument to explain why our claim does not presume the existence of a mind independent external world. We resolved our problem.

In the third chapter we introduced Ayer. We noted a major similarity between Ayer's distinction between meaningful/meaningless claims and our distinction between physical/nonphysical claims, namely that we both divided up non-tautological claims into two identical sets. We called this similarity the "similarity of division".

We used the proverb about a tree falling in the forest when no one is around to hear, to see whether we would classify such a claim the same way as Ayer would classify it and in the end we determined that we would both classify the claim the same way.

We then introduced Ayer's loose verificationist principle, which we agreed with, and then used the loose verificationist principle to explain why (3) and (4) were in fact physical claims. Whereas we had initially thought that (3) and (4) were, unlike (1), unique in not having clear truth conditions, we realized that no physical claims have perfectly clear truth conditions. We added a nuance to our claim. When we speak about

distinguishing between a world where a claim is true and world where claim is false, we merely mean to distinguish between a world where the claim is likely to be true and between a world where the claim is likely to be false. In light of this added nuance our issue with (3) and (4) disappeared. Both (3) and (4) are physical claims because we can distinguish between worlds where the claims are more or less likely to be true.

In the fourth chapter we begin to distinguish our claim from Ayer's claim. We introduce Ayer's anticipatory criterion and with it Ayer's rational probability. Given then Ayer's usage of probability, we presented a hypothetical case where Ayer would categorize as a nontautological claim differently than we would categorize that claim. In doing so we demonstrated that the similarity of division does not hold.

Further discussion

There are some other ways that we might potentially distinguish Ayer's distinction from our own. For example how would we classify claims in written fiction? How would we classify a case where a fictional person says that "there is a chair over there"? Now you might think that we have done exactly that in all of our examples. In all of our examples, however, we were proposing hypothetical claims and acting as if a real person said it.

Let's go a step further. Suppose a claim about a certain chair's location in Wonderland. That is a claim about a chair in a place that only exists in fiction. Would Ayer say that the claim is meaningless because we cannot predict any sensory experience? Would Ayer say that the claim is meaningful because we can use the claim to predict the sensory experience of seeing other words in the book; that is we can predict which letters we will see later in the book on the basis of the chair's location. For example we might predict that we will see words like "he sat down on the chair". It's not clear to me how Ayer would classify such a claim. Ayer does hold that poetry is often meaningful but it's not clear if poetry is totally fictional. I can think of a couple of different kinds of claims that could be made in fiction but for which I have no idea how Ayer would categorize them¹⁸; I'm not even sure how we would categorize them.

¹⁸ For example suppose a fictional person who says that there is a chair at a certain real location, like New Jersey for instance. Perhaps that claim is meaningful because we can in fact check if there is a chair in the non-fictional New Jersey.

But suppose then a fictional character in a fictional world much like our own save that this fictional world has an 8th continent. Suppose that this fictional character says that there is a chair on the 8th continent. We might say that this claim is meaningless because there is no such continent in real life which would allow us to test the claim. We might say that the claim is meaningful because there could in theory be such a continent and we would then be able to test the claim. We might say that the claim is meaningful because the existence of the continent is itself a testable part of the claim; we can test the claim by first checking if there is an 8th continent. We might even say that the claim is meaningful even though we can't check the claim because the fictional character could check the claim in his fictional world. We might say that the claim is meaningful because we ourselves can in fact check if the claim is true. For example in the case that all the fictional characters agree that there is no such chair we might say that the claim is likely false. So then were we in fact able to test the claim?

We could also suppose a fictional world which is exactly like our own world save for the fact that we never exist. Would claims made in this world be meaningful? We might say that claims made in this world are meaningless because we can never test them. But Ayer does say that claims about past events are meaningful even though we seemingly can't test them. Can we test claims about the past though? Some certainly but not all of them. So on and so forth. The point is that this discussion of claims made in fiction introduces many questions.

Also our treatment of Chalmers was a bit superficial. We set the table for a discussion but didn't really get into it. We didn't discuss the implications of our claim that "being inside of a simulation" is a meaningless phrase; for instance does it follow then that aren't any simulations? I don't think so. At the same time how can we say that there can be simulations but that being inside of a simulation doesn't mean anything? I'm not sure.

Reflecting now on the project I might have done a couple of things differently. I think that the part of the paper where we use Chalmers to explain why we don't assume the existence of a mind independent external world was too sloppy; while I think that that solution was elegant I don't think that the reasoning was explicit enough.

Also I might have liked to connect Ayer to something called "predictive processing". The basic idea of predictive processing is to think of the human brain as a prediction machine. Ayer seems to suggest something very broadly along those lines, "The hypotheses may therefore be described as rules which govern our expectation of future experience. There is no need to say why we require such rules. It is plain that on our ability to make successful predictions depends the satisfaction of even our simplest desires, including the desire to survive."(97): Ayer seems to be suggesting that predictions are our essential tool for interacting with the world. Predictive processing, if it's claims were true, would seem to support Ayer's emphasis on the importance of a claims predictive content.

The problem is that predictive processing doesn't really predict much. It's "postdicts" a lot but is currently pretty light on falsifiable predictions. I'm not sure if it really should predict much though. Perhaps it's just a simple way to organize many varying scientific claims under a single framework with far fewer assumptions. While writing this paper predictive processing was one of the main things I was thinking about but I can't really tell if it's some huge insight or bad, that is to say totally unfalsifiable, science.

We might also have spent more time discussing Ayer's view of meaning. Ayer says pretty much that for a claim to be meaningful there must be some way to test it's truth value. We can determine the truth value of tautologies and empirical claims so they are meaningful. But why must we be able to determine a claims truth value for the claim to be meaningful? I have some vague intuitions on the answer to this question but I still feel like I might be missing something here.

Bibliography

Ayer, A. J. (1952). *Language, truth and logic*. New York, NY: Dover Publ.

Chalmers, David J. (2005). The matrix as metaphysics. In Schneider, S. (2016). *Science fiction and philosophy: From time travel to superintelligence*. Chichester, West Sussex, UK: John Wiley & Sons. pp. 35.

Yudkowsky, E. (2007, July 28). Making Beliefs Pay Rent (in Anticipated Experiences).

Retrieved from

<https://www.lesswrong.com/posts/a7n8GdKiAZRX86T5A/making-beliefs-pay-rent-in-anticipated-experiences>

