

ABSTRACT OBJECTS IN THE CAUSAL ORDER

By

Justin Wun-Man Ngai

A thesis

Submitted to the Victoria University of Wellington

In fulfilment of the requirements for the degree of

Master of Arts

In Philosophy

Victoria University of Wellington

2013

Acknowledgements

I would like to say a huge ‘thank you’ to my supervisor Stuart Brock for his advice and guidance. Without him, this thesis would not be anywhere close to what it is today.

Thank you to Victoria University for giving me this opportunity to write this thesis. Thanks for the office, having a private spot to write has been helpful. Thank you to Simon Keller, Ramon Das, Sondra Bacharach and all the members of the faculty for giving me the chance to tutor. Thank you to all my fellow students for the friendship and great conversations we’ve had, especially Dan Burkett for being the best office-mate anyone could hope for.

Thanks to all my friends for all the support I’ve received. Hopefully I’ll be more sociable now that my Masters’ thesis is submitted.

Last but certainly not least, I would also like to thank my family for all the invaluable support I’ve received. Thanks to my father for being a reliable source of wisdom and encouragement. Finally, I would like to thank my mother for being so incredibly helpful and understanding.

Abstract

Abstract entities have long been viewed as entities that lack causal powers; that is, they cannot be constitutive of causes or effects. This thesis aims to reject this claim and argue that abstract objects are indeed part of the causal order. I will call this thesis ‘AOCO’ for short. In the first chapter I argue that other philosophers have committed themselves to the claim that some abstract objects have been caused to come into existence. In the second chapter, I argue that the best solution to Benacerraf’s problem is to concede that abstract objects have a causal influence on what we believe. In the third chapter I examine and evaluate objections to AOCO.

Contents

Acknowledgements	ii
Abstract	iii
Contents	iv
Introduction	1
1 – Background and Motivation	4
1.1. Meta-Ontology	4
1.2. Definitions of ‘Abstract’	11
1.2.1. The Way of Example	12
1.2.2. The Way of Conflation	13
1.2.3. The Way of Abstraction	14
1.2.4. The Way of Ostension	18
1.2.5. The Way of Negation	20
1.3. Abstract Objects in the Causal Sphere	22
1.3.1. Fictional Entities	22
1.3.2. Linguistic Types	24
1.3.3. Sets	26
1.3.4. Tropes	28
1.4. Opening the Floodgates	30
2 - Epistemology	31
2.1 The Benacerraf Argument	31
2.2 The Spurious Conclusion of the Benacerraf Argument	32
2.3 Deny the First Premise	34
2.3.1.a. The Causal Theory of Knowledge and Gettier Problems	35
2.3.1.b. Reliabilism - Truth Tracking Theory	38
2.3.1.c. Reliabilism - Armstrong’s Theory	40
2.3.1.d. Explanatory Theory	42
2.3.1.e. Internalism - Evidentialism	44
2.3.1.f. Contextualism	47
2.3.2. Steiner’s Objection	55
2.3.3 Objections with Positive Solutions	57
2.3.3.a. Wright’s Solution	58

2.3.3.b. Full-Blooded and Principled Platonism.....	59
2.4 Deny the Validity of Benacerraf's Argument.....	64
2.5. Deny the second premise.....	66
Chapter 3 - Metaphysics	70
3.1. Causal Theories	70
3.1.1. Regularity Theory	71
3.1.2. INUS Condition.....	74
3.1.3. Energy Transfer View	76
3.1.4. Counterfactual Theory	77
3.2. Azzouni's Objections	83
3.3. Arguments from Make-No-Difference	86
3.4 Counterfactuals with Necessary Abstract Entities	89
3.5. Nominalist and Traditional Platonist Objections.....	90
Conclusion	91
Final Conclusion.....	93
Bibliography.....	94

Introduction

Philosophers often take it for granted that abstract objects are not part of the causal order. To illustrate this, here is a small sample of philosophers making this claim:

“The negative way: abstract entities have no spatiotemporal location; they do not enter into causal interaction; they are never indiscernible from one another” (Lewis, 1986a, p.83).

“According to Platonists, mathematical objects are abstract: in other words, Platonists think of mathematical objects as neither causally active nor spatially located” (Liggins, 2006, p.135).

“There are reasons for supposing that if [abstract] objects exist then they lack causal powers and causal properties, or at least the causal power to influence human beings. We usually suppose that in order to exert a causal influence, an object must do so at some particular time and place, and this would not be possible for an object lacking a spatio-temporal location” (Cheyne, 2001, p.2).

“[It] is commonly assumed that such abstractions exist necessarily and that they cannot stand in causal relations” (Deutsch, 1991, p.210).

“[Abstract] objects, Platonistically construed, are neither causally active nor causally acted upon” (Drivers & Miller, 1999, p.308).

“Traditional Platonism is the realist ontology that recognises abstract objects, i.e., objects that are nonspatiotemporal and outside the causal order” (Linsky & Zalta, 1995, p.525).

“Crudely, we cannot, it appears, come into any kind of direct sensory contact with such ‘objects’; we cannot point at them, or spill our coffee on

them... their abstractness just consists, it appears, in a kind of causal impotence” (Wright, 1983, p.1).

“[Since] Platonists maintain that mathematical objects exist outside of spacetime, they endorse what we may call the principle of causal isolation (PCI), which says that there are no causal interactions between mathematical and physical objects” (Balaguer, 1998, p.110).

My goal in this thesis is to reject this assumption; that is, I claim that some abstract objects are in fact part of the causal order. I will call this thesis the ‘abstract objects in the causal order’ thesis or ‘AOCO’ for short. There have been other philosophers who have already postulated efficacious abstract objects (though in some cases, these philosophers do not want to call these entities abstract). Penelope Maddy for example, argues that sets are part the causal order,¹ whereas Linda Wetzel argues for efficacious linguistic types.² AOCO on the other hand goes beyond individual entities; it claims that we should reject the general claim that abstract objects are acausal. While some philosophers call these entities are concrete,³ I argue that we can still classify them as abstract. Furthermore, most philosophers take the term ‘Platonist’ to be a synonym of ‘abstract’. I will not use ‘Platonist’ to describe causal abstract objects because the term is metaphysically loaded. Since Plato describes his abstract objects as being in a Platonic heaven outside of space/time⁴ and also acausal, other versions of ‘Platonism’ seem to assume similar attributes. The entities I am interested in are different and so I will simply refer to them as ‘abstract objects’.

In Chapter 1, I will go through some definitions of what an abstract object is, as well as motivating my thesis by showing that philosophers have already postulated efficacious abstract entities. These entities do this by being part of causal effects. In chapter 2 I show how this argument can solve Benacerraf’s epistemic objection to Platonism – a problem that realists have historically struggled with. I will also argue that abstract objects should be part of causal effects, largely because of this. Finally, in chapter 3, I examine potential counterarguments to AOCO. Abstract objects are

¹ Maddy, 1992.

² Wetzel, 2009.

³ Such as Maddy (1992) and Bigelow (1988).

⁴ Plato, ND/2005, p.604-13.

part of the causal order. The traditional Platonist definition is outdated and should be replaced.

1 – Background and Motivation

There are two goals in this chapter, the first is to clarify what I mean by ‘abstract objects’, ‘existence’ etc. This will be dealt with in 1.1 and 1.2. After that, I will motivate AOCO by showing that different philosophers are individually committed to abstract objects being within the causal order. They do so by being part of causes, despite being abstract objects.

1.1. Meta-Ontology

Since we are having a discussion based on ontological issues, we need to ask what it means for something to exist. Most philosophers take existence and being as the same thing. For there to be an apple means that an apple exists. The two are logically equivalent and hence their meanings are the same. Philosophers such as Peter Van Inwagen have also argued that existence is univocal, meaning that the concept of existing is the same, regardless of the object that exists.⁵ He argues this point by drawing on the links between the concepts of numbers and existence. We can apply the same numbers to count very different types of objects, so existence is univocal.⁶ If there have been four major storms in our city, and four people in my family, I can say that the number of storms is the same as the number of people in my family. The concept of numbers is closely related to the concept of existence. To say that there are zero unicorns is to say that unicorns don’t exist. If I own exactly one pair of sneakers, then I have one more pair of sneakers than the number of unicorns in this world. This may sound odd but we can all agree that it is meaningful. Between the univocal nature of numbers and its link to existence, Van Inwagen believes we have good reasons to believe that existence is univocal. One may argue against Van Inwagen by pointing to all the different properties amongst different objects. My table and the American congress both exist, yet they have very different properties. Surely this shows that

⁵ Van Inwagen, 1998, p.236.

⁶ Van Inwagen, 1998, p.236.

things can exist in different ways. Quine has a response to this, he argues that these differences are external to the concept of existence itself.⁷ Imagine if I pointed to a flower and a car, then telling my friend that they are both red. My friend looks at me with a slightly puzzled look on his face. “Yes, I can see why you may think that they are both red, but surely they are not red in the same way! The flower is red because it is naturally that colour, whereas the car was just painted red”. My friend would be right in one sense but would nonetheless still be saying something odd. I can see that the reason that the flower is red is different from the reason that the car is, but that does not mean that they don’t share a colour. They are both red regardless of the differences my friend listed. In a similar light, different objects can have their existence exemplified in different ways, yet these are not strictly different ways of existing.

Famously, Quine has argued that the objects we claim to exist can be spotted through the use of an existential quantifier – ‘ \exists ’.⁸ ‘ $\exists x$ ’ means that there is one or more x , where ‘ x ’ is any object. Since being and existence are logically equivalent, ‘ $\exists x$ ’ means that x exists. Quine thus has a method for us to spot entities that we are ontologically committed to. First, we need to identify statements and theories that we believe to be true; we can then see what these theories are ontologically committed to. A theory is ontologically committed to Fs iff the theory entails or implies that there are Fs. Furthermore, as individuals, one is ontologically committed to Fs iff one believes a theory that is ontologically committed to Fs. What does it mean for a theory to entail or imply that there are Fs? Different philosophers say different things about this entailment or implication relation, for Quine, we investigate what a theory implies there is by a) putting the theory in the language of first order logic, b) listing all the theorems of the theory, and then c) looking for all statements that begin with the form “F exists”.⁹

For example, say one of the theorems is that 2 is an even number. In first order logic, that may be translated as “ $\exists x(Tx \& Ex)$ ” where “ Tx ” means that x is 2 and “ Ex ” means that x is an even number. As a statement, this can now be read as “ x exists, x is 2 and x

⁷ Quine, 1948, p.23.

⁸ Quine, 1960, p.242.

⁹ Quine, 1960, p.238-43.

is even". According to Quine, we are currently ontologically committed to the number 2 as an existing object. Given this commitment, we may or may not wish to continue postulating that these things exist. From here we have three options:

- i) Accept these entities as part of our ontology.
- ii) Stop committing to the theories that commit us to their existence.
- iii) Paraphrase these ontological commitments away.¹⁰

In other words we can either choose to accept these objects, change our theories and language so we don't refer to objects we don't believe in, or deny their existence. The paraphrasing idea can be explained like this: when we make statements such as 'there is a number between 2 and 4', if we don't want to say that the number 3 actually exists as an abstract object, we may choose to paraphrase this statement away. What we *really* mean when we say that 'there is a number between 2 and 4' is something like 'if you have more than two, but less than four whole things, you have three of them'. The new sentence is supposed to be logically equivalent to the original. We are now no longer ontologically committed to the existence of the number 3, but we're still expressing the same thing. Philosophers who believe in the existence of an object are called *realists*, whereas ones who denounce the existence of an object are called *anti-realists*. *Nominalists* are philosophers who don't believe in the existence of any abstract objects.

Though the Quinean method is well accepted in general, it is not without objections. William Alston for example objects to the paraphrasing option. His argument is as follows, the anti-realists according to Quine are supposed to paraphrase away ontologically problematic statements with logically equivalent ones that do not postulate these entities. However, how can two statements, one postulating the existence of an object, another denying it, mean the same thing? This seems to be a natural contradiction, either the two propositions say different things, in which they contain different ontological commitments, or they say the same thing, in which case they have the same ontological commitment and so the Quinean paraphrase has

¹⁰ Quine, 1960, p.129-34.

failed.¹¹ A reasonable anti-realist paraphrase is hence impossible, and not a true option for Quine's quandary.

Furthermore, Alston says that there is a fundamental problem with this process of admitting the existence of objects.¹² When I make propositions such as 'there is a good chance that he will come', am I really making any sort of ontological commitment? In everyday language, whether or not I am ontologically committed to these entities does not even come to mind. We certainly should not deem anyone making such an utterance an abstract realist by default. The only thing the paraphrase is good for is that it allows one to escape the *urge* to ask metaphysical questions that come with ontological commitments. That is, now that I am aware of my ontological commitments, I can replace the earlier proposition with say, "he will probably come". Now I will not have to urge to ask metaphysical questions such as if possibilities exist within space/time or if they are causal.¹³

To Alston, questions of what exists are 'less than parlour games'.¹⁴ Therefore, he is not only objecting to Quine, he is denying the usefulness of ontological debates in general. In response, all I can say is that Alston's interest simply differs from other philosophers, including myself, who partake in this debate. Given that we want to continue this discussion, perhaps there is still something to be said about the problem of paraphrasing. The solution is simple, rather than claiming that paraphrases retain the meaning of statements completely, we simply claim that it retains all meaning of the original, barring the objects we are ontologically committed to. This solution is so simple, it almost seems shallow or lazy. However there is no real reason why it should not be employed. After all, we know that this is what the Quinean paraphrases are supposed to give us. The point was never to retain all meanings of the proposition, including the ontology, which would make the exercise pointless. In response to Alston's second objection, I argue that just because we are not thinking of ontological issues when we make these abstract statements, that does not mean that we are not seriously ontologically committed to them. Or perhaps we're not, but that's something

¹¹ Alston, 1958, p.10.

¹² Alston, 1958, p.11.

¹³ Alston, 1958, p.17.

¹⁴ Alston, 1958, p.11.

we can sort out through Quine's methodology. Whether or not we are really ontologically committed to these entities due to our language is almost irrelevant.

In his 1980 paper *Ontological Commitment and Paraphrase*, Frank Jackson gives Quine another puzzle. Take the proposition "S is ontologically committed to K's" where S is a statement and K's are objects whose existences are in question. To put this proposition in a logical form would inevitably involve the existential quantifier '∃' followed by 'K'. It is possible for a nominalist to state that proposition while describing another philosopher being committed to abstract object K, but note that in saying that "S is ontologically committed to K's", he/she is automatically committed to the existence of K's, even before any debates over K occurred and despite nominalists believing in no abstract objects at all.¹⁵ Jackson then suggests that we should not look for ontological commitment through pure semantics and existential quantifiers, instead we should look at "one's preparedness to express one's sentences in terms of the semantic relation of *being true of* or *application*".¹⁶ For example, I am not ontologically committed to the existence of Harry Potter when I say 'Harry Potter is a fictional wizard', I am only committed if I agree that there is something that 'is a fictional wizard' applies to. An anti-realist of fictional characters for example would say no, she does not believe that this applies to any subject so she is not ontologically committed to Harry Potter. Jackson's objection to the Quinean method is very reasonable, since ontological commitments are about having references in the world, rather than purely linguistic notations. However, it is not strong enough to reject Quine's methodology. The existential quantifiers are only supposed to give us a starting point ontologically. Once again, we need to first notice that we are linguistically committed to the existence of Harry Potter and only then should we ask "OK, but are we *really* committed to this". The use of existential quantifiers really does not restrict us in the way Jackson thinks it does.

Quine's three options are not the only solutions; there is a fourth option for us that has been developed after Quine. This is the fictionalist option. Fictionalism can be used as a brand of anti-realism that does not completely undermine the usefulness of realism. For example, in his 2001 paper *Go Figure: A Path through Fictionalism*, Stephen

¹⁵ Jackson, 1980, p.305.

¹⁶ Jackson, 1980, p.310.

Yablo claims that we can make statements and theories that seemingly refer to abstract objects while remaining anti-realists. We do this by treating these statements or theories as hypotheticals or make-believe. For example,

‘Sherlock Holmes is a great detective’

should be read according to the fictionalist as:

‘Based on the Doyle novels, Holmes is a great detective’.

We can hence talk about Sherlock Holmes as though we are realists without really postulating the existence of fictional characters. In other words, we can make statements in the form of ‘ $\exists xFx$ ’ without truly asserting ‘ $\exists x$ ’.

There are many forms of fictionalism, each with its own differences but with the same spirit. A general division we can make between forms of fictionalism are the hermeneutic and revolutionary theories.¹⁷ According to the hermeneutic fictionalists, claims such as ‘Holmes is a great detective’ are *already* make-belief statements that are not really committed to the existence of abstract objects. Contrast this with revolutionary fictionalism; according to which speakers are truly committed to the existence of Sherlock Holmes for example. What makes them fictionalists is the fact that they think we should change this and treat these assertions as make-believe.¹⁸ For an example of hermeneutic fictionalism, see Gideon Rosen’s *Modal Fictionalism* (1990). Hartry Field’s *Science without Numbers* (1980) is no doubt the most well known example of revolutionary fictionalism.

Regardless of the fictionalist’s hermeneutic or revolutionary bends, he/she would still have a specific strategy when it comes to their form of fictionalism. Yablo gives us many options. I will give a quick survey of these strategies:

1. Instrumentalist fictionalism

¹⁷ These terms were originally dubbed by Burgess, 1983, p.96.

¹⁸ Refer to Burgess (1983) or the Stanford Encyclopaedia entry for more reading.

According to the instrumentalist, the statement quantifying over dubious objects should be read as make-believe. What is missing from this account however is how this statement *should* be read.¹⁹ This creates problems for the fictionalist, since it seems reasonable to ask them what the sentence *really* means, given that it is make-believe.²⁰

2. Meta-fictionalism

Let '*s*' be any statement asserting an abstract object. According to meta-fictionalism, to assert *s* is really just to say that *s* would be the right thing to assert given the make-believe world involved. For example, to assert that 'Sherlock Holmes is a great detective' is really just to assert that 'Sherlock Holmes is a great detective according to Doyle's novels'.²¹

3. Object fictionalism

According to this form of fictionalism, the statements referring to dubious objects are really just referencing non-controversial ones. The statement 'the number of cows is 3' is not really ontologically committed to the number 3, it is just another way of saying 'there are three cows'.

4. Reflexive fictionalism

Object fictionalism assumes that we introduce something just to talk about concrete objects, but that's only half the story. Sometimes we are just talking about concreta, but at other times we are actually talking about the abstract objects. Take the nominalist statement 'numbers do not exist'. Surely, it would be a wrong for a fictionalist about numbers to believe that this sentence is referencing concrete objects. Reflexive fictionalism states that one can switch between talking about abstract objects sincerely and fictionally.²²

¹⁹ Yablo, 2001, p.74.

²⁰ Yablo, 2001, p.74-5. Burgess, 1983, p.1.

²¹ Yablo, 2001, p.75-6.

²² Yablo, 2001, p.80-2.

5. Relative reflexive fictionalism

Yablo believes that reflexive fictionalism does not really benefit anyone. The nominalists would rather not talk about abstract objects at all, because according to them there aren't any. The Platonists on the other hand don't need fictionalism to talk about abstract objects. They believe that these entities exist, so why not just talk about them directly? Rather than slipping between concrete and abstract entities, Yablo suggests that we allow fictionalists to slip between engaging and disengaging the fiction. This way, the nominalist can escape talking about numbers for example when talking about philosophy, but still engage in the number-fiction when doing mathematics.²³ This is the crux of relative reflexive fictionalism.

Fictionalism can function as a powerful brand of anti-realism because it can maintain a lot of the explanatory power that realism provides without the ontological baggage. However it is important to note that one can be a fictionalist without being an anti-realist about the object in question. Just because we are engaging with make-belief x 's does not mean that x 's do not exist. Rather, it just means that we are agnostic about x 's. This gives fictionalism additional value. It is a very useful and versatile solution to Quine's quandary, one that allows us to escape from the ontological commitments of our accepted theories.

With this mind, I will now show how we are committed to certain abstract objects according to the Quinean method, and how these abstract objects are in fact part of the causal order.

1.2. Definitions of 'Abstract'

There are many ways of defining what exactly an abstract object is. David Lewis provides us with many ways of doing so.²⁴ We will discuss what he calls the way of

²³ Yablo, 2001, p.82-3.

²⁴ Lewis, 1986, p.81-6.

example, the way of negation, the way of abstraction and the way of conflation. Furthermore, we will look into Michael Dummett's way of ostension.²⁵ These methods are just supposed to be different options for us, depending on the purposes of the debate. Different entities will be categorised as abstract depending on our mode of definition. The term 'abstract' is ambiguous, but nonetheless there are examples of causal abstract objects regardless of which definition we adopt. We therefore have reasons to adopt AOCO – the thesis that abstract objects are part of the causal order.

1.2.1. The Way of Example

“First, the Way of Example: Concrete entities are things like donkeys and puddles and protons and stars, whereas abstract entities are things like numbers” (Lewis, 1986a, p.82).

According to the way of example, one doesn't even try to give an analytical explanation to the abstract/concrete distinction. Instead, one goes straight to the paradigm examples eg. 'donkeys are concrete, numbers are abstract etc.'. What justifies this lack of analysis is that there may not be any one concept that links all abstract objects. The advantage of this way is that what is abstract will be precisely what we paradigmatically call abstract. With the way of example, there will never be entities that we standardly believe to be abstract objects that get miscategorised. This cannot be said for all the other definitions, which we will discuss later. However, the downside to the way of example is precisely that it offers no explanation as to *why* something is identified as being abstract. For example, in this debate, if there is nothing that unifies abstract objects, it would be hard to find a reason why these entities cannot be part of the causal order. We may not want to call the way of example a definition at all, given that it does not try to explain the term and merely gives examples. Some paradigm cases of abstract objects are: properties, sets, propositions, numbers, and universals. Many of these paradigm cases have already been postulated as part of the causal order and I will explain that in depth later.

²⁵ Dummett, 1973, p.471-511.

1.2.2. The Way of Conflation

“Second, the Way of Conflation: the distinction between concrete and abstract entities is just the distinction between individuals and sets, or between particulars and universals, or perhaps between particular individuals and everything else” (Lewis, 1986a, p.83).

‘Conflation’ is defined as two objects or concepts merging until they seem to be one combined entity. The way of conflation involves the idea that all abstract entities can be ultimately reduced to sets or universals. It certainly seems unlikely in this day and age that philosophers would be mistaking one abstract object for another,²⁶ but there may be a more charitable interpretation to the way of conflation. Burgess and Rosen have an elegant explanation.²⁷ A philosopher may believe in the existence of a few abstract objects, yet still want to minimise her ontology as much as possible. Therefore she tries to reduce some abstract objects to others. For example, John Bigelow has argued that numbers are universals.²⁸ One may also be tempted to reduce numbers to sets. The way of conflation takes this idea and runs with it. Abstract objects are entities that can be reduced to universals or sets (or perhaps just cannot be reduced to particular individuals), whereas concrete objects can be reduced to particulars or individuals.

It is doubtful that all abstract objects can be reduced to sets or universals. Tropes for example are different from universals because the former is unique and the latter is not.²⁹ If anything, universals should be reducible to tropes, for tropes are finer grained than the general universals. It would also be difficult to reduce tropes to sets. Let’s apply this to an example and attempt to reduce the colour of my car (construed as a

²⁶ Rosen, 2012.

²⁷ Burgess & Rosen, 1997, p.19.

²⁸ Bigelow, 1988.

²⁹ Williams, 1953a p.4-6. Tropes are not paradigmatically abstract because they exist within space/time. However, some tropes do qualify as abstract under the ways of ostension and abstraction. This will be discussed further in 1.3.4.

trope of course) to a set. The set can only have one member, because tropes are specific. What would this member be? The obvious suggestion would be my car. However, my car has many different properties and this creates problems for us. Let's now try to reduce the speed of my car to a set. Once again, since we are talking about tropes, the set must only have one member. What is the one member of this set? My car? This is problematic. What is the difference between my car's speed and its colour? According to a set-theoretic view of tropes, these two properties are exactly the same! Both of them are one-membered sets, that one member being my car. Yet we would not want to accept this conclusion, surely the speed and the colour of my car are two separate properties. It should be obvious then that we cannot reduce tropes to sets. If tropes are abstract, this is problematic. Of course, one may simply choose to classify tropes as concrete and retain the way of conflation.

Burgess and Rosen have an additional objection to this way. It is one thing for all abstract entities to be reducible to universals or sets, it is quite another claim to say that they are abstract *because* they can be reduced in this way.³⁰ Why would being a universal or set in itself make something abstract? Furthermore, why are sets and universals abstract in the first place? Simply because they are universals or sets? That is circular and gives us no better explanation than the way of example.

The way of conflation is not a very popular one, but is still worth considering. I will argue later that sets are entities capable of causal interaction and thus, there are no problems with the way of conflation in conjunction with AOCO.

1.2.3. The Way of Abstraction

“The Way of Abstraction: abstract entities are abstractions from concrete entities. They result from somehow subtracting specificity, so that an incomplete description of the original concrete entity would be a complete description of the abstraction” (Lewis, 1986a, p.85).

³⁰ Burgess & Rosen, 1997, p.19.

This way of distinguishing the abstract from the concrete references how philosophers have historically defined abstract entities. Burgess and Rosen argue that this would be a mistake. Discussions about abstraction historically were about mental instantiations of abstract objects, or the mental process of abstraction.³¹ However, both the mental representations and the process of abstraction should not be confused with the objects themselves.

One way to make sense of the way of abstraction is to view it as the *products* of abstraction, where abstraction is the process of seeing two white things for example, and mentally creating the object ‘whiteness’. This view has some intuitive value, unfortunately it also depends on an outdated theory of mind.³² Alternatively, Burgess and Rosen suggest that the way of abstraction could be read as follows:

“not that abstracta *are* the products of a mental process of selective inattention, but that they are the kinds of objects that psychologically inclined philosophers of earlier times erroneously *took to be* such products” (p.19, 1997, Burgess & Rosen).

However, this is not much better. Why would the abstract or concrete status of objects be subject to the mistakes of past philosophers? That tells us more about errors in past ways of thinking than it does about the nature of entities.

That does not mean that the way of abstraction is doomed. Crispin Wright, Harold Noonan and Bob Hale take this idea and tweak it.³³ According to them, abstract entities are the referents of abstract ideas. Their models involve the concept of equivalence relations.

Take any property that we use in everyday language, such as ‘the *colour* of the table’. Let ‘*a*’ and ‘*b*’ be objects, let ‘*F*’ be the property in question, let ‘*Fa*’ be that property

³¹ Burgess & Rosen, 1997, p.17-8.

³² Wright, 1983, p.27. Rosen, 2012.

³³ Hale (1987), Wright (1983), Noonan (1978). I find Hale’s version the most well thought out and persuasive, so our discussion will be focused around his formulation.

in a , let ' Fb ' be that property in b and finally let and ' Rab ' be an equivalence relation between a and b that explains F . We can analyse properties in the following way:

$Fa = Fb$ if and only if Rab .

For example,

The direction of $a =$ the direction of b iff a is parallel to b .

We are hence giving a definition of the property through a relation. This can be done for any property that holds between different objects. Note that not all properties are abstract, some can be concrete. For example, the property of being a father can be formulated as

The father of $a =$ the father of b iff either a and b are siblings or a and b are step-siblings and the mother of $a \neq$ the mother of b .

The difference between abstract and concrete objects is a difference in the equivalence relation. More specifically, if there can be two instantiations of a property in two different places at the same time, then the property is abstract. Hale uses the example of types and tokens.³⁴ The first and the fourth word of this sentence are of the same type, but not the same token. From this, we can see that the two tokens of the word are not the same, but separate words. However, both words are of the same type and therefore, the same type can be exemplified at different places at the same time. Hale says that the way of abstraction is thus a cousin of the spatiotemporal distinction. He adds that this equivalence relation must be a grounding relation for the object, that is,

"R grounds F iff, for any statement of identity linking F -denoting terms, there is some statement to the effect that R holds among certain things, the

³⁴ Hale, 1987, p.56.

truth of which is (logically) necessary and sufficient for the truth of that statement of *F*-identity” (Hale, 1987, p.59).

For example, with the case of roundness again, we can explain the concept through two round objects and an equivalence relation – being the same shape. This relation can hold among different objects at different points in time, so roundness is abstract. Note that the spatial condition is essential.³⁵ Hale illustrates this point by showing how a concrete object can be misclassified as abstract. Imagine that we wanted to find out if fathers are abstract. According to the previous formulation, the grounding relation is siblinghood, which can hold between different people. This would misclassify fathers as abstract, according to Hale. However, if we analyse fathers for example with ‘*x* begat *a* and *y* begat *b* and *x* is the same human being as *y*’, where ‘*x*’ and ‘*y*’ are also people, then we can define fathers without having the one same equivalence relation between many different objects. It would hence fail to meet the spatial separation criterion. Since we need this criterion to be met for all of its grounding relations in order to call it abstract, it is correctly typed as concrete.³⁶ Another problem that Hale brings up is that sometimes, abstract objects have grounding relations bound over things that have no spatial location, such as concepts or other abstract objects. He simply fixes this problem by adding a further condition, somewhat forcefully into his thesis. He ends up with two different ways of expressing the criteria.

“*F* is an abstract sortal iff, for any *R* that grounds *F*, either
(i) *R* cannot hold between spatially located items at all,
or (ii) *R* can hold between things which are spatially, but not temporally,
separated” (Hale, 1987, p.61).

The condition added to deal with non-spatial objects is obviously condition (i). Take the concept of being a prime number. Whatever the equivalence relation is for this concept, it is not one that can hold among spatially located objects. It can only hold between numbers, which do not have spatial properties.

³⁵ Hale, 1987, p.59.

³⁶ Hale, 1987, p.58-60

There are consequences to defining the abstract through equivalence relations - not all objects that we may want to call abstract can be reduced to them. For example, how would we define fictional characters using equivalence relations? What objects would this equivalence relation hold between? The answer certainly is not clear, nor does it seem to be an intuitive model. Therefore, under this method of defining the abstract, fictional characters should be defined as concrete, which may not be a consequence we would want to accept. Regardless of how we want to formalise this definition or whether or not we agree with it, this method does not clash with efficacious abstract objects. Much like the way of conflation and abstraction, types are considered abstract here. I will make a case for types being part of the causal order and given that they are, abstract objects are part of the causal order under this definition.

1.2.4. The Way of Ostension

Michael Dummett gives an interesting method in which we may make the abstract/concrete distinction. His suggestion is that concrete objects can be pointed out ostensively, but this is not the case for abstract entities. For example we can point to a tree and make a reference to it, as we can for tables, chairs, people, shadows and an infinitely many number of concrete objects. Abstract objects however cannot be pointed out ostensively; instead we have to rely on its functional basis to 'point them out'. For example, we may describe a square as 'the shape of a window' or the number zero as 'the number of snakes in the New Zealand'. There needs to be a functional expression written in the form of 'the x of a '. Where ' x ' and ' a ' are objects, with ' x ' being the potentially abstract object in question.

x can be an object of ostension $\rightarrow x$ is concrete.

x cannot be an object of ostension & x can be picked out through functional expression $\rightarrow x$ is abstract.³⁷

³⁷ Dummett, 1973, p.481.

Dummett realises that not all concrete objects can be objects of ostension if we take ostension to literally involve a pointing gesture. It would be just about impossible to point out a colourless, odourless gas for example, yet it remains concrete.³⁸ This is why the first condition cannot be put as a biconditional. Furthermore, Dummett does not want colourless gasses for example to be misclassified as abstract. This is why the second condition is needed. We can see how certain abstract entities may not satisfy the antecedent of this second condition. It is unlikely that fictional character for example could be picked out through a functional expression. Therefore the second condition cannot be put as a biconditional either.

There are hence certain abstract objects that cannot be classified as concrete using the way of ostension. Hale comments on this part of Dummett's theory and argues that it is not a big problem. At the very least it sets up clear barriers between the abstract and concrete because being able to be expressed as a functional expression may necessitate that it cannot be an object of ostension.³⁹

David Lewis objects to the way of ostension in a footnote of *On the Plurality of Worlds*. He argues that even if this way can help us distinguish between the abstract and the concrete, it does not tell us anything about the nature of abstract and concrete entities. "It is like saying that snakes are the animals that we instinctively most fear – maybe so, but it tells us nothing about the nature of snakes" (Lewis, 1986a, p.82). This may be true, but it is still worth considering. For one, Lewis' way of example does even less in terms of telling us about the nature of the objects in question, yet it is still mentioned when philosophers ponder how to distinguish between the abstract and the concrete.

A more pressing problem is the fact that this analysis cannot define abstract objects through a biconditional. Dummett himself admits that this model can only give sufficient but not necessary conditions of the abstract and concrete. Nonetheless, this seems to make a very good distinction. As a result, any paradigmatically abstract object will be correctly defined. If any paradigmatically abstract objects are causal, they will be classified as such under Dummett's way of ostension.

³⁸ Dummett, 1973, p.481.

³⁹ Hale, 1987, p.47.

1.2.5. The Way of Negation

“The negative way: abstract entities have no spatiotemporal location; they do not enter into causal interaction; they are never indiscernible from one another” (Lewis, 1986a, p.83).

The way of negation is one of the most popular methods to make the abstract/concrete distinction. The thought is that abstract objects lack certain qualities which concrete objects have, and that is what makes them different. In fact the most common method of making the abstract/concrete distinction is to argue that abstract objects do not have spatiotemporal properties, and that they have no causal powers. It is important to note that philosophers don't tend to *define* abstract objects first and foremost as being non-*efficacious*; they define them as being non-spatiotemporal and only then do they make the inference that they are acausal.

We will start with the spatiotemporal condition. It makes no sense to ask where the number 2 is or when beauty was first created. This does not necessarily mean that the number 2 somehow exists in a Platonic heaven outside the spatiotemporal realm, it simply makes no sense to talk about where a number is. Because of their non-spatiotemporal nature, abstract objects are supposedly acausal, in every sense of the word. They can neither be causes nor effects.⁴⁰ Why are abstract objects acausal? The standard reason given is that in order for something to have any causal influence it must exist within space/time. How else is something supposed to interact with other objects?⁴¹ This link between not having spatiotemporal properties and being acausal is certainly not a clear one. There is a debate on whether or not objects are required to be in space/time in order to cause anything. This will be discussed further in chapter 3. We shall hence focus on the non-spatiotemporal condition when discussing the way of negation, as it does seem to be the more central one.

⁴⁰ Burgess & Rosen, 1997, p.20.

⁴¹ Cheyne, 2001, p.2.

It is noteworthy that the indiscernibility condition is rarely mentioned amongst philosophers, despite the way of negation being widely accepted. Though the literature explains why one would want to adopt the claim that abstract objects are acausal based on their non-spatiotemporal nature, it is hard to find a similar reason for indiscernibility. Many things do not exist in space/time, yet we can tell them apart. Neither Harry Potter nor Sherlock Holmes have these spatiotemporal properties for example, yet we rarely mistaken them for each other. There are many other properties which are unique to them, for example, one is attributed as a wizard and the other as a detective. They have different properties, thus they are discernible. Perhaps the way of negation is better viewed as a collection of three isolated claims. We could just argue that indiscernibility, being non-spatiotemporal and acausal are different ways of separating the abstract and the concrete.

The way of negation in terms of being non-spatiotemporal and acausal is the most popular way of distinguishing the abstract from the concrete, probably because of how many paradigm cases it ‘correctly’ identifies. This is especially true if we drop the temporal condition. Not surprisingly, there are only so many counterexamples to the way of negation. Fictional characters for example are created at some point in time according to creationists, yet they are abstract.⁴² If we are hesitant to call fictional characters and other similar objects concrete, we may not want to adopt the way of negation.

The term ‘abstract’ is ambiguous, so I will not argue for one way of making the distinction over another. For our purposes, the definition will not matter much, as there are causal entities under each of the proposed definitions for being abstract.

⁴² Van Inwagen, 1977.

1.3. Abstract Objects in the Causal Sphere

Why should we accept AOCO? This may seem like a controversial thesis but causal abstract entities have already been argued for by philosophers. These abstract objects are consequences of concrete events and will be discussed in depth. For each method of defining the abstract we can find different abstract entities within the causal sphere.

1.3.1. Fictional Entities

Fictional Characters such as Harry Potter, Sherlock Holmes and Peter Parker are often described as being abstract entities. We make references to fictional entities all the time. For example, I may say that “Harry Potter is a wizard” or “Sherlock Holmes is a brilliant detective”. I believe these statements to be true, so according to Quine, I am now ontologically committed to the existence of fictional characters. As Van Inwagen puts it,

“[These] sentences, if they are translated in the obvious way into the language of formal logic, will yield sentences that begin with ‘ $\exists x$ x is a character and...’. Therefore, anyone who believes that what these sentences say is literally true and who accepts what seems to be the obvious formal translations of these sentences, accepts the thesis that there are fictional characters” (Van Inwagen, 1983, p.73).

Realists of fictional characters believe that these entities exist and these realists are often creationists – they believe that fictional characters were created by their authors in every sense of the word. Of course, just because realists believe in the existence of these characters, that does not mean that they believe we can actually go meet Sherlock Holmes or Harry Potter. These realists are not delusional people who think these characters are not simply part of fiction. What realists believe is simply that

when we make claims about fictional characters, there is in fact a subject in which we are referring to.

“Mrs. Gamp appears to have incompatible properties. For consider the properties: being a woman [and] having been created by Dickens. . . . [O]n the theory I am proposing, Mrs. Gamp *has* only the second of these properties” (Van Inwagen, 1977, 308).

Since fictional characters are created by their authors, these entities are examples of abstract objects which are within the causal sphere. Creationism is the dominant view within realists of fictional characters. It would be an odd view to claim that fictional characters were not created at some time. How is it feasible to claim that Harry Potter existed even before the birth of J.K. Rowling? It would be odd to claim that fictional characters somehow exist necessarily. In fact, if any fictional character exists necessarily, all fictional characters, even ones that will never be spoken of or written about would exist necessarily. Here is the reason, if an author does not take any part in the creation of a fictional character, how would we distinguish fictional characters that were *actually* created from ones that were not? Surely, we do not want to suggest that the characters that are used and will be used in actual works of fiction just happen to perfectly align with the fictional characters that exist within some abstract realm. If realism is true for fictional characters, creationism seems like the most reasonable position.

This means that there are efficacious abstract entities within some definitions of the abstract. Fictional characters are classified concrete by some definitions, but not all. They lack spatiotemporal locations and are therefore abstract according to the way of negation. It seems feasible to reduce fictional characters to universals as they can be viewed as objects that are exemplified at different locations in the same time. If this is an adequate account, fictional characters will be abstract under the way of conflation. The way of ostension is a little more problematic. It is possible to point to Harry Potter on the silver screen, but it is not possible to point to him in a book. Does that make him a possible object of ostension or not? This is a mystery. Finally, the way of abstraction has fictional entities as concrete, as it would be difficult to pick out some

equivalence relation from Sherlock Holmes or Peter Parker. Nevertheless, fictional characters serve as a strong case of abstract entities being part of the causal order.

1.3.2. Linguistic Types

Other examples of abstract objects within the causal sphere are linguistic types. If someone asks whether or not the ninth and the eleventh words of this sentence are the same, how would you answer? In one sense they are the same – they are both the word ‘the’, but in another sense they’re not - they are two separate splodges of ink found on different points on the page. This illustrates the difference between a linguistic type and a token. In the example, we have one type but two tokens. They are two examples of the one word type – ‘the’.⁴⁴

We are ontologically committed to linguistic types. Take statements such as “How do you spell ‘addiction’?” or “the word ‘apple’ is pronounced with a short-A sound”. We often have statements we believe in where types play the role of the object linguistically. In chapter 1 of her book *Types and Tokens: On Abstract Objects*, Linda Wetzel provides an extraordinarily long list of examples of this form involving linguistic types. She also agrees that we are committed to their existence based on Quine’s methodology.

“Type talk is *pandemic*. It is not occasional; it is not unusual; it is the norm...We have to face the responsibilities posed by such talk of types: either concede that types exist, or give a systematic semantics for claims apparently referring to types. This book attempts to make the case for the greater plausibility of conceding that they exist” (Wetzel, 2009, p.21-22).

Wetzel argues that types are within our causal sphere. Words mean what the members of the linguistic community want those words to mean. The word ‘cake’ for example has a certain meaning and is pronounced in certain ways because the English speaking

⁴⁴ Wetzel, 2009, p.xi.

world decided at some point that it would be that way and this has been practiced over a period of time. Therefore there is no mystery as to how we interact with such objects.⁴⁵

Furthermore, after we decide on the meanings of words, we obviously go and use them. Wetzel is therefore proposing an abstract object that can be both part of causes and effects. This model for linguistic types has many appeals. First of all it allows realists to account for the role the linguistic community plays in the formation of words. It just seems obviously true that words are defined and changed by people who use them. The traditional Platonist model simply gives us unintuitive results. To say that the properties of words, including their existence is completely independent and isolated from us is a proposition that leaves us scratching our heads.⁴⁶ This model allows us to explain language as an ongoing and organic process. This is a brand of realism that involves abstract objects being very active causally.

It is important to note that linguistic types are classified as abstract under every definition we've looked at. Once again, these entities are abstract according to the way of example. They lack spatiotemporal properties and are therefore abstract under the way of negation. Like fictional characters, they may perhaps be reduced to universals. Therefore, we get similar results with the way of conflation. The way of ostension has word types categorised as abstract. Though it is possible to pick out a word token ostensibly, we cannot do the same to a word type. We would need to use a functional expression eg. 'the word that appears on the first line of *Hamlet*'. In reference to the way of abstraction, we can define word types through equivalence relations as follows: let W1 and W2 be different word tokens. Then W1 is the same word type as W2 iff W1 and W2 can be swapped without there being any major differences in meaning, according to the linguistic community. This is of course only one rough and unpolished way in which followers of Noonan and Hale could analyse word types. I am certain that they could analyse linguistic types in a more adequate fashion. It is enough for our purposes to just show that this sort of formalisation could be convincing. Linguistic types therefore act as a strong example case for every common definition of the abstract.

⁴⁵ Wetzel, 2009, p.123.

⁴⁶ Wetzel, 2009, p.123.

1.3.3. Sets

Sets are another example of abstract objects within the causal sphere. A set is simply a collection of different objects eg. five different crayons may form the set of crayons on my table. Note that sets can be made up of all sorts of objects; my loyalty and my officemate's loyalty may form the set of loyalties in our office. Penelope Maddy has demonstrated how the use of sets could help us explain mathematical phenomena.

“Now suppose you take a set theoretic perspective and again ask why multiplication is commutative. Here an answer is forth-coming: because if A and B are sets, then there is a one-to-one correspondence between the Cartesian products $A \times B$ and $B \times A$ ” (Maddy, 1981, p.499).

We therefore have theories we believe to be true that are based on the use of sets. By the Quinean method, we are currently ontologically committed to the existence of sets.

As long as concrete objects form part of a set, this abstract object can be part of the effect caused by only concreta. For example, if I put all my crayons in a bag, the set of crayons now have the property of being in my bag. When my sister was born in 1994, the set containing all of my family members gained a new member. Sets therefore seem to be good examples of abstract objects within the causal order.

One may object that this assumes that sets are located in the same place as its members. That is fair enough, but we have a strong argument for efficacious sets, even if this is not true. In 1994, my only sister was born, so since 1994, the set of my sister(s) has existed containing only one member. Here's the question: what was the set like before 1994? Did it even exist? There are three possible answers for this:

- a) The set was exactly the same. It existed as a one-member set.

- b) The set existed but had no members. It was the null-set \emptyset .
- c) The set did not exist.

For us to make sense of (a), we would need to accept a four-dimensionalist theory of time. A presentist account leads to odd consequences. The set of {my sister(s) at a particular point in time before 1994} would only have one member, despite me being an only child then. This just seems wrong, because the presentist cannot make references to any sisters I will have in the future. They therefore cannot explain how I could possibly have a sister while being an only child. A four-dimensionalist account however has no such problems. They explain (a) as follows, let us assume that I will only ever have one sister. If we looked at a timeline and searched for the number of sisters I will ever have, we would find that I only have one sister. The set of my sisters will hence always be one-membered. This account assumes a four-dimensionalist theory of time however, which would not appeal to everyone. Here is a bigger problem: If anyone asked me before 1994 how many members there were in the set of my sister(s), I would say none. However, according to this view, I would be wrong. In fact, as long as my parents have the ability to reproduce, I cannot know the answer to this question. It simply seems like more intuitive to say that the set gained a member when my sister was born.

Options (b) and (c) on the other hand both involve causal changes. If we accept (b) as the desired view, we would have the set of my sister(s) changing from the null-set to a one-member set when my sister was born. This is a causal story.⁴⁷ If we accept option (c) instead, we would be claiming that the set of my sister(s) did not exist before 1994, but that changed since then. In essence we would be proposing sets coming into existence and hence adopting a creationist model of sets. Much like the creationist theory of fictional characters, this would be considered a case of causation.

Sets are abstract according to the way of example. Whether or not they will be classified as such with the ways of negation and ostension depend on your theory of sets. Sets under a mereological model would be defined as concrete, whereas they would be abstract with other models. Things are more complicated with the way of

⁴⁷ If it is not a causal change, what sort of change is it? This would remain mysterious.

abstraction. Do all members of a set share the right forms of equivalence relations? Perhaps, but the answer is not clear.

1.3.4. Tropes

Tropes are properties, but unlike universals, they are abstract particulars. Examples of tropes are the particular colour of my car, Barack Obama's charisma and the shape of my football. These should be compared to non-particular abstract entities such as redness, charisma and circularity in general.⁴⁸ The concept of tropes originated from D.C. Williams in two papers from 1953. Rather than making a sharp distinction between the abstract and the concrete, he argues that concrete objects can be parts of tropes.⁴⁹ The particular redness of my car for example at least partially consists of my actual car, although we may not want to say that the concrete objects are the only components of tropes. Since these concrete parts exist within space/time, we can say that tropes at least partially exist within space/time. In *On the Elements of Being: II*, Williams *explicitly* states that tropes are part of the causal order.

“[The] actual strength of a girder is what holds the bridge up; the heat of an atomic explosion scorches a city; and in fact, as we have observed, all of the efficacies we know are of abstracta on abstracta” (Williams, 1953b, p.184).

The redness of the car makes me buy it; the shape of the football caused it to bounce a certain way, Obama's charisma got him elected. Williams happily assigns these causal roles to tropes. This efficacious nature has some intuitive appeal. If these entities exist, these causal stories all seem reasonable, at least in everyday language. Williams' tropes are therefore examples of abstract objects that are compatible with AOCO.

⁴⁸ Williams, 1953a, p.4-7.

⁴⁹ Williams, 1953a, p.7.

Are tropes really abstract? This is a difficult question to answer. Tropes aren't paradigmatically abstract, so they are concrete under the way of example. The way of negation would classify tropes as concrete, as they all exist within space/time. However, it is not as clear with the other ways. Different tropes will be categorised differently. For example, specific colours and shapes would have to be concrete under the way of abstraction because these singular tropes cannot hold between multiple concrete objects. However, what about the specific influence of Microsoft? Surely this is a trope that holds in more than one singular place and time. Therefore colour tropes will be concrete under the way of abstraction, but Microsoft's influence as an entity for example will be abstract. The way of abstraction isn't alone in its piecemeal categorisations of tropes; the way of ostension is the same. It is reasonable to claim that we can point to the redness of a car or the specific shape of a ball, so these tropes will be concrete here. Microsoft's influence however certainly cannot just be picked out ostensively and must be picked out through a functional expression (eg. the influence of Microsoft in American economics). With the way of conflation, we have already discussed whether tropes can be reduced to universals. It is more reasonable to say that universals are reducible to tropes, rather than other way round. It is likely then that tropes will be concrete according to the way of conflation. Nevertheless, we can still say that some tropes are abstract according to the ways of abstraction and negation. Furthermore, Williams claimed that concreta may not be the only building blocks of tropes. If there are non-concrete building blocks, surely these elements would be abstract. Despite the philosophical tradition, we have reasons to believe that at least some tropes can be considered abstract under certain well known definitions.

1.4. Opening the Floodgates

We hence already have many cases of causal abstract objects under the different listed definitions of being abstract. Therefore, regardless of how we define the term, one has to acknowledge the presence of these proposed cases. The listed examples are all cases in which concrete events can cause changes within abstract objects. There are already decent examples of abstract objects being part of the causal order. One may argue that these examples are not enough; what we are really interested in is whether or not abstract objects can cause changes themselves. Well, given that they are already in the causal sphere, why not? Is there really any reason why this causal arrow should only point one way? I will argue for abstract entities being part of causes in the next chapter. Adopting AOCO solves many problems. For example, it allows us to explain phenomena using abstract objects. Statements such as ‘his pride lead to his demise’ and ‘that word made the poem smoother’ are perfectly sensible ones, but we cannot make sense of them literally without causal, abstract entities. The advantages are not just linguistic; we often *do* believe that properties for example are responsible for changes in our world. For example, we may believe that greed is the chief motivator of capitalism. Furthermore, AOCO solves an important epistemic issue raised by Benacerraf. This will be discussed in depth in chapter 2. In conclusion, abstract objects are already postulated as consequences of causal phenomena, allowing abstract objects causal powers solves many problems, and finally there isn’t any reason why we shouldn’t allow abstract objects these causal powers given that they are already in the causal sphere. I will argue for abstract objects being causes of knowledge in the next chapter. For now, we should accept that abstract objects can be part of causal effects.

2 - Epistemology

2.1 The Benacerraf Argument

Abstract objects can be causes of knowledge. This has some intuitive value. For example, I know I have to buy 5m of tin for my triangular roof because of Pythagoras' theorem. This however has not been the view philosophers have historically chosen. I will show why we should not accept this claim and why we should accept that abstract objects can cause knowledge. This will then naturally be motivation for us to accept AOCO – the thesis that abstract objects are part of the causal order.

A very popular argument made against realists originated from Benacerraf, on the topic of mathematical entities. Mathematical entities were normally taken to be abstract entities taken in the traditional sense – objects which lack spatiotemporal properties and hence acausal. Benacerraf then asked if these objects contain no causal powers, how can we know anything about them? Surely, if we know about anything about the number 3 (taken as an abstract entity), there must be a causal relationship between us and that number. In his own words,

“If, for example, numbers are the kinds of entities they are normally taken to be, then the connection between the truth conditions for the statements of number theory and any relevant events connected with the people who are supposed to have mathematical knowledge cannot be made out”
(Benacerraf, 1973, p.673).

This argument has been extended beyond just mathematical entities and is now used to argue against the existence of abstract objects in general. If abstract entities are acausal and knowledge requires a causal link between ourselves and the object, how

could we have any knowledge of abstract objects? The argument can be summarised as follows:

- 1) We can only know about things that are in the causal order.
- 2) Abstract objects are not part of the causal order.
- 3) Therefore, we can't know anything about abstract objects.

Other philosophers have noted the strength of Benacerraf's objection, including Cresswell (2010), Chihara (1990), Liggins (2006), Maddy (1981), Bigelow (1988) and many more. In fact, it would be no stretch to say that Benacerraf's objection has become the most prominent argument against realism about abstract objects.

My main claim in this chapter is (i) that once we accept AOCO, we have an easy solution to Benacerraf's problem and (ii) that other purported answers to Benacerraf are inadequate. I will first elaborate on why we should not accept Benacerraf's conclusion. In 2.3, I will explore the option of denying (1) and adopting another causal theory. I will argue that we should keep the causal theory but adjust it slightly. I will also discuss Steiner's objection to the causal theory of knowledge. Furthermore, there are certain solutions to Benacerraf's problem which require us to deny or seriously amend the causal theory of knowledge. These theories are positive answers to Benacerraf in the sense that they try and show how we can have non-causal cases of abstract knowledge. I will discuss these in 2.3. 2.4 will be devoted to whether or not the argument is valid – an option that Max Cresswell explores but one that we should not accept. Finally I conclude that the best option against Benacerraf is to object to (2) and accept AOCO.

2.2 The Spurious Conclusion of the Benacerraf Argument

There are at least two problems to Benacerraf's conclusion that we cannot know about abstract objects:

1. It is counter-intuitive and clashes with propositions/theories we believe to be true.

David Lewis for example has argued that we cannot accept this conclusion when it comes to mathematical entities. If anything, we know a lot more about mathematics than we do about epistemology.⁵⁰ Surely, it is hubris to make great changes to the current formulations of mathematics just because they are based on abstract objects. This reasoning can be extended to cover other abstract objects. Take properties, I know that my plastic cup shares a property with my shirt – the property of being blue (or at least, I seem to know). I also know that Harry Potter is a fictional character and that J.K. Rowling is not. Having to sacrifice these knowledge claims does not seem attractive. Notice that none of these abstract truths would seem like controversial statements of knowledge outside the philosophy classroom, so they have intuitive value.

One may object with the following argument: these are not cases of abstract knowledge; they are simply if-then-ist statements. *If* abstract objects exist *then* these statements would be true, but that does not mean we know anything about abstract objects. This argument seems powerful, but should we be convinced by it? Phenomenologically, it certainly does not *feel* like we're making if-then-ist statements. Imagine if we asked some mathematicians whether 2 plus 2 equals 4 and they all say yes. Then we ask them, 'do you just mean that 2 plus 2 equals 4, or that *if* numbers exist *then* 2 plus 2 would equal 4?'. How would the mathematicians respond? They would certainly all be confused. We don't seem to be making if-then-ist propositions at all; the assertions are not hypothetical. It certainly seems like a reach for us to label these abstract statements if-then-ist. If a) we are stating facts and b) these facts seem to involve abstract objects, then surely it is reasonable to say that we have abstract knowledge.

2. Denying that we have abstract knowledge gives the abstract realist an unnecessary disadvantage.

⁵⁰ Lewis, 1986a, p.108-9.

This point is perhaps less important than the previous one but it is still worth stating. Imagine that you were approaching these ontological debates for the very first time and you wanted to weigh up the pros and cons (without bias) of being an abstract realist, what advantages would abstract objects give you? Well, once again, it seems to be an intuitive and simpler position. When I say that ‘I know 2 plus 2 equals 4’, it seems like I am quantifying over abstract entities. If mathematical entities exist, this is literally true. On the other hand, if nominalism is true, we would need to repackage all our mathematical knowledge, much like Hartry Field did.⁵¹ Abstract realism allows for abstract knowledge, which gives us a more intuitive account of these cases. However if Benacerraf is correct, abstract realism loses this advantage. Realism would in fact be in *worse* shape for bloating our ontology without giving us any form of epistemic advantage. We cannot simply ignore Benacerraf’s argument, but perhaps there is a way around it. If Benacerraf is wrong and abstract knowledge is plausible, this would completely change our ontological debates.

2.3 Deny the First Premise

One option the realist has against Benacerraf is to deny the causal condition of knowledge. Recall that Benacerraf’s problem only arises because abstract objects do not meet the criteria required by the causal theory, therefore if we reject premise 1, the problem fades. However, there are problems with this. I will show the historical reasons for postulating such a condition as well as why alternative theories of knowledge cannot answer these questions adequately. There are other reasons for objecting to the causal theory of knowledge, ones that do not give alternative theories themselves. These objections will equally be dealt with.

⁵¹ Field, 1980.

2.3.1.a. The Causal Theory of Knowledge and Gettier Problems

The traditional theory of claims that there are three necessary conditions for knowledge – justification, truth and belief. Together, these conditions are sufficient. This theory has been embraced for a long time until Edmund Gettier raised some famous objections to it. The argument is simple; there are cases where something is a justified, true, belief without it being a case of knowledge. For example, Smith and Jones are two men being interviewed for the same job. Smith has strong reasons to believe that Jones will get hired (the owner of the company told him so). Furthermore, he knows that there are ten coins in Jones' pocket. Together, he forms the belief that the man who will get the job has ten coins in his pocket. However, assume that Smith was mistaken and he got the job instead. As a matter of chance, Smith also has ten coins in his pocket; this means that he has a true belief that the man who gets the job has ten coins in his pocket. This is a justified true belief, yet we wouldn't say that Smith *knew* it. Therefore we should amend or reject the JTB account (Gettier, 1963).

The reason Gettier cases work in the JTB account is that the connection between the knowledge and the belief is not strong enough. In other words, the reasons Smith believed what he did were different from the reasons that made it true. Smith's belief was true, but it was just a coincidence. Whatever amendment we make to our theory of knowledge, it would need a strong link between belief and truth.

The causal condition would obviously provide us with such a link. One could explain Smith's case by the fact that there was no causal link between Smith's belief and the facts. Let 'S' be someone who is capable of knowledge, let '*p*' be a proposition, let '[*p*]' be the fact that *p*, and finally, let 'B*p*' be S's belief that *p*. Put more formally, the causal theory of knowledge states that

(CT) S knows that *p* if and only if B*p* was caused by [*p*].

There are many objections to the casual theory, but I believe they can all be overcome. In particular, we will discuss the issue of extended causal connections and the problem of inductive knowledge.

There are various counterexamples to the causal theory of knowledge. For example, I'm sitting at home and I hear the rain falling. Whenever it rains, the roads become wet and whenever that happens, there are an increased number of accidents on the freeway. Since I know this from hearing the rain fall, I come to believe that there will be more accidents on the freeway. Note that my belief that there will be more accidents on the freeway was not directly caused by the fact that there will be more accidents on the freeway. The causal theory of knowledge needs to be able to account for these extended and complicated causal stories. Furthermore, there's the problem of logical consequences. In 1871, Mendeleev released the first periodic table. He noted that there were gaps and hence postulated the existence of elements that were not yet discovered. Germanium is such an example. Despite never having any causal interactions with this element, Mendeleev knew many facts about it such as its melting point, atomic mass, oxide activity etc.⁵² The problem of logical consequences also appears in cases of *a priori* knowledge. I may come to know *a priori*, without any *a posteriori* aids, that all triangles have three sides. How can this be? The fact that all triangles have three sides is supposedly acausal, so the fact that *p* could not have caused the belief that *p*. Therefore we have many counterexamples to the causal theory of knowledge. Any respectable theory of knowledge should account for such cases.

Colin Cheyne acknowledges these problems and hence suggests that we should amend the causal condition. He has very specific types of causal connections in mind and calls these 'k-causal connections'.

Let '*q*' and '*r*' be other propositions, let '*[q]*' be the fact that *q*, let '*[r]*' be the fact that *r*, let *Bq* be S's belief that *q*, and finally let *Br* be S's belief that *r*. *[p]* is k-causally connected to *Bp* iff either:

⁵² Cheyne, 2001, p.108

- a) $[p]$ is a cause of Bp , or
- b) $[p]$ and Bp has a common cause, or
- c) there is some $[q]$ such that $[p]$ is a logical consequence of $[q]$, and $[q]$ is the cause for Bp , or
- d) there is some $[q]$ such that $[p]$ is a logical consequence of $[q]$, and $[q]$ and Bp have a common cause, or
- e) there is some $[q]$ and $[r]$ such that $[q]$ and $[r]$ are k-causally connected to Bq and Br , respectively, and $[p]$ is a logical consequence of $[q \& r]$.⁵³

We can hence adjust the causal theory to accommodate different types of knowledge. Using condition (b), we can explain the rain example. Imagine a branching causal chain. First of all, the fact that it is raining caused my belief that it is raining, then that caused my belief that the roads are wet, which finally caused my belief that there will be more accidents. At the same time, the fact that it is raining caused the fact that the roads are wet, which caused there to be more accidents. Therefore, there was ultimately one common cause between the fact that there are more accidents on the road and my belief of it. The germanium and triangle cases can be explained by (c). We will start with the former. There are some general rules of atomic structure. The fact that these rules exist caused Mendeleev to believe that they exist. This then caused his belief that germanium has certain properties. At the same time, the fact there are these general atomic rules has the logical consequence of germanium following these rules and hence having certain properties. By (c), this story is a k-causal one. The triangle case is simpler. This is a standard definition of triangles: a shape that has three connected and straight sides. From this definition, we come to believe that all triangles have three sides. From the definition of a triangle, there is also a logical consequence that says that all triangles have three sides. Therefore we can also explain away this counterexample to (CT) with condition (c) of the k-causal theory. Given that we can deal with these counterexamples by adopting the k-causal theory, I propose adopting it rather than dropping the causal theory of knowledge altogether.

(KTK) S knows that p iff Bp was k-causally connected to $[p]$.

⁵³ Cheyne, 2001, p.67

K-causation may adequately explain our intuitions about causation, but are there better options? Of course the causal theory of knowledge is only one way of answering the Gettier problems. I will now go through the alternative theories, seeking both necessary and sufficient conditions of knowledge. I will test to see whether these theories could answer the same questions better.

We will start with different variants of reliabilism. Reliabilist theories are externalist in nature. As long as the epistemic justifications are reliable in obtaining knowledge, the reliabilist does not demand any form of internal justification from the knower. That is, as long as the methods in obtaining knowledge are reliable, the subject could have no idea how he/she came to have this knowledge, yet it would be justified. For example, a psychic has random and accurate visions about future (which she believes). She has absolutely no clue how this clairvoyant process works, but according to the reliabilist, this is sufficient to be called knowledge. The causal theory of knowledge is one type of reliabilism, though not all reliabilists demand a causal condition.

2.3.1.b. Reliabilism - Truth Tracking Theory

According to the truth tracking theory, one's belief that p and the fact that p has the same truth conditions. In worlds where p is true, one believes that p ; and in worlds where p is false, one does not believe that p . For a belief to be counted as knowledge, one cannot have that belief without that belief being true. More formally,

Let 'S' be someone who is capable of knowledge, and let ' p ' be a proposition.

(TTT) S knows that p if and only if S believes that p and in the closest possible worlds where p is true, S believes that p ; and in the closest possible worlds where p is false, S does not believe that p .

Robert Nozick describes it as follows:

“To know is to have the belief that tracks the truth. Knowledge is a particular way of being connected to the world, having a specific real factual connection to the world: tracking it” (Nozick, 1981, p.178).

Unlike the old JTB account, TTT has a better time dealing with Gettier cases. Take the Smith and Jones case, would Smith still believe that the man who gets the job has ten coins in his pocket (call this proposition p) if it wasn't true? The closest possible world in which p is not true is one where Smith does not have ten coins in his pocket. In this world, Smith still has the false belief that p . Therefore according to TTT, the Gettier cases do not qualify as knowledge.

Note that TTT is very similar to the causal theory of knowledge. In fact, if we compare it with a counterfactual theory of causation, especially one that involves possible worlds, the two theories are almost exactly the same. According to a counterfactual theory of causation for, event p causes event q iff p and q both occur, and had p not occurred, q would not have occurred either.⁵⁴ Now take just about any true proposition that someone believes in, if it qualifies as knowledge under TTT, it would qualify as causation under Lewis' definition as well!

However, one does not have to adopt a counterfactual theory of causation, in which case, TTT does not necessarily describe a causal process. This leads to a subtle difference between TTT and a causal theory of knowledge – the causal theory requires the connection between truth and knowledge to be a causal one, whereas TTT doesn't. Imagine a non-causal process that links the fact that p to the belief that p – the clairvoyant story will suffice. The clairvoyant has a non-causal process that allows her to track the truth. Let us say that when she receives visions of a murder, a murder always happens in her town and when she does not receive a vision, no murders take place. According to the causal theory of knowledge, the clairvoyant would not know that murders were or were not taking place, whereas with TTT, it would.

⁵⁴ Lewis, 1973, p.560.

What are the advantages that TTT has over the causal theory? It seems as though the only cases where TTT is preferable are ones which involve some sort of non-causal process that suffices for knowledge. Importantly, TTT employs a counterfactual framework without it being classified as causation. That means we can explain mathematical truths (for example) using abstract objects but without employing AOCO!⁵⁵ However, there are challenges for the truth tracking theorist. They need to a) explain what these non-causal processes are and b) defend a theory of causation other than the counterfactual theory, such that whatever is described under TTT isn't necessarily causal. There are very good reasons to accept the counterfactual theory over rival theories; this debate will be discussed in depth in chapter 3. Regardless, a convincing argument for other theories of causation needs to be made Nozick and his followers in order for TTT to be an attractive option.

2.3.1.c. Reliabilism - Armstrong's Theory

The next reliabilist theory we shall consider is Armstrong's reliable indicator account. He says that for our beliefs to be considered knowledge, they have to be reliable indicators of truth. Armstrong makes an analogy with a thermometer – the thermometer can reliably tell us what temperature it is; analogously, our processes used in gaining knowledge should also be reliable. More formally,

Let 'S' be someone who is capable of knowledge, let ' p ' be a proposition, and finally, let ' Bp ' be S's belief that p .

(AT) S knows that p if and only if Bp was brought upon by a law-like connection in nature such that whenever S believes that p , then p is true.⁵⁶

⁵⁵ Thanks to Bill Fish for this point.

⁵⁶ Armstrong, 1973, p.168.

There are some important points about Armstrong's version of reliabilism that needs to be covered. First of all, note that being reliable does not entail being infallible. A thermometer does not accurately measure the temperature all the time, yet it is reliable. Likewise, our methods for obtaining knowledge do not need to be 100% accurate; it simply needs to have a high success rate. At this point the reader may wonder if there are any important differences between the TTT and Armstrong's theory and whether or not they are really just the same. This is the difference: Armstrong's theory states that there needs to be a specific method or law of nature that acts as a reliable bridge between fact and belief whereas TTT does not require such a process. For example, imagine a man who through an infinite amount of mysterious processes always comes to know next week's lottery numbers. Sometimes the numbers appear to him through an arrangement in his tealeaves, sometimes he receives an anonymous letter with all the numbers inside, other times he simply has a vision which tells him what all the winning numbers are. According to the TTT, the man knows what the winning lottery numbers will be. However according to Armstrong, this does not qualify as knowledge since there is not a single reliable process that would count as justification.

A problem with this theory is that we then need a definition or some other explication of what a reliable process is. Colin Cheyne argues that we should take the thermometer analogy further and say that it is a *causal* connection that makes a reliable process.⁵⁷ Armstrong himself has denied that that we should use a causal description. By his own words, "[it] is an ontological condition. It is not, however, a *causal* connection" (Armstrong, 1973, p.169). Cheyne is worried about Armstrong's causal scepticism, and rightly so. If these reliable connections are not at least part of a causal chain, then what are they? Much like TTT, this leaves Armstrong's account with a puzzle he needs to solve. The causal account at least gives us a clearer account of what this mysterious process should be. The fact that these are causal cases seem to explain why these are reliable cases of knowledge, whereas the reliable indicator theory offers no such explanatory advantage. Once again, Armstrong does not believe that these reliable processes are causal, so we cannot even use the causal explanation within the reliable indicator account.

⁵⁷ Cheyne, 2001, p.86-87.

Another problem with Armstrong's account is that it does not always deal with Gettier problems better. Though it can with some cases, the Jones and Smith story for example is one that this account struggles with. What process lead Smith to believe that the man who will be hired has coins in his pocket? It was the fact that he was told Jones will get the job from the interviewer. For the most part, this seems like a pretty reliable process for knowledge about an interview. Defenders of Armstrong may object, saying that getting false information from an interviewer that is prone to trick the interviewees is not a reliable process. This seems desperate. How fine-grained do our processes need to be? It seems more reasonable to say that obtaining information from an interviewer is a reliable process but not an infallible one. The information may be true in most cases, but not all. So called 'reliable-processes' naturally invite exceptions to the rule and hence can never completely rule out counterexamples. Given these problems, we should refrain from adopting this alternative theory as the causal condition deals with these problems better.

2.3.1.d. Explanatory Theory

Alan Goldman⁵⁸ has another theory of knowledge from his book *Empirical Knowledge* (1998). He states that the fact that p must be a prominent reason for the belief that p . This is further explained using probabilities and possible worlds.⁵⁹ Together, we get the following theory of knowledge:

Let 'S' be someone who is capable of knowledge, let ' p ' be a proposition, let ' $[p]$ ' be the fact that p , and finally, let ' Bp ' be S's belief that p .

(ET) S knows that p if and only if either $[p]$ makes Bp 'significantly' more probable or $[p]$ explains Bp .

⁵⁸ Not to be confused with Alvin Goldman.

⁵⁹ Goldman, 1988, p.25-32

Note that this theory is separate from Armstrong's. Just because $[p]$ makes Bp significantly more probable does not mean that there is a reliable process between the two. Let us imagine that the conspiracy theorists were right and aliens have once visited the earth, leaving behind only small traces of evidence. It is reasonable to say that the fact that the aliens visited made the belief that aliens visited the earth significantly more probable. So let p here be the proposition that the aliens visited the earth. We would not want to call this a link between $[p]$ and Bp a reliable process, since the majority of the population still do not believe that aliens have visited.

How well does the explanatory theory of knowledge stand up to criticism? First of all, there is the question of *how* $[p]$ would make Bp more probable. Can we not explain this with causal chains? Goldman insists that this ambiguity is a strength of his theory. There are some defining features in all cases of knowledge; however they differ in degrees of stringency according to their contexts. "What counts as empirical knowledge in a casual conversation might not count in a court of law, in a philosophy classroom, or in a physics laboratory" (Goldman, 1988, p.21). Cheyne raises a more important worry – precisely what counts as 'significantly more probable'?⁶⁰ Take the weight of an atom for example. We may know the weight of it, but if it were just a little bit heavier, so little that we cannot measure the difference, would the probability of our beliefs changing really be that significant? We can take this worry even further: Goldman himself raises this example: there is a hospital in which the equipment is powered from the public utility, but also has a backup generator. To say that the public utility explains the running of the equipment would be insufficient, because in counterfactual cases, the backup generator would power the equipment. Therefore, the probability of the equipment being able to run is not significantly raised by the public utility.⁶¹ Goldman himself does not see this as a problem because this is an adequate account of explanation. A good theory of explanation should show that the equipment did not run solely on the public utility. What Goldman does not realise however, is that his theory is still not adequate as a theory of knowledge. Take the popular example of the father learning that his son is not a murderer.⁶² In this case, the father hears about a murder on the radio news. From its contents, he induces and believes

⁶⁰ Cheyne, 2001, p.89.

⁶¹ Goldman, 1988, p.33-4.

⁶² Cheyne (2001) for example talks about this at p.90.

that his son cannot be the murderer because of the time the murder took place, where it was etc. However, let us assume that the father would have believed that his son was innocent regardless of the news. He may believe this because of his fatherly instincts coupled with the fact that his son is only 14 (and hence unlikely to be able to formulate such a murder). In this case, the fact that the son did not commit the murder can only be a partial and inadequate explanation to the father's belief that his son is innocent. However it seems perfectly feasible for us to say that the father knew that his son was innocent given the news on the radio. We seem to have a counterexample to the explanatory theory of knowledge. Furthermore, can we have explanations and raises in probability without any sort of causal chain whatsoever? This is not something we should easily accept. Given that we adopt Cheyne's concept of k-causal chains it seems unlikely that we can find cases of a fact either raising the probability of a belief or explaining it without using any causal concepts. With the prevalence of these worries, we should not embrace this theory over the causal theory of knowledge.

2.3.1.e. Internalism - Evidentialism

Evidentialism is a very different approach to reliabilism. Reliabilism employs an external account of justification, meaning that as long as the belief is brought about by a reliable process, it can bring about knowledge. In comparison, the evidentialist has an internal account of justification, meaning that any sort of justification involved in obtaining knowledge must at least be evident to the knower. Richard Feldman and Earl Conee provided one of the first accounts of evidentialism. According to them,

“Doxastic attitude *D* toward proposition *p* is epistemically justified for *S* at *t* if and only if having *D* toward *p* fits the evidence *S* has at *t*” (Conee & Feldman, 1985, p.15).

Knowledge is one of these doxastic attitudes.

Let 'S' be someone who is capable of knowledge, let ' p ' be a proposition, let ' Bp ' be S's belief that p , and finally let T be the moment in time when S formed her belief that p .

We can hence reformulate this account specifically for knowledge as

(EVT) S knows that p if and only if Bp fits with S's evidence at T and p is true.

For example, if I walk into a room and come to believe that the lights are on, I can trace my belief to certain pieces of evidence. I may have noticed that the room is well lit and that the light switch is in the 'on' position. My belief that the lights are on is justified by my evidence and this evidence is crucial to knowledge, according to the evidentialists. Since evidentialism is an internalist theory of knowledge, it disagrees fundamentally with the reliabilists' theories when it comes to justification. Imagine that a man had a thermometer device unknowingly implanted in his head such that he would receive constant beliefs about what temperature it was. If he has no idea how these beliefs are formed then he would not have knowledge according to the evidentialists, whereas the reliabilists would be happy saying that the man in fact knows how hot or cold it is.

However, EVT alone is not enough, because it still has issues with Gettier problems. Smith's evidence lead to the conclusion that the man who will get hired has ten coins in his pocket, yet this is not a case of knowledge. The reliabilists hence need a further condition that prevents Gettier cases. Conee & Feldman for example argue that there can't be any piece of evidence that S can find at T which would undermine p .⁶³ So let ' E ' be a piece of evidence, adding this to EVT gives us:

(EVT*) S knows that p if and only if p is true, Bp fits with S's evidence at T, and there are no E 's available at T that would undermine p .

Internalists often defend their own position (when contrasted with externalist theories of knowledge) by trying to raise epistemic questions surrounding externalism. In

⁶³ Conee & Feldman, 2004, p.93

Laurence Bonjour's 1980 paper *Externalist Theories of Empirical Knowledge*, he raises various examples involving a clairvoyant. In one of the stronger cases for example he asks us to imagine this clairvoyant coming to believe the fact that the president is in New York. This clairvoyant believes that he is a clairvoyant but has no reason to. He has on many occasions tried to prove his clairvoyance but has failed every time (his powers malfunctioned under those circumstances). Is he justified in believing that the president is in New York? Bonjour says not.⁶⁴ Conee and Feldman raised other counterexamples in *Internalism Defended*. Bob and Ray both read in a forecast that it is warm today. Bob then walks outside and feels for himself that it is indeed warmer today. According to Conee and Feldman, Bob then experienced a "mental change which so to speak 'internalized' the actual temperature. Ray had just the forecast to rely on" (Conee & Feldman, p.236). This is supposed to show that Bob is more justified than Ray because of internal processes that externalists don't mention.

The first case highlights a common mistake about objections from internalists; it simply muddles two different concepts –

1. how to decide whether an existing belief counts as knowledge and
2. how to put ourselves in the best position to receive knowledge.

This is a point that Alvin Goldman makes very clear in his 1980 paper *The Internalist Conception of Justification*.⁶⁵ The problem is that the externalist wants to focus on the former problem whereas the internalist wants to focus on the latter. In the clairvoyant case, the externalists insist that the clairvoyant knows whatever it is that his clairvoyant powers allowed him to know, but the internalists object to these types of cases with pragmatic objections. How could the clairvoyant justifiably make decisions based on his psychic beliefs over rational evidence? How can he say that he knows anything at all? They are questions under point 2, but they are *not* questions of the first kind. The internalists' objections to these cases really don't have much to do with whether or not the psychic has obtained knowledge. Allow me to phrase the usual internalist objections in externalist terms.

⁶⁴ Bonjour, 1980, p.72.

⁶⁵ The point is made in p.37-39.

- i) “How does the clairvoyant know that he knows?”
- ii) “How would the psychic deal with a piece of knowledge that doesn’t seem justified?”
- iii) “What if this new knowledge seems to contradict other things that he knows?”... etc.

In other words, we should look at these problems as separate questions to be asked *after* we have established that the psychic does in fact does have knowledge in these case. There is no *a priori* reason to believe that one cannot have knowledge that they cannot explain or justify verbally. If we are truly just asking question 1 – “is this a case of knowledge?”, then we should only be judging these cases as a neutral observer. We should *not* be asking questions aimed at the believer, because questions such as the three above are further epistemic questions separate from whether or not something counts as knowledge.

Furthermore, Conee and Feldman’s point can be answered by the externalists. Can we not explain Bob’s superior justification in externalist terms? Sure we can, feeling the heat for yourself is a more reliable process of telling the temperature than reading about it on a forecast. The internalist has no advantage in this example over the externalist. In conclusion the internalists really don’t have an advantage over the externalist theory of knowledge when it comes to the questions that matter.

2.3.1.f. Contextualism

One theory of knowledge has been becoming popular in recent years; that theory is contextualism. The other theories all have one uniform standard as to what counts as knowledge, what counts as knowledge in everyday life should be the same as what counts as knowledge in philosophy. Contextualism objects to this. What counts as knowledge should be context-dependent. Though it may be wrong for us to make claims such as ‘I know I have hands’ in a debate with a sceptic, it seems odd to say

that such a claim would be problematic under normal circumstances. I will begin with Dretske's form of contextualism then compare it with Lewis and Cohen's. Ultimately, contextualism should not replace the causal theory of knowledge. I will explain my reasons for this.

It all starts with Dretske's account. Let 'S' be someone who is capable of knowledge, let ' p ' be a proposition, and let ' q ' be another proposition. There is a widely accepted principle in philosophy that states that

If subject S knows that p and that p entails q , then S also knows that q .

Call this the 'epistemic principle of closure' or 'EPC' for short. Dretske denies EPC.⁶⁶ Sceptical arguments often rely on a *modus tollens* form of EPC. If EPC is correct, that logically entails the following:

S knows that p entails q , if S does not know that q , then S does not know that p .

For example, say we are at a zoo, looking at some zebras. Now take the following claims:

- 1) I know that I am looking at some zebras.
- 2) If I don't know whether or not these 'zebras' are really just cleverly disguised donkeys, then I don't know that I am looking at some zebras.
- 3) I don't know whether or not these 'zebras' are really just cleverly disguised donkeys.⁶⁷

This seems like a paradox, the three propositions cannot all be true. The sceptic then denies premise 1 and robs us of such a basic knowledge claim. Dretske on the other hand denies premise 2, which is actually just EPC applied. His reasoning is that epistemic operators such as 'I know that', 'I believe that' etc. do not *penetrate*. This means that the truth values of the propositions are not necessarily the same before and

⁶⁶ Dretske, 1970, p.1010.

⁶⁷ Dretske, 1970, p.1016.

after the inference.⁶⁸ The operator did not penetrate because we were talking about different contexts. Dretske explains that an epistemic alternative will only penetrate if the latter context was part of the relevant alternatives of the original situation.⁶⁹

So exactly what are these relevant alternatives? The relevant alternatives could be explained as the reasonable ways you could be wrong. When I believe that I am watching a zebra in a zoo, maybe what I am looking at is really just a sign involving a picture of a zebra. Alternatively, maybe my son is just holding up his toy zebra and I just mistook it for a real one. These are all relevant alternatives given the context. What is *not* a relevant alternative is the possibility of a cleverly disguised donkey. The chances of this happening are incredibly slim. In everyday normal contexts, we cannot consider this as a relevant epistemic possibility.⁷⁰ The context changes if we are in a philosophy classroom and not the zoo. Suddenly, sceptical possibilities like the zebra being a cleverly disguised horse *are* relevant. So the standards required for knowledge change with context.

Dretske explains what makes an alternative relevant or irrelevant in his 1981 paper *The Pragmatic Dimension of Knowledge*. First of all, language can give us hints as to which alternatives are relevant.⁷¹ One's choice of words is such an example. If I said 'your neighbour found your joke funny' it would seem that I was making a comment about how funny your joke was. However if I said 'a man I didn't recognise found your joke funny', perhaps I was enquiring about *who* found your joke funny. Word stress can do the same thing. 'Your *neighbour* found your joke funny' has different connotations from 'your neighbour found your joke *funny*'. In the first instance, the relevant alternatives would be your doctor or your wife finding the joke funny for example. In the latter, the relevant alternatives would be your joke being lame or only mildly amusing. Your choice of words can also show *how* you came to believe something. 'I *heard from my hairdresser* that America has started peace talks' may imply that this information came from an unreliable source. If your original belief was that America wants to keep the war going, perhaps this is enough to make the possibility of peace talks irrelevant.

⁶⁸ Dretske, 1970, p.1007.

⁶⁹ Dretske, 1970, p.1023.

⁷⁰ Dretske, 1970, p.1022.

⁷¹ Dretske, 1981, p.373-4.

Apart from language, Dretske also talks about how the importance of a belief can influence your set of relevant alternatives.⁷² Say you believed that your tap-dancing shoes were in the closet. If you were just wondering if you could show off your dancing skills to your friends this weekend, the consequences of you being wrong are not that severe. Therefore we can have relatively low requirements for knowledge and have less relative alternatives. However, if you were preparing for the world tap-dancing championships tomorrow, the consequences would be much more serious. Therefore we should have higher standards of knowledge and have more relevant alternatives.

Finally Dretske asks whether an alternative can just be too remote to be relevant. He says not, since there are always cases where the seemingly unlikely situation just happens to be the truth. Having evidence for a belief being a possible alternative is not a necessary condition for being a relevant alternative, but it may be sufficient.⁷³ Importantly, Dretske concludes that what qualifies as a relevant alternative should not be dictated by what one regards as such, instead, it should be something within the objective reality.⁷⁴

So concludes the first Dretske's formulation of contextualism. This account has been influential. Philosophers such as David Lewis, Stewart Cohen and Keith DeRose have all contributed to contextualism afterwards. However, Dretske's form of contextualism has also been challenged by precisely these philosophers for different reasons.

Recall that Dretske's answer to the sceptic's paradox was to deny EPC. Denying the closure principle is not a popular solution amongst many philosophers. Stewart Cohen notes that the contextualist can respond to the paradox without denying EPC.⁷⁵ Here is the paradox again:

- 1) I know that I am looking at some zebras.

⁷² Dretske, 1981, p.374.

⁷³ Dretske, 1981, p.376.

⁷⁴ Dretske, 1981, p.377.

⁷⁵ Cohen, 1999, p.65-7.

- 2) If I don't know whether or not these 'zebras' are really just cleverly disguised donkeys, then I don't know that I am looking at some zebras.
- 3) I don't know whether or not these 'zebras' are really just cleverly disguised donkeys.

By definition, the contextualist believes that standards of knowledge changes with context. Now let's ask ourselves, are the standards of knowledge the same for 1 and 3? They're not. In 1, we are talking in a casual, everyday sense, and hence have low standards for knowledge. In 3, we are approaching sceptical problems as philosophers, so obviously our standards for knowledge would be higher. We can hence object to different propositions in different contexts. In normal situations, we will reject 3, whereas when we are doing epistemology, we will reject 1. Either way the closure principle is not threatened.

A bigger disagreement between Dretske and other philosophers is what governs the relevancy of an alternative. Before we go into these rules, it's important to note that according to Lewis' contextualism, S knows that p iff p fits all the possibilities not eliminated by S's evidence.⁷⁶ This has certain implications on the importance of the rules dictating which alternatives are relevant. Here are Lewis' conditions for relevancy:

1. Actuality

What actually happens can never be an irrelevant alternative. If S knows that p then p is true. This is not controversial.⁷⁷

2. Belief

This is another condition that fits well with the old JTB account. Whatever S believes will be a relevant alternative.⁷⁸ Note that this is different with Dretske's account. Dretske claimed that there is an objective truth as to what counts as a relevant

⁷⁶ Lewis, 1996, p.552.

⁷⁷ Lewis, 1996, p.553.

⁷⁸ Lewis, 1996, p.555.

alternative, so *S* believing that *p* should not be a necessary condition for *p* being relevant, according to him.

3. Resemblance

If an alternative *q* saliently resembles a relevant one *r*, then *q* is relevant too. Any alternative that saliently resembles actuality in terms of the evidence pointing towards it is also relevant. This is an important rule, since it is what (apparently) gets Lewis' contextualism out of Gettier problems. Take the Smith case again, what were the relevant alternatives? One relevant alternative was that the employer was lying coupled with the possibility that Smith did *not* have ten coins in his pocket. This is an alternative that saliently resembles his original belief, so it is relevant. However in this possibility, *p* (that the man who gets the job has ten coins in his pocket) does not hold. Therefore by Lewis' definition, Smith did not have genuine knowledge.⁷⁹

4. Reliability

Certain processes such as perception, personal testimony and memory are generally reliable processes. Therefore in general, they should play a role in an alternative's relevancy (eg. if I remember reading that the zookeeper likes to disguise horses as zebras, the possibility of the 'zebra' being a horse should be a relevant alternative).⁸⁰

5. Representation & Explanation⁸¹

These are two rules that Lewis only talks about briefly. We should believe that our samples are representational and our inferences to the best explanation are true. There are bound to be problems with these rules, the most logical explanation may not be true; the truth may seem random to us. There are more important rules however that have drawn more attention.

6. Conservatism

⁷⁹ Lewis, 1996, p.556.

⁸⁰ Lewis, 1996, p.558.

⁸¹ Lewis, 1996, p.559.

There are certain possibilities that are ignored because of society's conventions. Having a donkey disguised as a zebra for example, no matter how much it resembles actuality, just seems implausible, so we ignore it. We are allowed to ignore these types of alternatives.⁸²

7. Attention

According to Lewis, the relevant worlds are a subject to our conversational contexts.⁸³ For example, say we are discussing whether or not we knew that we were looking at a zebra. We start by chatting conversationally. At this point, we are using a normal, everyday conception of knowledge, so possibilities such as the zebra really being a donkey are not relevant. Suddenly, a philosopher joins us and starts making sceptical arguments. Suddenly, cleverly disguised donkeys become relevant alternatives.

The rule of attention is a controversial one. Richard Feldman (1999) has objected to this with a counterexample. Take any normal knowledge claim eg. that I have hands, when we become convinced by sceptical arguments, we now think that we don't know such things at all. Furthermore, when we look back to moments before we became sceptics, we don't think that we were right in both occasions. Rather, we think we were wrong before and now we are right.⁸⁴ Therefore, it really does not seem like we shifted contexts at all. Rather, one belief simply replaced another. Cohen agrees with Lewis on the rule of attention and objects with Feldman. His has two arguments. First of all contextualism can be construed as a type of error theory. It may not *seem* to us that we were crossing contexts, but according to the contextualists, we really were. Furthermore, we could take the contextualist claim as a normative rather than a descriptive one. We may not be crossing epistemic contexts when we switch from our everyday conversations to sceptical debates; but perhaps we should.⁸⁵ Feldman's phenomenological arguments can hence be answered.

⁸² Lewis, 1996, p.559.

⁸³ Lewis, 1996, p.559.

⁸⁴ Feldman, 1999, p.107.

⁸⁵ Cohen, 1999, p.81.

There are more pressing problems for contextualism – Gettier problems to be precise. David Lewis has argued that the rule of resemblance along with his definition of knowledge can together rule out Gettier problems. This is not the case and Cohen (2006) has an argument for it. Cohen of course is a contextualist himself, but he, like most contextualists, doesn't claim that he can solve the Gettier problems. In his 2006 paper titled *Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery*, he raises an objection to Lewis. Recall that the rule of resemblance requires the resemblance to be salient, but that, according to Cohen is what gives Lewis' account troubles. He asks us to imagine a special type of Gettier case: subject S is looking at a hill and he thinks he sees a sheep. However, that 'sheep' is really just a rock that looks like a sheep from a distance. Additionally, just out of pure luck, there is a sheep right behind the rock. Does S know that there is a sheep on the hill? No he does not. Lewis would say that his account can deal with this case because the possibility of the 'sheep' being a rock, and a sheep being behind that rock resembles S's belief. Therefore it is a relevant alternative that cannot be ignored. So far so good. Here's the trick, another person, call him 'A', is standing right next to S. A is wondering whether or not S knows that there is a sheep on the hill. Furthermore, A is not aware that S only sees a sheep-shaped rock, so the possibility that S is in a Gettier situation is not salient to him. Therefore, A can say that S knows that there's a sheep on the hill.⁸⁶ The contextualist therefore cannot deal with certain Gettier problems.

Ultimately, the truth of contextualism does not necessitate the falsity of the causal theory. Let's assume that the contextualists are right and we adopt the claim that standards of knowledge are context dependent. What then? What counts as justification (or some equivalent to it) according to this account? If we look past Lewis' relevant alternatives semantics, what we have is essentially an evidentialist account. However as we have seen, evidentialist accounts have no advantage over causal accounts. Can we not still apply causal notions to contextualism? Yes we can. Perhaps different contexts calls for different standards of knowledge, but these standards can all involve causal concepts. Given that a) the objections to the causal

⁸⁶ Cohen, 1998, p.297.

theory could be solved and that b) its competitors have even greater problems, the causal theory remains our best theory of knowledge.⁸⁷

2.3.2. Steiner's Objection

In response to Benacerraf's problem, Steiner chooses to object to the first premise and the causal theory of knowledge. Unlike the philosophers considered in the previous section who give an alternative theory of knowledge, Steiner objects on the grounds that there isn't a reasonable formulation of the causal theory that doesn't make any reference to causal abstract objects. He goes through different formulations of the causal theory and eventually settles with his conclusion. My objection to Steiner is that he made a crucial mistake in his second attempt at formulating the causal theory. I will go through all the formulations he considers. Of course, I have no problems with accepting a causal theory of knowledge that involves causal abstracta, but we don't need to make such a point against Steiner. For all of Steiner's attempts, let 'S' be someone who is capable of knowledge, let ' p ' be a proposition, let '[p]' be the fact that p , and finally let ' χ ' be a sentence.

(1) S cannot know that p unless [p] causes one's belief that p .

This makes references to facts, which are abstract objects.⁸⁸ For a fact to cause knowledge is hence to postulate a causal abstract object. Steiner therefore tries to formulate the causal theory without making references to abstract entities:

(2) S cannot know that χ is true, unless χ must be used in causal explanation that χ is true.

The problem with this formulation is that in its current form, it cannot undermine the force of mathematics with reference to abstract objects. Let p be the proposition that the axioms of number theory and or analysis are true; given that we believe that they are true, Steiner argues that there must be a theory that causally explains why we

⁸⁷ For more on contextualism, see Brendel and Jäger (2004).

⁸⁸ Steiner, 1973, p.59.

believe these axioms to be true. This theory will inevitably involve the axioms. We would then have a case where the axioms are used to causally explain their own truth. 2 is met and so we cannot object to mathematical Platonism through this version of the causal theory of knowledge.⁸⁹ Perhaps what we need is then a version that makes no references to causal explanations:

- (3) S cannot know anything involving F unless one's beliefs about F were caused by F (or F 's).

Steiner notes that this formulation is problematic because it has objects as its causal relata. This is unusual, we tend to think of causal relata as events.⁹¹ So we need to reformulate the causal theory as

- (4) S cannot know anything involving F unless one's beliefs about F were caused by an event (or events) which F (or F 's) participate(s) in.

The problem with this formulation is that we can have everyday cases of knowledge that do not meet this requirement. Steiner's example is one involving extinct species. Let's say a now extinct mammal created a footprint in the jungle. Many years later, the zoologists discovered this footprint and came to know many things about this animal because of it. Here we have knowledge about F 's but the latter events don't actually involve F 's.⁹² Steiner states that this is an example of an event causing a condition and that condition being partly responsible for another event, which in turn caused something else.⁹³ Our problem stems from the fact that the story is put in terms of conditions. Talk about conditions causing is talk about causal explanation, but if we wanted to translate this back to talk about causal explanations, we end up with the same structure as (2). This is problematic. This is where Steiner concludes that a nominalist construal of the causal theory is impossible.

Steiner makes a mistake in arguing against version 2. 2 does not specify what it means to be 'used in a causal explanation'. Any theory that involves the mathematical

⁸⁹ Steiner, 1973, p.60-1.

⁹¹ Steiner, 1973, p.61.

⁹² Steiner, 1973, p.62.

⁹³ Steiner, 1973, p.62.

axioms for example may use its *truth* to infer certain things, but what it does not do is show the axiom itself interacting with us causally. Surely, it is the latter of causation that we are interested in; that is the crux of Benacerraf's epistemic problem.⁹⁴ As long as the causal theorist provides us with a version of the causal account that specifies this type of causal interaction, they can avoid Steiner's objection. We will add this condition to KTK.

Let ' p ' be a propositions let ' $[p]$ ' be the fact that p , let Bp be S 's belief that p , and finally let Op be all the objects involved in p (identified through the Quinean method).

(KTK#) S knows that p iff Bp was k-causally connected to $[p]$, and that all the components of Op are capable of causally interacting with us.

Of course, we could also just choose to adopt (1) or (2). AOCO theorists such as myself have no problem accepting a theory like that. Either way, Steiner's objections therefore do not pose any real threat to the causal theory of knowledge.

2.3.3 Objections with Positive Solutions

We have examined alternative accounts to the causal theory of knowledge. Other philosophers such as Wright, Zalta, Linsky and Balaguer deny the traditional causal theory without proposing other theories of knowledge. Instead, they give positive accounts, showing precisely how we can come to know about acausal abstract objects. These accounts are very impressive, but ultimately are not accounts we should adopt.

⁹⁴ There will be more on this line of thinking in 2.4.

2.3.3.a. Wright's Solution

Crispin Wright actually gives us a way of learning about abstract objects without needing them to be efficacious. Wright objects to the causal theory of knowledge, or at least, he objects to the standard formulations of the theory of knowledge. For the causal theory to accommodate his solution to Benacerraf, substantial changes need to be made. Recall that the standard theory of knowledge requires the subject to causally interact with the object of his knowledge. Wright suggests weakening this condition. He claims that for example, it may be enough for us to interact with the causally necessary states of affairs sufficient to the proposition in question.⁹⁵

Wright's solution once again makes use of equivalence relations. Recall that Wright himself wants to define abstract objects as equivalence relations, so this is no surprise. This is his epistemic solution: we can know about abstract objects by interacting with the concrete components that the entity is equivalent with. Wright gives a famous example involving the universal of having a certain direction. Two lines, X and Y, pointing to the same direction is logically equivalent to X and Y being parallel. With this in mind, imagine someone who has mastered the use of sentences involving directions and parallel lines. Benacerraf's problem troubles her – exactly how did she come to learn about directions? She then realises that she came to this understanding through the relevant equivalence. We can hence learn about abstract entities through equivalence relations, gaining knowledge about them through concrete entities which partake in these relations.⁹⁶ If Wright is correct, we would then have a way of knowing about causal entities without causally interacting with them.

Wright's argument is one that nominalists don't like; Hatry Field particularly argues in a direct response that the abstract object adds no epistemic value to this story. In other words, I could know that two lines are parallel without abstract objects even existing.⁹⁷ Field himself uses this point to argue against Wright's Platonism, which is unfortunately out of place within our discussion. Once again, what we want are

⁹⁵ Wright, 1983, p.85.

⁹⁶ Wright, 1983, p.87.

⁹⁷ Field, 1984, p.641.

arguments for or against the knowledge of abstract objects. Questions such as these about the properties of abstract entities should precede questions about their existence. Perhaps we can adjust Field's argument for our purposes. Do these equivalence relations really give us knowledge about abstract objects at all, or simply of concreta? Wright's solution hence opens the door to as many questions as it could answer. On one hand, it seems to allow us to accept theories involving abstracta; on the other hand, we're not even sure if this theory of abstract objects tells us anything about abstracta at all. There is also the familiar problem of whether all abstract objects can be reduced to equivalence relations. If not, then Wright's solution can only help us so much.

2.3.3.b. Full-Blooded and Principled Platonism

Mark Balaguer, Bernard Linsky and Edward Zalta have also all tried to give realist accounts that get around Benacerraf's problem. Balaguer does this through his theory of 'full-blooded Platonism'. Linsky and Zalta have a similar account with what they call 'principled Platonism'. Both of these accounts try to evade Benacerraf's problem by a) denying the causal theory of knowledge, and b) showing that we can have non-causal cases of knowledge by postulating a plenitude of abstract objects. Although the existing objections to these theories can be solved, I will propose further reasons for rejecting FBP and PP.

Balaguer's FBP is a model that states that "all the mathematical objects which possibly *could* exist actually *do* exist, or perhaps that there exist mathematical objects of all kinds" (Balaguer, 1995, p.304). This model was created to get around Benacerraf's epistemic challenge. From FBP, any mathematical proposition that is consistent therefore successfully picks out part of the abstract realm. Unlike Wright, Balaguer outright denies the causal theory of knowledge.⁹⁸ His model of realism allows us to have knowledge about abstract entities without there being a causal connection between ourselves and the object.

⁹⁸ Balaguer, 1995, p.303.

Note that this view naturally generates an infinite amount of mathematical objects, many of which we haven't even thought of. There may be different sets of natural numbers for example because they differ according to different mathematical axioms. Whenever we refer to a mathematical object, say the number 1, we will inevitably refer to a myriad of different abstract entities that are all the number 1, but with slightly different properties. Balaguer does not view this as a problematic consequence of his theory.⁹⁹ Furthermore, Balaguer states that if we accept a correspondence theory of truth, then any consistent mathematical proposition is true! This seems to lead to contradictions, because different mathematical theories may contradict each other. Let 'C' be a mathematical theory. How can both C and not C be true? Balaguer finds this unproblematic. These are different theories referring to different regions of the mathematical realm, so there is no real contradiction (Colin Cheyne has more to say about this and I shall discuss this later).

Cheyne argues that FBP displays an odd relationship between the numerical terms and the objects they are supposed to reference. He makes the distinction between intentional and extensional assignments. In extensional assignments, we can (even arbitrarily) pick out a group of objects and give them a term. Cheyne uses the example of arbitrarily picking certain sheep from a mob and calling them 'slithy' sheep. On the other hand, an example of intentional assignment would be picking out the merino sheep from the mob. In intentional assignments, we already understand the term, and that meaning then picks out the set. With extensional assignments, the set picks out the meaning of the term. The worry with FBP is that the assignments of number terms to abstract objects here seem purely extensional, when Balaguer seems to think that the assignment should be intentional in the rest of the paper.¹⁰² Cheyne claims furthermore that we *would* be able to intentionally assign numbers to their numerical terms if we knew what it meant to have number-like properties, but we don't. Numbers exist outside of space/time, so their properties are out of reach. He makes an analogy to how colour terms would be meaningless if the human race never developed sight.¹⁰³ Since we cannot come to contact with number properties, we can't

⁹⁹ Balaguer, 1995, p.315.

¹⁰² Cheyne, 2001, p.179.

¹⁰³ Cheyne, 2001, p.179.

understand them, and hence, we cannot pick out numbers through intentional assignments.

FBP has the tools to answer this challenge. Do we really not know anything about the properties of numbers? Take the number 1, we know that it is the smallest positive integer; we know that it is half of 2; we know that we can add it to 2 to make 3... it seems that we know an awful lot about something we can never know about. The reason we know about the properties of numbers is that it is part of a consistent system of arithmetic. For our system to be consistent, numbers simply have to be a certain way. Given the conditions surrounding the consistency of arithmetic, we come to know properties of numbers, even if we have no causal connections with the actual entities. Since we can know these things about numbers, we can pick them out through intentional assignment. Cheyne's challenge can be resolved.

Cheyne has another argument which is also based on the worry that we cannot distinguish between different mathematical objects. He starts by showing how a Platonist like Balaguer may formulate natural-number arithmetic:

“(P1) Zero is a natural number.

(P2) Every natural number has a successor, which is also a natural number.

(P3) Zero is not the successor of any natural number.

(P4) Different natural numbers have different successors.

(P5) [The induction postulate]” (Cheyne, 2001, p.175)”.

Now compare this formulation to an analogous but intentionally ridiculous one that Cheyne created:

“(D1) Genie is a djinn.

(D2) Every djinn has a meta, which is also a djinn.

(D3) Genie is not the meta of any djinn.

(D4) Different djinns have different metas.

(D5) [A suitable induction postulate]” (Cheyne, 2001, p.184).

‘Djinn arithmetic’ is equally consistent, as it was created by mirroring the structure of its predecessor. According to FBP, djinn arithmetic exists, and according to a correspondence theory of knowledge it is true! Of course, this is not a consequence we would want to accept. Djinn arithmetic is a ridiculous creation, miles away from any respectable formulation of arithmetic. Due to the nature of FBP we get the unpleasant consequence of any consistent model being mathematical truth. This is a great objection by Cheyne. The FBP-ist needs some sort of clause that avoids these senseless models from being vacuously true.

On the other hand, Linsky and Zalta’s principled Platonism has all the advantages of FBP but without this problem. PP introduces a new concept that sets it apart from other forms of realism – the concept of encoding. Most forms of realism explain the properties of an abstract object through the concept of exemplification; PP distinguishes between exemplification and encoding. The most important principles surrounding the concept of encoding are:

- 1) Every condition of properties has an abstract object that encodes precisely these properties.
- 2) If an abstract object encodes a property, it does so necessarily.
- 3) Two abstract entities are identical iff they encode the same properties.¹⁰⁴

The properties that are encoded in an abstract object are precisely the ones that make it unique. For example, the property of redness is encoded in the universal of red. What is exemplified in the universal are properties such as existing outside of space and time ie. properties that the abstract object has but that plays no part in distinguishing itself from other abstract objects of the same type. The first principle makes PP a model that postulates an abundance of abstract objects, much like FBP. From 1, we can say that for every condition of properties, there is an abstract object that exists which encodes those properties.¹⁰⁵ This allows PP knowledge of abstract objects without needing any causal interactions, much like FBP, as any knowledge over properties would inevitably be knowledge about abstract objects. Like Balaguer,

¹⁰⁴ Linsky & Zalta, 1995, p.536.

¹⁰⁵ Linsky & Zalta, 1995, p.537.

Linsky and Zalta explicitly deny the causal theory of knowledge, because their account shows how abstract knowledge can be acausal.¹⁰⁶

Cheyne argues that PP gets around the djinn arithmetic challenge because of the differences between PP and FBP. Where FBP states that any consistent mathematical proposition exists, PP only states that there is an abstract object that encodes every condition. This is an important distinction because ‘being the meta of a djinn’ for example is a meaningless statement. What would it even mean for an object to have this property? How can anything exemplify a meaningless attribute? Therefore, despite its consistency, there really isn’t any such property. We therefore do not have to worry about any old property being true or playing roles in fields that they are not related to¹⁰⁷ and PP does not have any problems regarding this objection.

Given these strengths of PP, shall we accept it as our solution to Benacerraf’s problem? Although this would work with certain abstract objects such as mathematical entities, it is less appealing when applied to others. First of all, not all abstract objects are properties. Sets, fictional characters and events for example are all other types of abstract objects. Even if we assume that these entities can be convincingly construed as properties, we still have problems with contingent abstract objects. Let us take a creationist view of fictional characters and apply PP to it. Even before J.K. Rowling created Harry Potter, the character existed according to PP, because there is an abstract object that encodes any property that could exist. This seems odd, if Harry Potter always existed then what happens when the author creates it? Remember that we are assuming a creationist model here, so it seems incredibly odd to say that Harry always existed. Notice that although PP has no problems with the djinn arithmetic case, it still struggles with this one. In order for PP to be applicable to cases such as fictional characters, we must insist that any abstract object that exists must exist as necessary properties. Not only would we have to insist that these abstract objects are properties, we would have to give up on all models which view any abstract objects as contingent. These are both good reasons for us to reject PP since not all abstract objects fit into this mould. These models are really no help to the abstract realist.

¹⁰⁶ Linsky & Zalta, 1995, p.547.

¹⁰⁷ Cheye, 2001, p.191.

2.4 Deny the Validity of Benacerraf's Argument

Max Cresswell argues against the validity of Benacerraf's objection. He claims that there is a fallacy of equivocation.

“What I will try to show in the present article is that premise (2) relies on an equivocation in ‘part of the causal order’” (Cresswell, 2010, p.250).

Benacerraf describes ‘being part of the causal order’ in two different ways. When he states that only things which are part of the causal order can lead to knowledge, he is talking about general causal *interaction*. We know about things because we can interact with them in certain ways. However, when Benacerraf talks about abstract objects being outside of the causal order, he is talking strictly about *existence*. When philosophers say that numbers are outside of the causal order, all they are really saying is that it exists necessarily. It is still possible for such objects to interact with us causally. Cresswell makes this point using propositions as an example. For someone to know something is for person *a* and proposition *p* to have some form of knowledge relation. Statements of knowledge then come into causal explanations eg. ‘If I had know that you were coming I’d have baked a cake’.¹⁰⁸ The *truth* of the proposition is contingent and can be a subject of interaction within the causal order. However, the *existence* of the proposition is necessary. Traditional realists about propositions would say that the proposition has always existed and outside of space/time.

“If I cause an accident it is not normally my *existence* which caused the accident – it is something I did, something which is true of me. If so, questions about what kind of existence an entity has need not tell us anything about what other kinds of things can be true of that entity” (Cresswell, 2010, p.252).

¹⁰⁸ Cresswell, 2001, p.251.

So truth and existence are separate concepts, the former is contingent and the latter is necessary. The truth of a proposition can be part of causal interactions. Although Cresswell does not make this point explicitly, it is strongly suggested that abstract objects such as propositions are therefore part of the causal order in terms of its *interactions*, but they are not part of the causal order in the sense that they do not come in and out of *existence*. Benacerraf did not realise this distinction.

If Cresswell is right, this is a deathblow to Benacerraf. Benacerraf's argument should really read like this

- 1#) We can only have knowledge of things that we can interact with causally.
- 2#) Abstract objects exist necessarily, they do not come in and out of existence; therefore the events of these objects coming into existence are not within the causal order.
- 3#) Therefore, we cannot have knowledge of abstract objects.

This is obviously invalid. Once we see that Benacerraf's original argument contained the fallacy of equivocation, the premises no longer entails the conclusion. We need a further premise such as 'we cannot causally interact with entities that do not come in and out of existence' for the argument to be structurally valid.

Cresswell states that if Benacerraf's argument has any force at all, it requires an underlying metaphysical premise – what Benacerraf himself calls a PMP or 'Princess Margaret Premise'. For example, Benacerraf may claim that we can analyse everything without the aid of abstract objects, so we should not postulate them. Alternatively, perhaps we should only postulate the existence of things we can perceive. However, Benacerraf needs such an argument and this is not present in his epistemic objection; so his argument, as it currently stands, does no real work.¹⁰⁹

This looks grim for Benacerraf, but perhaps things are not as bad as it looks. Let us assume that he disagrees with the objection. It's not that there is an equivocation on

¹⁰⁹ Cresswell, 2010, p.250-258.

‘being part of the causal order’, it’s that Benacerraf (along with other metaphysicians) believe abstract objects are acausal on every front. According to this view, propositions cannot partake in causal interactions. This certainly seems to fit with the traditional way of viewing abstract objects. According to the defenders of Benacerraf, this is what the argument *should* be like:

- 1#) We can only have knowledge of things that we can interact with causally.
- 2#) Abstract objects both exist necessarily and cannot be interacted with causally.
- 3#) Therefore, we cannot have knowledge of abstract objects.

Supporters of Cresswell will no doubt object to this, “Isn’t it obvious that the truths of propositions are within the causal order?” they would say. Perhaps they’re right, but that doesn’t show that there is an equivocation, that shows that abstract objects are indeed part of the causal order. I therefore conclude that denying the structure of Benacerraf’s problem is not our best solution.

2.5. Deny the second premise

We have tried different answers to Benacerraf’s epistemic problem but none of them seem appealing. Cresswell attacks the validity of the argument, but that can be easily avoided. Denying the causal condition also seems like a fruitless endeavour. The causal theory of knowledge stands as the strong favourite among theories of its kind and its status seems hard to shake. Wright, Linsky and Zalta all offer impressive arguments but they all have their shortcomings.

The realist seems to be caught in a compromising position; on the one hand she wants to continue postulating the existence of abstract objects for whatever reasons, on the other hand, she realises postulating acausal entities would lead to major epistemic issues. The realist cannot both have her cake and eat it, which in effect makes anti-realism a more attractive position, on the face of it. However, this is not the case.

There is one option the realist has left. I propose that we should all adopt AOCO. Given my proposal, we can accept Benacerraf's argument without sacrificing the existence of abstract entities. The causal model of abstract objects thus arms the realist with a simple and effective solution to this epistemic problem.

In chapter 1, I argued that abstract objects are part of causal effects. Now we should accept that they can also be part of causes. I propose that abstract objects can do this at least in terms of giving us knowledge. We at least have some reasons to allow abstract entities into the causal order given that there are philosophers who have allowed for their abstract objects to be efficacious.

Just in case the reader is still not convinced I will strengthen my argument by listing philosophers who have postulated abstract objects as part of causes. We will start with Gödel's mathematical axioms. Axioms are the building blocks of mathematics. There are different axiom theories which philosophers and mathematicians propose. These axioms are supposed to provide a base in which all mathematical propositions and truths could be derived from. If one fully accepts the first sentence of this paragraph, he/she is committed to the existence of axioms through Quine's methodology (eg. Mx where x is an object and $Mx = 'x$ is a building block of mathematics').

Kurt Gödel claimed that these axioms not only exist as abstract objects, they can be perceived by us through intuition.

“He compares the axioms of logic and mathematics with the laws of nature and logical evidence with sense perception, so that the axioms need not necessarily be evident in themselves, but rather their justification lies (exactly as in physics) in the fact that they make it possible for these “sense perceptions” to be deduced; which of course would not exclude that they also have a kind of intrinsic plausibility similar to that of physics” (Gödel, 1951, p.127).

In other words the fact that we have mathematical knowledge means that we should have a method in gaining that knowledge. This gives us evidence in believing both the existence and the causal nature of mathematical axioms. We come to gain knowledge

of these axioms through mathematical intuition. This intuition is best explained through a comparison with sensory perception. Both of these faculties provide us with “elementary indisputable evidence”.¹¹⁰ This helps explain why many people can have diverse opinions on many topics, yet a unified account of mathematics.

Furthermore, John Bigelow in his book *The Reality of Numbers* (1988) has argued for a physicalist conception of numbers. He argues that numbers are universals. According to Bigelow, universals are defined as things that can be wholly present in multiple places at the same time. The universal of beauty for example is wholly present in a flower in New Zealand and the Mona Lisa in the Louvre. Universals have a problem with location. So where is the beauty in the Mona Lisa? Is it in the corner? Is it in the background perhaps? Despite these problems, Bigelow argues that universals nonetheless still have locations – they are located where particulars exemplify these properties, but we cannot say much more than that.¹¹¹

Given that numbers are universals thus construed, they by definition have spatiotemporal properties. How else could a universal be at different places in the same time? Since numbers can exist within space/time, they can be causal in a way traditional Platonic numbers cannot. It is this very causal nature of universals and numbers which made Bigelow hesitate in calling them ‘abstract’.¹¹²

Though Bigelow did not want to call his numbers abstract, we should ask the question: why not? Is it only because he sees them as causal entities? A more natural option is simply to view these entities as causal abstract objects. Given all the positive consequences of AOCO, this certainly has its values.

Let’s not forget about Linda Wetzel’s linguistic types either. It’s not only natural for us to be able to be causally influenced by words for example; it is equally natural for us to causally influence them. In her own words:

¹¹⁰ Gödel, 1951, p.128.

¹¹¹ Bigelow, 1988, p.22.

¹¹² Bigelow, 1988, p.26.

“They may be specified one at a time, as when we refer to the word ‘eleemosynary’, or all at once, as when a linguist specifies the sentences of English by means of a set of recursive rules. So there really is no mystery about how spatiotemporal creatures like ourselves come to know about these abstract linguistic objects, what properties they have, and so on” (Wetzel, 2009, p.123).

In conclusion we should accept that abstract objects are part of the causal order, as parts of both causes and effects. AOCO is our best solution to Benacerraf’s problem, a problem which has been historically problematic for the realists. With it out of the way, we can have a healthier debate between abstract realism and nominalism. Causal abstract objects have also been proposed by many philosophers in the past, so there is undeniably some support for it. We should accept the causal nature of abstract entities as soon as possible.

Chapter 3 - Metaphysics

There is a paradigm within philosophy that abstract objects are not part of the causal order, meaning that they cannot be part of causes nor effects. A big reason is that many philosophers define abstract objects that way, but I've tried to show that there are reasons to think otherwise. We now need to ask ourselves if this assumption holds water. The aim of this chapter is to answer any metaphysical objections to AOCO: the thesis that abstract objects are part of the causal order.

In 3.1, I will consider four different causal theories and ask if they are consistent with AOCO. In the remaining sections, I will look for more general objections to AOCO. In section 3.2, I will consider direct objections to AOCO. However, since publications on this question are few and far between, I will draw arguments from another debate – the conversation surrounding whether abstract objects can make a difference (a broader claim than mine). We will see if any of these arguments can be transitioned to our debate in 3.3. In 3.4, we will investigate a problem regarding a special subset of abstract objects – ones that are necessary – and discuss whether these are problematic. Finally we will examine broad objections from nominalists and traditional Platonists in 3.5. Ultimately, none of these arguments should stop us from accepting AOCO.

3.1. Causal Theories

There are many causal theories that have been proposed. My aim in this part is to find whether any of these theories have problems with AOCO. Most of these theories have no problem with my proposal and only a minority – ones that aren't particularly popular – are incompatible. I will analyse these causal theories accordingly.

3.1.1. Regularity Theory

Traditional regularity theories originate from Hume. The basic idea of regularity theories is that causation involves a certain type of event necessarily following another type of event. In Hume's own words:

“We may define a cause to be an object followed by another, and where all the objects, similar to the first, are followed by similar to the second.”
(Hume, 1748/1975, p.76).

When Hume talks about necessity, he is talking about *nomic necessity*.¹¹³ Something is nomically necessary iff it is necessary given by our natural laws. Hence regularity theories tend to involve conditions involving rules of nature.

Let ‘*c*’ be the name of an event *c*, let ‘*e*’ be the name of an event *e*, let ‘*L*’ be some nonempty set of true law propositions, and finally let ‘*F*’ be some relevant set of background propositions about matters of particular fact.

c causes *e* iff

- 1) *c* occurs,
- 2) *e* occurs,
- 3) *c* is an event of type *C*
- 4) *e* is an event of type *E*
- 5) *L* and *F* jointly imply that whenever events of type *C* occur, events of type *E* follow,
- 6) *L* and *F* jointly do not imply that an event of type *E* occurs, and
- 7) *F* alone does not imply that whenever an event of type *C* occurs, an event of type *E* follows.

¹¹³ Hume, 1748/1975, p.62.

This is a standard formulation of a regularity theory. However, Hume claims that a spatiotemporal contiguity condition should also be an essential component.¹¹⁴ Firstly, causes and effects have to be spatially contiguous (i.e. next to each other) and this can be seen through experience. Even when we think events at a distance are causally connected, this can be explained through causal chains. Contiguity is hence essential to causation. Causes must also occur before the effect. This is important for Hume, otherwise we could not distinguish the cause from the effect. Therefore Hume's regularity theory requires an eighth condition:

8) *c* and *e* are spatiotemporally contiguous.

A strength of the regularity theory is that it has no problems with redundant causation unlike the counterfactual theory of causation for example. Regularity theories allow for multiple causes leading to the same effect. An analysis of what causes someone to drink a cold beverage for example would have multiple sufficient causes such as being thirsty, being in the heat, wanting a drink to accompany lunch etc., and each of these conditions can necessitate the effect. Regularity theories can hence accommodate cases with more than one possible cause.

An event is a cause of *e* if and only if it is one of the sufficient conditions. Any *c* that necessarily leads to *e* is a sufficient condition for *e*. This is a strength not all causal theories have and it gives us more reason to respect the regularity theory. It also has an intuitive appeal. If one thing caused another, the latter had to happen. Outside of the philosophy classroom, this seems like a decent way of describing causation.

Does the regularity theory have a problem with efficacious abstract objects? Surprisingly, even with the contiguity condition, we can have certain types of abstract objects being efficacious. Certain tropes for example should be allowed. Let's say that 'the weakness of the bridge made the travellers fall'. We can certainly imagine cases where given the weakness of the bridge (a trope), the travellers *had* to fall. Other abstract objects however will be rejected as causal entities. One reason for this is the contiguity condition. It is hard to imagine how linguistic types viewed as universals

¹¹⁴ Hume, 1739/1978, p.73-8.

for example can ever satisfy the contiguity condition. Linguistic types have no spatial properties, so by nature, they can't be contiguous.

However does the regularity theory really need these spatiotemporal limitations? Hume's own arguments on this matter are fairly weak. He claims that spatial contiguity is necessary in causation because all our everyday causal cases follow that pattern. This argument is not very persuasive, as frequency should not imply necessity. In other words, Hume has made the observation that all our every day causal stories involve objects that are next to each other. However it is a big jump from this observation to the claim that spatiotemporal contiguity is necessary within causation. For example, can we not imagine a possible world in which I can move objects that are far away from me, without any long causal links existing within the space between us? This is possible, even if can only occur in a world that is distant to us. A good theory of causation should involve these possibilities, as these possibilities do not clash with the concept of causation.

Unfortunately, rejecting the contiguity condition does not help us much. There is an additional hurdle for philosophers who accept AOCO – the condition of nomic necessity. Take fictional characters for example, can they nomically necessitate anything? This seems unlikely. If we say that 'the fact that Harry Potter went to Hogwarts caused Suzy to buy a wizard hat', are we saying that the latter had to happen given the laws of nature? No we are not. In fact it is very difficult to imagine a situation where fictional characters would nomically necessitate anything. We can say the same about linguistic entities. Therefore regularity theories are compatible with a few causal abstract objects, but not all of them. This could be worrying for anyone who accepts AOCO but are anti-realists about tropes for example.

A problem with the regularity theory is that not all examples of causation involve the effect being nomically necessitated. The fact that I forgot to bring an umbrella may cause me to get soaked by the rain, but that does not mean that me being soaked was necessitated by the fact that I didn't bring an umbrella. John Mackie has an adjusted version of the regularity theory that deals with this.

3.1.2. INUS Condition

There is a close cousin to the standard regularity theory which still tries to explain causation in terms of necessary or sufficient conditions. This theory includes the concept of INUS conditions, put forth by John Mackie.

An INUS condition is an *insufficient but necessary part of an unnecessary but sufficient condition*. For any given effect, there can be many unnecessary but minimally sufficient conditions. In the case of a forest fire, a fire could be caused by a match being lighted and thrown on the ground, but a fire could also be started by a heat wave; or a teenager with a flamethrower; and so on. All of these minimally sufficient but unnecessary conditions can be combined to form one necessary and sufficient condition such that event *e* occurs iff

- 1) *d* occurs and *f* occurs, or
- 2) *g* occurs, *h* occurs and *i* occurs, or
- 3) *j* occurs, *k* does not occur, *l* occurs etc.

where '*d*', '*f*', '*g*' ... etc. are all individual events.

Take any one of these sufficient but unnecessary conditions, *d* and *f* both occurring for example, *d* occurring here is an insufficient but necessary part of both *d* and *f* occurring. *d* occurring here is an INUS condition of *e*.

In the formula above, we can obviously call the entire antecedent a cause of *e*. However, most of what we call 'causes' are not complex disjuncts of this sort. When we say that something's a cause, we often mean that it is an INUS condition. In our example involving the forest fire, the lighting of the match was a cause of the fire. This was obviously not a necessary and sufficient condition of the forest catching fire, but it was an INUS condition. Note that Mackie does not think that these cases exhaust are the only examples of causation,¹¹⁵ but since the term 'causation' can be

¹¹⁵ Mackie, 1965, p.247.

ambiguous, we will only focus here on the type of causes that Mackie is concerned with.

Continuing with our original terminology, let ' c ' be the name of an event c , let ' e ' be the name of an event ' e '; furthermore, let ' X ' stand for the unnecessary but sufficient condition that contains c (eg. if $c = \alpha$, $X = [\text{both } d \text{ and } f \text{ occurring}]$), and let ' Y ' be the combination of all the other unnecessary but sufficient conditions ($[g \text{ occurring, } h \text{ occurring, and } i \text{ occurring}] \& [j \text{ occurring, } k \text{ not occurring, } l \text{ occurring}] \dots \text{ etc.}$). The INUS theory can be formulised as follows:

c causes e iff

- 1) c is at least an INUS condition of e ,
- 2) all the components of X (including c) occurred, and
- 3) every disjunct in Y that does not involve c was absent on this occasion.¹¹⁶

In our example, we can therefore say that the lighting of the match was a cause of the fire.

An advantage of Mackie's theory is that he allows for causes that weren't universally necessary for the effect. Note that the lighting of the match was not a necessary or sufficient condition for the fire starting. The fire would not have started if the match was lit but was *not* thrown into the ground, yet given that it *was* thrown on the ground, we would still want to call the lighting of the match a cause. The standard regularity theory cannot account for these types of causes as it requires a cause to necessitate the effect. This gives Mackie's theory some advantages.

Would this theory of causation be compatible with abstract objects being part of the causal order? Yes it would be, at least for certain types of abstract objects. The fact that Harry Potter is a wizard, the prevalence of pop culture stores in town, Suzy's love for Harry Potter etc. form a minimally sufficient condition for her to buy a wizard hat. It follows that she would not have bought this wizard hat if Harry was not a wizard, so

¹¹⁶ Mackie, 1965, p.247.

the fact that Harry is a wizard is an INUS condition. This fact about a fictional character hence counts as a cause of Suzy buying the wizard hat (hopefully the reader will grant me that facts about fictional characters are also abstract). We can imagine other cases of the sort with other abstract objects.

3.1.3. Energy Transfer View

There is a theory of causation that is inconsistent with AOCO more directly than the regularity theory. David Fair's analysis of causation states that causal connections can be explained by a flow of energy or momentum from the cause to the effect. This is a theory of causation that could potentially directly contradict AOCO.

“I will argue that physical science *has* discovered the nature of the causal relation for a large class of cases. As a first approximation, it is a physically-specifiable relation of energy-momentum flow from the objects comprising cause to those comprising effect. The causal relation of ordinary language seems likely to be *reducible* to that relation” (Fair, 1979, p.220).

Although other causal theorists like to specify their causal relata as propositions or events, Fair is not so specific. Instead, he introduces the distinction between *A-objects* and *B-objects*. *A-objects* are the combination of events, actions, mental states and facts that comprise any cause A, *B-objects* are the same but for any effect B.¹¹⁷ For the sake of semantic consistency, I will use the terms *c-objects* and *e-objects* instead, with *c* and *e* being whatever that manifest these objects respectively. Here is an example, take the proposition ‘John shooting the gun caused Bob to die’; using our terminology, *c* is John shooting the gun and *e* is Bob dying, that much is clear. Now furthermore, John's mental states, the pulling of the trigger, the bullet itself etc. are the *c-objects*. Likewise, Bob's heart stopping, his loss of consciousness, the fact that

¹¹⁷ Fair, 1979, p.236.

he is no longer alive etc. are the e-objects. We can say that c-objects are the different components of *c*, and the same with e-objects and *e*.

According to Fair's energy transfer view,

c causes *e* iff there are physical redescriptions of *c* and *e* such that transfers in energy or momentum occur from the c-objects to the e-objects.

It is important to note that Fair himself never claimed to have found a theory that covers all cases of causation. The energy transfer view however does cover a large proportion of causal stories.¹¹⁸ This theory will appeal to those who are especially scientifically minded. Many would argue that causation involving a chain of energy or momentum is the only form of causation. What better way to explain physical causation than through physical means?

Can Fair's theory of causation account for abstract objects being efficacious? Abstract objects obviously contain no mass, therefore any talk about energy flow or momentum would not apply to it. However, since Fair never claims to cover all cases of causation with his energy transfer view, it is at least possible for Fair's view to coexist with AOCO. This is a surprising result, but an encouraging one.

3.1.4. Counterfactual Theory

The counterfactual theory is by far the most prominent analysis of causation today. This theory is naturally built around counterfactuals. The motivation behind this theory is intuitive – if one event caused another, perhaps the second event would not have occurred if the first event never happened.¹¹⁹ This helps us explain the force of causation.

¹¹⁸ Fair, 1979, p.220.

¹¹⁹ Lewis, 1973, p.562-3.

Let ‘*c*’ be the name of an event *c* and ‘*e*’ be the name of an event *e*. According to the counterfactual theory of causation, we can say that

c causes *e* iff

- 1) Both *c* and *e* occur; and
- 2) If *c* had not occurred, *e* would not have occurred either.

David Lewis and Robert Stalnaker use possible worlds to explain counterfactuals. There are an infinite amount of possible worlds, such that any possibility that our universe could be or could have been is exemplified. However, only one world is actual – our world, though this does not mean that our world is metaphysically unique. We just happen to be in this world, much like how I happen to be in my room, which does not grant my room any sort of special status. All other worlds are ranked in terms of closeness i.e. some worlds are more like our world than others.¹²⁰ However, Stalnaker and Lewis account do have important differences. Stalnaker believes that there is such a thing as the *closest* possible world,¹²¹ whereas Lewis doesn’t. He just believes that worlds are loosely ranked in terms of closeness,¹²² so the concept of the one closest possible world is not meaningful to him. This leads to different treatment for 2 in CTC.

Let ‘@’ be the actual world, Lewis states that

2L) there is a world in which both *c* & *e* do not occur that is closer to @ than any world in which *c* does not occur and *e* does.¹²³

Whereas given Stalnaker’s treatment of possible worlds,

2S) in the closest world to @ in which *c* does not occur, *e* does not occur either.

¹²⁰ Lewis (1973, 1986), Stalnaker, (1968).

¹²¹ This is strongly suggested in Stalnaker, 1968, p.102-5.

¹²² Lewis, 1973, p.560.

¹²³ Lewis, 1973, p.560.

Therefore it's *possible* for *c* not to cause *e*, but that possibility is one that is far away and alien to our situation.

There are plenty advantages to a counterfactual theory of causation. First of all, unlike the regularity theory, effects do not need to be nomically necessitated. It has an advantage over Fair's energy transfer view in the fact that not all causal stories involve strongly physical components. For example, the anger of a nation caused the civilians to revolt, yet there is no physical energy flow between the civilian's anger and the revolution. It is also a methodologically simpler theory than Mackie's theory; this makes the counterfactual theory an appealing one.

Let us take a look at the objections to the counterfactual theory of causation. One of the most popular objections is the problem of redundant causation. There are two types of cases we have to deal with – causal overdetermination and pre-emption. Let us start with pre-emption. These are cases where an event *c* seems to cause event *e*, but had *c* not occurred, another event *d* would have caused *e* to occur later. For example, both Billy and Suzy throw rocks at a bottle. Suzy's rock hits the bottle first and breaks it, but if Suzy's rock hadn't hit the bottle, Billy's rock would have broken it nonetheless. This is a problematic case with the counterfactual theory of causation because Lewis and his supporters would want to say that Suzy caused the bottle to break, yet in the closest possible worlds that Suzy does not throw a rock, the window still breaks. The counterfactual test hence fails.¹²⁴

Lewis solves this problem by arguing that Suzy breaking the window and Billy breaking the window are two distinct events.¹²⁵ Since redundant causation only occurs when the same event has multiple causes, cases of pre-emption do not qualify. The Billy and Sally story for example is only an example of two separate causal stories.

The problem of causal overdetermination is trickier. These cases involve two events *c* and *d* seemingly causing an effect *e* simultaneously.

¹²⁴ McDermott, 1995, p.528.

¹²⁵ Lewis, 1986b, p.195.

Two electrons, C1 and C2 fire simultaneously and stimulate neuron N at the same time. Did C1 cause stimulation in N? Well, had C1 not occurred, C2 would have stimulated N nonetheless, so the counterfactual test fails. What makes cases of causal overdetermination different from pre-emption is that the counterfactual theorist cannot argue that C1 and C2 cause separate effects. Therefore, these cases require another type of solution. Martin Bunzl argues that we should be suspicious to whether or not there are genuine cases of overdetermination.¹²⁶ With this current example, C1 and C2 together cause a distinct effect, one that is different from the effect that C1 or C2 alone would have produced.¹²⁷ Therefore had C1 not occurred, the neuron would be stimulated by a completely different event. We can do something similar with all the other proposed cases of causal overdetermination, therefore cases of overdetermination can be explained away. Of course, one may not be persuaded by Bunzl. One may find this argument pedantic and insist that C1 and C2 together do not cause a distinct effect.

McDermott shows that Bunzl events don't always solve the problem of causal overdetermination. Let us change the previous case slightly; two electrons, C1 and C2 fire simultaneously and stimulate neuron N at the same time. Furthermore, neuron N responds in precisely the same way here as what would have happened if only C1 or C2 have fired. Therefore the stimulation the neuron receives from (C1 & C2), C1 alone, and C2 alone are not distinct. We therefore cannot explain away this type of causal overdetermination in the way Bunzl wants to.

Lewis himself has an odd way of dealing with cases of causal overdetermination. He argues that in these tougher cases, it is not clear what our intuitions say. Some would want to say that only Sally throwing the rock was a cause in the window breaking, others may argue that both Sally and Billy throwing the rock were causes. Some may even say that neither were causes. In these hard cases, the theory of causation can say whatever it wants. It gets to decide what counts as causes and what does not. "Such cases can be left as spoils to the victor" (Lewis, 1986b, p.194).

¹²⁶ Bunzl, 1979, p.134-50.

¹²⁷ Bunzl, 1979, p.143.

Lewis' response is decent, but it leaves certain puzzles. Many philosophers would intuitively believe that both C1 and C2 were causes of the stimulation of N. However, they cannot explain this intuition right now with the counterfactual theory of causation. I will make an attempt in amending this. Additionally, it would be nice if this solution gives us an alternative solution to cases of pre-emption. I will hence propose a model of redundant causation – a concept that involves both pre-emption and overdetermination.

A good causal analysis of the C1 and C2 story should not just include one of the electrons, but both, showing the total causal story. Part of adequately describing this example is showing that either C1 or C2 would have caused the same stimulation by themselves. Furthermore, we have to show that if none of the causes involved in causing the stimulation of N occurred, (that is, C1 and C2), N would not have occurred. This much should be obvious. These are all conditions that should be involved in a model of causal overdetermination.

Let '*c*' be the name of an event *c*, '*e*' be the name of an event *e*, and finally, let '*X*' be the combination of two or more non-*e* events. I therefore propose the following model of redundant causation:

Effect *e* was the result of redundant causation iff there was a combination of events *X* such that

1. The events in *X* and *e* occurred,
2. If the events in *X* had not occurred, *e* would not have occurred, and
3. If any one or more of the events in *X* had occurred, *e* would still have occurred.

With this amendment, McDermott's case, as well as all other examples of overdetermination can now be accounted for. We will now add this to our counterfactual theory of causation:

c causes *e* iff either

- 1) Both c and e occurred, and
 - 2) If c had not occurred, e would not have occurred either,
- or
- 3) c was part of X , and
 - 4) e was the result of redundant causation due to the events involved in X .

There are other counterexamples involving the problem of transitivity. There is a man in an electric chair and two men with switches in front of them. When the switches are placed in the same position, the man in the chair receives a shock. The first man does not want the man in the chair to receive a shock, but he nonetheless has to make a move and so flicks the switch to the right. The second man actually wants to electrocute the man in the chair, so he also flicks his switch to the right and the delivers a shock.¹²⁸ The first man caused the second man to flick the switch to the right, and the second man caused the man in the chair to get electrocuted, yet it is counter-intuitive to say that the first man – who did not want any torture to happen, caused the man in the chair to get electrocuted.

The way this counterexample works is that it takes advantage of the fact that Lewis allows for causal chains. Let ' f ' be an event f , ' g ' be an event g , and ' h ' be an event h . If f causes g and g causes h , then g causes h . We are hence given cases where f causes g , g causes h , yet we are not prepared to say that f causes h . There are ways to fix this problem using counterfactuals. Did the first man cause the man in the chair to get electrocuted? Putting everything aside, if we use the counterfactual test, we would say no. Regardless of what the first man does, the second man would've moved his switch in a way that electrocutes the man in the chair. Therefore, anything the first man does is causally irrelevant. We can simply add a condition to the rules involving causal chains to prevent this from happening.

If f causes g and g causes h , then f causes h unless it is proven that f cannot cause h once we remove the transitivity of causal chains.¹²⁹

¹²⁸ McDermott, 1995, p.532.

¹²⁹ Lewis for example argues for this in 1973, p.563.

Many of these counterexamples to the counterfactual theory of causation can be solved with counterfactual tools. Therefore they should not pose serious threats to the counterfactual theory of causation.

Does this theory have any issues with abstract objects being part of the causal order? I see no reason why there it would be. Abstract objects can form part of the antecedent in counterfactual stories, much like concrete objects. If Harry Potter did not go to Hogwarts, Suzy would not have bought her wizard hat. The most popular causal theory today therefore can accommodate efficacious abstract objects.

Most of the causal theories covered allow for at least some abstract objects to be part of the causal order. The only exception to the rule here is the energy transfer view, which is not a widely accepted causal theory. Even then, we have seen that it is at least possible for AOCO and Fair's view to coexist. I therefore conclude that AOCO does not clash with our theoretic conceptions of causation.

3.2. Azzouni's Objections

There have not been many publications on my topic. The two publications I could find are from Benjamin Callard and Jody Azzouni, one arguing for AOCO (roughly) and another objecting to it. Callard has argued a very closely related position to mine. He states that there is no reason *a priori* that abstract objects cannot be part of the causal order.

“Conceptually speaking, there is, on reflection, simply no problem at all with the idea of something being affected by an entity radically dissimilar from it. In particular, there is no problem with the idea of abstract objects effecting changes in us; there is no conceptual difficulty with the idea that they impart energy to our brains, and that they do this ‘at a distance’, *i.e.*, without the benefit of contiguity relations” (Callard, 2007, p.354-5).

He focuses on the possibility of causation from a distance and argues that though one can argue against it through empirical means, there is no reason *a priori* that this is not possible.¹³⁰ Azzouni responded in a paper that challenges Callard to show how such causation could happen. For Azzouni, the mere logical possibility is not enough. It is not acceptable to try and explain phenomena with answers that create more metaphysical problems.¹³¹

Azzouni argues that it is mysterious how the same abstract object can have different spatiotemporal relations to different people. He asks us to imagine that he has a perfect duplicate. Azzouni thinks of a number and after some time, in a different place, his duplicate thinks of the number as well. The same number therefore shares different relational properties between Azzouni and his duplicate. However, these differences are purely spatiotemporal. This is problematic, because numbers do not have any spatiotemporal properties, and hence, there should not be any differences in spatiotemporal relations between different people and the number. Yet, this is the sort of causal interaction that AOCO leads to (though neither Callard nor Azzouni use my terminology).¹³²

This objection in its current form is easy to dismiss. The problem with this argument is that there is a difference between the properties of an object and the properties of the thoughts regarding that object. Suppose a group of doctors were discussing their brain-dead patients. These doctors hence think about people that are incapable of thinking. By Azzouni's logic, this story would be full of contradictions, but of course it is not! The reason is simple - there is a difference between the properties of the subject and the properties of the *thoughts about* subject. Bringing this back to Azzouni's original counterexample, just because numbers do not have any spatiotemporal properties, that does not mean that the thoughts about numbers cannot.

Luckily for Azzouni, his counterexample can be adjusted to deal with this type of attack. Two renowned writers create scripts based on Sherlock Holmes. The first writer writes a script for the Sherlock Holmes movies, played by Robert Downey Jr.

¹³⁰ Callard, 2007, p.347-356.

¹³¹ Azzouni, 2008, p.401.

¹³² Azzouni, 2008, p.398.

In another place and another time, the second writer creates a script for the television series – Sherlock. The work of fiction was therefore exemplified at different places and at different times. How can something without spatiotemporal properties such as fictional characters hold different spatiotemporal relations with different objects, when these relational properties should be the same?

To this I have two answers. First of all there is some general confusion that comes with saying that abstract objects exist outside of space of time. This can be true without meaning that there exists a separate realm or a dimension where Sherlock Holmes for example resides in. All it means is that we can find the Sherlock Holmes anywhere, at any time. Therefore the abstract object can be non-spatiotemporal and yet we can interact with it in different places and at different times.

In case the reader is not impressed by my last argument, here is another. There have historically been many causal connections which we acknowledge, yet we had no viable explanation for them at the time. An ancient astronomer may know that the Earth orbits the sun for example, yet have no idea how this occurs. Nonetheless, he was right; it just meant that there were further puzzles to answer afterwards. Likewise, this mystery that comes with how these entities cause should not be blown out of proportion. If abstract objects do in fact exist, I have provided reasons why they are part of the causal order. If one is not satisfied as to how this causal process happens, that simply means we have more work to do. That does not mean that abstract objects do not have any causal powers. Azzouni's challenge towards the AOCO can therefore be answered.

3.3. Arguments from Make-No-Difference

My main causal theory does not seem to clash with any of the main causal theories and Azzouni's objection can be dealt with. Unfortunately, there has not been much discussion about this question among philosophers, so we are limited when it comes to objections. Similar questions however have been made on whether an abstract object can *make a difference*. This is a broader question as something can make a difference without it doing so causally. Ultimately, I do not believe there is a good argument against AOCO because a) the arguments from MND cannot be applied to our discussion and b) whether or not an abstract object plays a causal role cannot be shown definitely. I will list the points made for the make-no-difference argument, change them to fit the context of our debate and judge whether or not we should take these threats seriously.

The following arguments are all from Mark Balaguer and objected by Alan Baker. The first what Baker calls the 'blinking out argument'. Assume that Platonic entities exist (Balaguer was talking about mathematical entities, specifically), now imagine if all these entities suddenly disappeared. The physical world would stay exactly the same. This suggests that abstract objects make no difference in our world.

“We can think of it this way: if all the objects in the mathematical realm suddenly *disappeared*, nothing would change in the psychical world; thus, if empirical science is true right now, then its nominalistic content would remain true, even if the mathematical realm disappeared; but this suggests that if there never existed any mathematical objects to begin with, the nominalistic content of empirical science could nonetheless be true” (Balaguer, 1998, p.132).

Baker argues that it is problematic to think of Platonic entities coming in and out of existence. These entities exist necessarily; it makes no sense to say that the number 5 came into existence. Therefore the premise of the blinking out argument is flawed and

cannot be adjusted to argue against the make-no-difference argument.¹³³ Baker's point is a strong one, but it is not one we can use in our debate. Baker and Balaguer are debating the properties of a certain brand of abstract objects – ones that cannot come in and out of existence. However, our debate covers entities such as fictional characters and tropes, which *can* come in and out of existence. Therefore, Baker's argument cannot be applied here. However, perhaps we now have a simple solution to the blinking out argument. If the strength of the bridge disappeared, the bridge would probably fall. It is hence easier to see how a contingent abstract object would influence the world causally. Can we now show that all abstract objects do have causal powers? Unfortunately not. Although we can make sense of how certain abstract entities are causal, this does not apply to all abstract objects. Abstract objects that exist necessarily are especially problematic because we have problems understanding the counterfactual involved. This will be explained more later.

The second argument is named the 'irrelevance argument'. Balaguer argues that abstract objects cannot make a difference by referencing their usage in science.¹³⁴ Baker reformulates Balaguer's argument thusly:

- 1) Scientists apparently 'know' that mathematical objects are acausal.
- 2) So science does not assign any causal powers to mathematical entities.
- 3) So science predicts that the behaviours of the physical world are not dependent upon the existence of mathematical entities.
- 4) This suggests that what science says about the physical world could be true even if mathematical entities did not exist.¹³⁵

In other words, abstract objects don't seem to be crucial to scientific explanations. Baker sees two problems with this argument, the first between point 2 and point 3, the next with the strength of conclusion 4. With the former, point 2 specifies a causal dependence, whereas point 3 says the physical aren't dependent of the mathematical in 'any way at all'. Baker shows that the physical can depend on the mathematical without any causal relations, for example the proposition 'the physical system S is 40

¹³³ Baker, 2003, p.249.

¹³⁴ Balaguer, 1998, p.133.

¹³⁵ Baker, 2003, p.250.

degrees Celsius' depends on the number 40 playing the 'Celsius relation to the physical system', yet there is no causal dependence between the two. With the latter, conclusion 4 only says that the physical world *could* be the same if the mathematical world was any different, but that is not enough. That does not entail that mathematical objects don't make a difference. It may be the case that they *do* make a difference, but the physical world may be the same through some other means if mathematical objects didn't exist.¹³⁶

The latter point is fair; the irrelevance argument does not prove that abstract objects do not make a difference. What it does is it leaves us with an agnostic conclusion – we cannot tell if abstract objects make a difference or not. What the MND theorists need is some sort of Ockham's razor argument to go with it. This seems like a plausible option, especially if we want to use this argument in the causal debate; it is metaphysically simpler if we don't have to show how abstract objects are involved in causal chains. The other point Baker raised however unveils a problematic issue. Balaguer is accused of arguing for MND based on abstract objects not being causally responsible and this may be a dangerous move, but let us pause for a moment. The fact that this argument requires abstract objects to be causally inept is alarming. Premise 2 assumes that abstract objects are acausal, but given that our aim is to prove or disprove the causal powers of abstract objects, we surely cannot just assume that they don't! We therefore cannot use the irrelevance argument in our causal debate because to accept the premise would be a circular argument. Arguments from MND simply don't transition well to our debate. There is however, a problem that is raised by Baker (2003) in the MND debate that may give AOCO trouble. I will explain this argument next, as well as a similar problem involving counterpossibles.

¹³⁶ Baker, 2003, p.250-1.

3.4 Counterfactuals with Necessary Abstract Entities

Let's say that we wanted to test if abstract objects that exist necessarily such as numbers have any causal powers. Once we try to apply Lewis' counterfactual test, we realise that this is simply impossible to prove using counterfactuals. There are two reasons for this, one including counterpossibles and the other involving counterlegals. I will go through both problems.

Let us begin with counterpossibles. A counterpossible is a conditional with a necessary antecedent. We can take numbers here for example and set up a counterpossible counterfactual:

If Pythagoras' theorem is true, the bridge will be sturdy.

Let us try and apply Lewis' counterfactual test. We will start by denying the antecedent. Unfortunately, this is as far as we can go. We cannot deny the antecedent because the antecedent is necessarily true. There are no possible worlds in which Pythagoras' theorem is false; it is true in all possible worlds. Therefore we cannot apply Lewis' test. Whether or not these objects are causal therefore remain a mystery.

Let us suppose, just hypothetically, that we can indeed apply the counterfactual test. For example, Pythagoras' theorem may not exist in all possible worlds, so there are worlds in which we can deny the antecedent. We would still be in a sticky situation because of the nature of counterlegals. A counterlegal is a counterfactual that involves a change of our natural laws. A weak counterlegal is one that involves a natural law changing at some point in time, whereas a strong counterlegal involves us imagining a natural law that was always different. When we are imagining worlds in which Pythagoras' theorem is false, we are hence dealing with a strong counterlegal. Baker argues that these strong counterlegal situations involving necessary abstract objects are neither true nor false because we simply do not know what denying the antecedent

would entail.¹³⁷ The only things we can know in strong counterlegals are analytic truths based on the antecedent. Of course, propositions involving concrete objects such as whether or not our bridges will fall do not follow by definition, given that Pythagoras' theorem is false. Therefore we cannot say that these abstract objects are causal or causal.¹³⁸

This is neither an argument in favour of or against AOCO. It simply states that we should be agnostic about when it comes to certain abstract entities. There are still plenty of contingent abstract objects that we can prove to be causal, so AOCO remains intact.

3.5. Nominalist and Traditional Platonist Objections

AOCO will not appeal to nominalists for the simple reason that it is a realist thesis. The standard nominalist objection is that we simply do need to postulate abstract entities to explain our world. Ockham's razor states that we should not postulate abstract objects without needing to. We can explain our world without postulating abstract objects, therefore we shouldn't. I have three responses to this. First of all, notice that AOCO is not specific towards any particular type abstract object. Therefore, as long as we allow any abstract object within our ontology, AOCO should be considered. It is simply unrealistic for any nominalist to suggest that there are no abstract objects at all. Nominalists have tried to show how everything can be paraphrased while maintaining a functional metaphysics, but many of their attempts have been disappointing. Even Hatry Field could only paraphrase away a limited amount of mathematical statements. As long as we have to postulate some abstract objects, AOCO remains viable. Secondly, notice that this type of argument is not specific to ontological debates surrounding abstract objects. Common-sense anti-realists have argued similarly that tables and chairs do not add any explanatory value either, when we already have smaller metaphysical components like atoms and

¹³⁷ Baker, 2003, p.260-2.

¹³⁸ For more on counterlegals, see Kvat (1996).

molecules. This is therefore really just an argument for an ontology composed of simples. In the same sense that we can still postulate tables and chairs despite already committing to atoms and molecules, we can commit to abstract objects despite committing to concrete entities. Finally, if it pleases the nominalists, view AOCO as a discussion that should occur *before* ontological debates. We can view AOCO as a methodological if-then-ist thesis. That is, *if* abstract objects do exist, I suggest that they should have causal powers. Whether or not abstract objects exist is another matter but before we enter this debate, we need to define what properties abstract objects have. Any debates over the existence of abstract entities should involve them being causal; that is all AOCO needs to say. Nominalism hence should not directly clash with AOCO.

The second objection is one that traditional Platonists and nominalists alike can make. They can argue that having abstract objects within the causal order is simply more complicated than having them being causally inert. We should not overcomplicate our metaphysics without good reasons, thus we should reject AOCO. My response is that after extended investigations, it simply makes more sense for these entities to be causal. Word types for example seem to obviously be decided by the linguistic community, and it seems equally obvious that this change then affects the way we communicate. It would be odd for us to have a view of causally efficacious linguistic types and this is equally true with certain other abstract objects. Furthermore, we have discussed all the different benefits of AOCO. There are good reasons for us to accept AOCO over the traditional Platonist model.

Conclusion

The aim of this chapter was to investigate any potential metaphysical problems with AOCO. We first looked at whether different theories of causation would be incompatible with this position. Although some theories cannot accompany abstract objects functioning within the causal order, the most prominent theory – the counterfactual theory, has no problems with it. We then looked at Azzouni's direct

objection regarding non-spatiotemporal objects having different spatiotemporal relations and concluded that it could be answered. We tried to apply arguments from the make-no-difference debate, but these arguments either could not transition to ours or left us with an open question. Furthermore, we investigated questions concerning whether necessary abstract objects can be part of causation and decided that it was either not a problem, or simply an open question. Finally, traditional nominalist arguments do not pose a threat to AOCO. Until there are further metaphysical objections to the claim that abstract objects can be part of the causal order, it seems that we can answer any metaphysical queries that are raised.

Final Conclusion

Abstract objects have long been assumed to be acausal, but as I have shown there are good reasons to reject this assumption. Firstly, there are already abstract objects postulated by philosophers that are part of the causal order. My claim also offers a simple solution to Benacerraf's epistemic problem, which has been a thorn on the realists' side for a long time. Finally, we looked at metaphysical objections to my AOCO including direct ones from Azzouni and showed that they can be dealt with. I thereby conclude that abstract objects are in fact part of the causal order, and that nominalists should at least view these entities as the ones they should deal with, as they pose a stronger threat than traditional Platonic objects.

Bibliography

Alston, W.P. (1958). Ontological Commitments. *Philosophical Studies*, 9, 1-2, pp. 8-17.

Armstrong, D. M. (1973). *Belief, Truth and Knowledge*, Cambridge: Cambridge University.

Azzouni, J. (2008). A Cause for Concern: Standard Abstracta and Causation. *Philosophia Mathematica*, 16, 3, pp. 397-401.

Baker, A. (2003). Does the Existence of Mathematical Objects Make a Difference?. *Australasian Journal of Philosophy*, 81, 2, pp. 246-264.

Balaguer, M. (1995). A Platonist Epistemology. *Synthese*, 103, 3, pp. 303-25.

- (1998). *Platonism and Anti-Platonism in Mathematics*. Oxford: Oxford University Press.

Benacerraf, P. (1973). Mathematical Truth. *The Journal of Philosophy*, 70, 19, pp. 661-79.

Bigelow, J. (1988). *The Reality of Numbers*. Oxford: Clarendon Press.

Bonjour, L. (1980). Externalist Theories of Empirical Knowledge. *Midwest Studies in Philosophy*, 5, 1, pp. 53-74.

Brendel, E., & Jäger, C. (2004). Contextualist Approaches to Epistemology: Problems and Prospects. *Erkenntnis*, 61, 2-3, pp. 143-72.

Brock, S. (2010). The Creationist Fiction: The Case against Creationism about Fictional Characters. *Philosophical Review*, 119, 3, pp. 337-64.

Bunzl, M. (1979). Causal Overdetermination. *The Journal of Philosophy*, 76, 3, pp. 134-50.

Burgess, J. (1983). Why I am not a Nominalist. *Notre Dame Journal of Formal Logic*, 24, 1, pp. 93-105.

Burgess, J., & Rosen, G. (1997). *A Subject with No Object*. Oxford: Clarendon.

Callard, B. (2007). The Conceivability of Platonism. *Philosophia Mathematica*, 15, 3, pp. 347-56.

Cheyne, C. (2001). *Knowledge, Cause, and Abstract Objects: Causal Objections to Platonism*. Dordrecht: Kluwer.

Chihara, C. S. (1982). A Godelian Thesis Regarding Mathematical Objects: Do They Exist? and Can We Perceive Them?. *The Philosophical Review*, 91, 2, pp. 211-27.

Cohen, S. (1998). Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery. *Australasian Journal of Philosophy*, 76, 2, pp. 289–306.

- (1999). Contextualism, Skepticism, and the Structure of Reasons, *Philosophical Perspectives*, 13, pp. 57–89.

Conee, E., & Feldman, R. (1985). Evidentialism. *Philosophical Studies*, 48, 1, pp. 15-34.

Cresswell, M. (2010). Abstract Entities in the Causal Order. *Theoria*, 76, pp. 249-65.

Davidson, D. (1967). Causal Relations. *The Journal of Philosophy*, 64, 21, pp. 691-703.

Deutsch, H. (1991). The Creation Problem. *Topoi*, 10, 2, pp. 209–25.

- Dummett, M. (1973). *Frege – Philosophy of Language*. London: Trinity.
- Dretske, F. (1970). Epistemic Operators. *The Journal of Philosophy*, 67, 24, pp. 1007–23.
- (1981). The Pragmatic Dimension of Knowledge. *Philosophical Studies*, 40, 3, pp. 363–78.
- Drivers, J., & Miller. A. (1999). Arithmetical Platonism: Reliability and Judgement-Dependence. *Philosophical Studies*, 95, 3, pp. 277-310.
- Fair, D. (1979). Causation and the Flow of Energy. *Erkenntnis*, 14, 3, pp. 219-50.
- Feldman, R. (1999). Contextualism and Skepticism. *Philosophical Perspectives*, 13, pp. 91–114.
- Field, H. (1980). *Science without Numbers*. Princeton: Princeton University Press.
- (1984). Frege's Conception of Numbers as Objects by Crispin Wright. *Canadian Journal of Philosophy*, 14, 4, pp. 637-62.
- Gettier, E. (1963). Is Justified True Belief Knowledge?. *Analysis*, 23, 6, pp. 121-3.
- Gödel, K. (1951). Russell's Mathematical Logic. In P.A. Schlipp (Eds), *The Library of Living Philosophers – The Philosophy of Bertrand Russell* (pp. 125-51). New York: Tudor.
- Goldman, A.H. (1988) *Empirical Knowledge*. Berkeley: University of California Press.
- Hale, B. (1987). *Abstract Objects*. Oxford: Blackwell.
- Hume, D. (1975). *Enquiry concerning Human Understanding* (3rd ed.). Oxford: Clarendon. (Original work published 1748).

- (1978). *A Treatise of Human Nature* (2nd ed.). Oxford: Clarendon. (Original work published 1739).

Jackson, F. (1980). Ontological Commitment and Paraphrase. *Philosophy*, 55, 213, pp. 303-15

Kvart, I. (1996). *A Theory of Counterfactuals*. Indianapolis: Hackett

Lewis, D. (1973). Causation. *Journal of Philosophy*, 70, 17, pp. 556-67.

- (1986a). *On the Plurality of Worlds*. Oxford: Blackwell.
- (1986b). *Philosophical Papers: Volume II*. Oxford: Oxford University Press.
- (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74, 4, pp. 549–67.

Liggins, D. (2006) Is There a Good Epistemological Argument against Platonism?. *Analysis*, 66, 2, pp.135–41.

Linsky, B. & Zalta, E. (1995). Naturalized Platonism versus Platonized Naturalism. *Journal of Philosophy*, 92, 10, pp. 525-55.

Mackie, J.L. (1965). Causes and Conditions. *American Philosophical Quarterly*, 2, 4, pp. 245-264.

Maddy, P. (1981). Sets and Numbers. *Noûs*, 15, 4, pp. 495-511.

- (1992). *Realism in Mathematics*. Oxford: Clarendon Press.

McDermott, M. (1995). Redundant Causation. *The British Journal for the Philosophy of Science*, 46, 4, pp. 523-44

Noonan, H. W. (1978). Count Nouns and Mass Nouns. *Analysis*, 38, 4, pp. 167-72.

Plato (2005). Parmenides. In M. Cohen, P. Curd and C.D.C Reeve (Eds), *Readings in Ancient Greek: From Thales to Aristotle – Third Edition* (pp. 604-13). Cambridge, Indianapolis: Hackett Publishing Company. (Original work published ND).

Quine, W.V. (1948). On What There Is. *The Review of Metaphysics*, (pp. 21-38), Cambridge, MA: Harvard University.

- (1960). *Word and Object*. Cambridge, MA: The M.I.T. Press.

Rosen, G. (1990). Modal Fictionalism. *Mind Association*, 99, 395, pp. 327-54.

- Abstract Objects. *The Stanford Encyclopedia of Philosophy*. Retrieved 18 March, 2012 from <http://plato.stanford.edu/archives/spr2012/entries/abstract-objects>.

Stalnaker, R. (1968). A Theory of Conditionals. *American Philosophical Quarterly*, 5, 2, pp. 98-112.

Steiner, M. (1973). Platonism and the Causal Theory of Knowledge. *The Journal of Philosophy*, 70, 3, pp. 57-66.

Van Inwagen, P (1977). Creatures of Fiction. *American Philosophical Quarterly*, 14, 4, pp. 299-308.

- (1983). Fiction and Metaphysics. *Philosophy and Literature*, 7, 1, pp. 67-77.
- (1998). Meta-Ontology. *Erkenntnis*, 48, 2-3, pp. 233-50.

Wetzel, L. (2009). *Types and Tokens: On Abstract Objects*. London: MIT.

Williams, D. C. (1953a). The Elements of Being: I. *The Review of Metaphysics*, 7, 1, pp. 3-18.

- (1953b). The Elements of Being: II. *The Review of Metaphysics*, 7, 2, pp. 171-92.

Wright, C. (1983) *Frege's Conception of Numbers as Objects*. Aberdeen: Aberdeen University Press.

Yablo, S. (2001). Go Figure: A Path Through Fictionalism. *Midwest Studies in Philosophy*, 25, 1, pp. 72-102.